



Informatica® PowerExchange  
10.0

# CDC Guide for Linux, UNIX, and Windows

© Copyright Informatica LLC 1993, 2022

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](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.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://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://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: 2022-06-13

# Table of Contents

<b>Preface .....</b>	<b>11</b>
Informatica Resources. ....	11
Informatica Network. ....	11
Informatica Knowledge Base. ....	12
Informatica Documentation. ....	12
Informatica Product Availability Matrixes. ....	12
Informatica Marketplace. ....	12
Informatica Global Customer Support. ....	12
 <b>Part I: PowerExchange CDC Introduction.....</b>	 <b>13</b>
 <b>Chapter 1: Change Data Capture Introduction.....</b>	 <b>14</b>
PowerExchange CDC Overview. ....	14
Change Data Capture. ....	15
Change Data Extraction and Apply. ....	15
PowerExchange CDC Data Sources. ....	16
DB2 for Linux, UNIX, and Windows Data Sources. ....	16
Microsoft SQL Server Data Sources. ....	17
Oracle Data Sources. ....	17
i5/OS and z/OS Data Sources with Offload Processing. ....	18
PowerExchange CDC Components. ....	19
PowerExchange Listener. ....	19
PowerExchange Logger for Linux, UNIX, and Windows. ....	19
PowerExchange Navigator. ....	20
PowerExchange CDC Architecture. ....	20
PowerExchange Integration with PowerCenter. ....	22
Environment Variable Incompatibilities Between PowerExchange and PowerCenter. ....	23
Summary of CDC Implementation Tasks. ....	24
 <b>Part II: PowerExchange CDC Components.....</b>	 <b>25</b>
 <b>Chapter 2: PowerExchange Listener.....</b>	 <b>26</b>
PowerExchange Listener Overview. ....	26
Customizing the dbmover Configuration File for CDC. ....	26
CAPI_CONNECTION Statements. ....	27
CAPI_SRC_DFLT Statement. ....	28
CAPT_PATH Statement. ....	30
CAPT_XTRA Statement. ....	30
Starting the PowerExchange Listener. ....	31
Starting the PowerExchange Listener on Linux or UNIX. ....	31

Starting the PowerExchange Listener on Windows. . . . .	31
Stopping the PowerExchange Listener. . . . .	33
Displaying Active PowerExchange Listener Tasks. . . . .	33

## **Chapter 3: PowerExchange Logger for Linux, UNIX, and Windows. . . . . 35**

PowerExchange Logger Overview. . . . .	35
PowerExchange Logger Tasks. . . . .	37
PowerExchange Logger Files. . . . .	37
CDCT File. . . . .	37
PowerExchange Logger Log Files. . . . .	38
Lock Files. . . . .	39
Message Log Files. . . . .	40
File Switches. . . . .	40
PowerExchange Logger Operational Modes. . . . .	41
Continuous Mode. . . . .	41
Batch Mode. . . . .	42
PowerExchange Logger Usage Considerations. . . . .	42
Logging Data from Remote z/OS or i5/OS Sources. . . . .	43
Configuring the PowerExchange Logger. . . . .	43
Enabling a Capture Registration for PowerExchange Logger Use. . . . .	43
Customizing the PowerExchange Logger Configuration File. . . . .	44
Customizing the dbmover Configuration File for the PowerExchange Logger. . . . .	56
Using PowerExchange Logger Group Definitions. . . . .	60
Starting the PowerExchange Logger. . . . .	63
PWXCL Syntax and Parameters. . . . .	64
How the PowerExchange Logger Determines the Start Point for a Cold Start. . . . .	67
Cold Starting the PowerExchange Logger . . . . .	67
Managing the PowerExchange Logger. . . . .	68
Controlling and Stopping PowerExchange Logger Processing. . . . .	68
Monitoring the PowerExchange Logger. . . . .	69
Determining If the PowerExchange Logger Captured Changes. . . . .	71
PowerExchange Logger Verbose Messages. . . . .	72
PWXUCDCT Commands for Maintaining the PowerExchange Logger CDCT and Log Files. . . . .	72
Backing Up PowerExchange Logger Files. . . . .	74
Re-creating the CDCT File After a Failure. . . . .	75

## **Part III: PowerExchange CDC Data Sources. . . . . 76**

### **Chapter 4: DB2 for Linux, UNIX, and Windows CDC. . . . . 77**

DB2 for Linux, UNIX, and Windows CDC Overview. . . . .	77
Planning for DB2 CDC. . . . .	78
Prerequisites. . . . .	78
Required User Authority. . . . .	78

DB2 Datatypes Supported for CDC. . . . .	79
DB2 CDC Considerations. . . . .	80
Configuring DB2 for CDC. . . . .	81
Configuring PowerExchange for DB2 CDC. . . . .	82
Configuring PowerExchange CDC without the PowerExchange Logger. . . . .	82
Configuring PowerExchange CDC with the PowerExchange Logger. . . . .	82
Creating the Capture Catalog Table. . . . .	83
Initializing the Capture Catalog Table. . . . .	83
Customizing the dbmover Configuration File for DB2 CDC. . . . .	84
Using a DB2 Data Map. . . . .	89
Task Flow for DB2 Data Map Use. . . . .	89
Managing DB2 CDC. . . . .	90
Stopping DB2 CDC. . . . .	90
Changing a DB2 Source Table Definition. . . . .	90
Reconfiguring a Partitioned Database or Database Partition Group. . . . .	91
DB2 for Linux, UNIX, and Windows CDC Troubleshooting. . . . .	92
Workaround for SQL1224 Error on AIX. . . . .	93
 <b>Chapter 5: Microsoft SQL Server CDC. . . . .</b>	 <b>94</b>
Microsoft SQL Server CDC Overview. . . . .	94
Planning for SQL Server CDC. . . . .	95
SQL Server CDC Prerequisites. . . . .	95
Required User Authority for SQL Server CDC. . . . .	95
SQL Server Datatypes Supported for CDC. . . . .	96
SQL Server CDC Operational Considerations. . . . .	97
Extracting Data for Multiple Publication Databases. . . . .	98
Configuring SQL Server for CDC. . . . .	99
Configuring PowerExchange for SQL Server CDC. . . . .	99
Configuring PowerExchange CDC without the PowerExchange Logger. . . . .	100
Configuring PowerExchange CDC with the PowerExchange Logger. . . . .	100
Customizing the dbmover Configuration File for SQL Server CDC. . . . .	101
Managing SQL Server CDC. . . . .	106
Disabling Publication of Change Data for a SQL Server Source. . . . .	106
Changing a SQL Server Source Table Definition. . . . .	106
Changing the MULTIPUB Parameter Setting After Running Extractions. . . . .	107
 <b>Chapter 6: Oracle CDC with LogMiner. . . . .</b>	 <b>109</b>
Oracle CDC with LogMiner Overview. . . . .	109
Planning for Oracle CDC with LogMiner. . . . .	110
Oracle CDC with LogMiner Implementation Considerations. . . . .	110
Oracle CDC with LogMiner Restrictions. . . . .	112
Oracle Datatypes Supported for CDC. . . . .	113
Performance Considerations. . . . .	114

Oracle Configuration for CDC with LogMiner. . . . .	114
Configuration Readme File and Script Files. . . . .	114
Configuring Oracle for CDC. . . . .	115
Configuration in an Oracle RAC Environment. . . . .	119
PowerExchange Configuration for CDC with LogMiner. . . . .	120
Configuring Oracle CDC with LogMiner - Without the PowerExchange Logger. . . . .	120
Configuring Oracle CDC with LogMiner - With the PowerExchange Logger. . . . .	121
Customizing the dbmover Configuration File for Oracle CDC. . . . .	122
Managing Oracle CDC with LogMiner. . . . .	137
Stopping Oracle CDC with LogMiner Processing for Selected Tables. . . . .	137
Stopping Oracle CDC with LogMiner Processing Temporarily. . . . .	138
Changing a Source Table Definition Used in Oracle LogMiner CDC. . . . .	138
 <b>Chapter 7: Express CDC for Oracle. . . . .</b>	 <b>139</b>
Express CDC for Oracle Overview. . . . .	139
PowerExchange Express CDC for Oracle Benefits. . . . .	140
PowerExchange Express CDC for Oracle Architecture. . . . .	141
PowerExchange Express CDC Configuration Overview. . . . .	145
Gathering Information About the CDC Environment. . . . .	145
Planning Considerations. . . . .	147
PowerExchange Express CDC for Oracle Restrictions. . . . .	147
Oracle Datatypes Supported for Express CDC. . . . .	148
Operational Considerations. . . . .	150
Performance Considerations. . . . .	153
RAC Considerations. . . . .	153
ASM Considerations. . . . .	154
Oracle Data Guard Physical Standby Databases as Sources. . . . .	155
Oracle Multitenant Pluggable Databases as Sources. . . . .	157
Summary Task Flow. . . . .	158
Configuring Oracle for Express CDC. . . . .	159
Specifying an Archive Log Destination. . . . .	159
Enabling ARCHIVELOG Mode. . . . .	159
Creating an Oracle User and Granting User Privileges. . . . .	160
Creating an ASM User (ASM only). . . . .	161
Enabling Oracle Minimal Global Supplemental Logging. . . . .	162
Configuring PowerExchange for Express CDC. . . . .	162
Configuring the dbmover Configuration File. . . . .	164
Customizing the PowerExchange Express CDC for Oracle Configuration File. . . . .	173
Managing PowerExchange Express CDC for Oracle. . . . .	184
Monitoring PowerExchange Express CDC for Oracle. . . . .	184
Adding a Capture Registration. . . . .	185
Stopping CDC Processing for a Table. . . . .	186
Stopping CDC Processing Temporarily. . . . .	186

Changing the Structure of an Oracle Source Table . . . . .	186
<b>Chapter 8: Remote Logging of Data. . . . .</b>	<b>188</b>
Remote Logging Overview. . . . .	188
Remote Logging of Data from i5/OS or z/OS Sources. . . . .	189
Remote Logging of Data from Linux, UNIX, or Windows Sources. . . . .	190
Requirements for Capture Registrations. . . . .	193
Security Settings for Data from z/OS Sources. . . . .	193
Configuration Tasks for Remote Logging. . . . .	194
Customizing the PowerExchange Logger Configuration File for Remote Logging. . . . .	194
Customizing the dbmover Configuration File on the System to Which Data Is Logged. . . . .	198
Customizing the dbmover Configuration File on the PowerCenter Integration Service System. . . . .	199
Configuring Capture Registrations for the PowerExchange Logger. . . . .	199
Configuring PowerCenter Connection Attributes for Extracting Data from the Log Files. . . . .	199
Example of Remote Logging from a z/OS Data Source. . . . .	200
Example of Remote Logging from a DB2 for i5/OS Data Source. . . . .	202
Example of Remote Logging from an Oracle CDC with LogMiner Data Source. . . . .	204
<b>Part IV: Change Data Extraction. . . . .</b>	<b>207</b>
<b>Chapter 9: Introduction to Change Data Extraction. . . . .</b>	<b>208</b>
Change Data Extraction Overview. . . . .	208
Extraction Modes. . . . .	209
PowerExchange-Generated Columns in Extraction Maps. . . . .	210
Uses of BI and CI Fields in Extraction Maps. . . . .	214
Restart Tokens and the Restart Token File. . . . .	215
Multiple-Source Processing in CDC Sessions. . . . .	216
Commit Processing with PWXPC. . . . .	218
Tuning Options. . . . .	219
<b>Chapter 10: Extracting Change Data. . . . .</b>	<b>220</b>
Overview of Extracting Change Data. . . . .	220
Security Considerations for Extracting z/OS Data. . . . .	221
Task Flow for Extracting Change Data. . . . .	222
Testing an Extraction Map. . . . .	222
Configuring PowerCenter CDC Sessions. . . . .	224
Changing Default Values for Session and Connection Attributes. . . . .	224
Configuring Application Connection Attributes. . . . .	225
Examples of Controlling Commit Processing. . . . .	233
Recovery and Restart Processing for CDC Sessions. . . . .	235
PowerCenter Recovery Tables for Relational Targets. . . . .	236
PowerCenter Recovery Files for Nonrelational Targets. . . . .	237
Application Names. . . . .	238



Restart Processing for CDC Sessions by Start Type. . . . .	238
Creating Restart Tokens for Extractions. . . . .	241
Displaying Restart Tokens. . . . .	241
Configuring the Restart Token File. . . . .	242
Restart Token File Statements. . . . .	243
Example Restart Token File. . . . .	245
<b>Chapter 11: Managing Change Data Extractions. . . . .</b>	<b>246</b>
Starting PowerCenter CDC Sessions. . . . .	246
Cold Start Processing. . . . .	247
Warm Start Processing. . . . .	247
Recovery Processing. . . . .	248
Stopping PowerCenter CDC Sessions. . . . .	248
Stop Command Processing. . . . .	249
Terminating Conditions. . . . .	249
Changing PowerCenter CDC Sessions. . . . .	250
Examples of Adding Sources and Creating Restart Tokens. . . . .	250
Recovering PowerCenter CDC Sessions. . . . .	252
Example of Session Recovery. . . . .	253
<b>Part V: Monitoring and Tuning. . . . .</b>	<b>255</b>
<b>Chapter 12: Monitoring CDC Sessions. . . . .</b>	<b>256</b>
Monitoring Overview. . . . .	256
Monitoring CDC Sessions in PowerExchange. . . . .	256
Read Progress Messages. . . . .	257
Extraction Statistics Messages. . . . .	257
Multithreaded Processing Statistics. . . . .	258
PowerExchange Listener DISPLAY ACTIVE or LISTTASK Command. . . . .	259
PowerExchange Listener DISPLAYSTATS Command. . . . .	259
PowerExchange Logger for Linux, UNIX, and Windows Monitoring Statistics . . . . .	261
Monitoring CDC Sessions in PowerCenter. . . . .	263
Session Log Messages. . . . .	264
Performance Details in Workflow Monitor. . . . .	264
Viewing Performance Details in Workflow Monitor. . . . .	266
<b>Chapter 13: Tuning CDC Sessions. . . . .</b>	<b>268</b>
Tuning Overview. . . . .	268
PowerExchange DBMOVER Statements for Tuning CDC Sessions. . . . .	269
PowerCenter Connection Attributes and Session Properties. . . . .	272
PowerCenter Connection Attributes for Tuning CDC Sessions . . . . .	272
PowerCenter Session Properties for Tuning Buffer Memory. . . . .	275
CDC Offload Processing. . . . .	275

Rules and Guidelines for CDC Offload Processing. . . . .	276
Enabling Offload Processing for CDC Sessions. . . . .	276
Example of CDC Offload Processing with an Oracle Source. . . . .	277
Multithreaded Processing. . . . .	277
Rules and Guidelines for Multithreaded Processing. . . . .	278
Enabling Multithreaded Processing for CDC Sessions. . . . .	278
<b>Appendix A: DTL__CAPXTIMESTAMP Time Stamps.....</b>	<b>279</b>
Time Stamps That Are Reported in the DTL__CAPXTIMESTAMP Field by Data Source. . . . .	279
<b>Index. . . . .</b>	<b>281</b>

# Preface

The *PowerExchange CDC Guide for Linux, UNIX, and Windows* describes how to configure, implement, and manage PowerExchange change data capture (CDC) on Linux, UNIX, and Windows systems.

This guide covers the following PowerExchange data sources for CDC:

- DB2 for Linux, UNIX, and Windows
- Microsoft SQL Server
- Oracle, including PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle

**Note:** If you use the remote logging or offloading feature, some PowerExchange CDC processing for DB2 for i5/OS and z/OS data sources can also run on Linux, UNIX, or Windows.

Before implementing change data capture, verify that you have installed the required PowerExchange components.

## 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.

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.
- 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](mailto: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](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](mailto: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 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>.

# Part I: PowerExchange CDC Introduction

This part contains the following chapter:

- [Change Data Capture Introduction, 14](#)

# CHAPTER 1

## Change Data Capture Introduction

This chapter includes the following topics:

- [PowerExchange CDC Overview, 14](#)
- [PowerExchange CDC Data Sources, 16](#)
- [PowerExchange CDC Components, 19](#)
- [PowerExchange CDC Architecture, 20](#)
- [PowerExchange Integration with PowerCenter, 22](#)
- [Environment Variable Incompatibilities Between PowerExchange and PowerCenter, 23](#)
- [Summary of CDC Implementation Tasks, 24](#)

## PowerExchange CDC Overview

PowerExchange change data capture (CDC) works in conjunction with PowerCenter to capture changes to data in source tables and replicate those changes to target tables and files. This guide describes PowerExchange CDC for relational database sources on Linux, UNIX, or Windows operating systems.

These sources are:

- DB2 for Linux, UNIX, and Windows
- Microsoft SQL Server on Windows
- Oracle on Linux, UNIX, or Windows

After materializing target tables or files with PowerExchange bulk data movement, you can use PowerExchange CDC to synchronize the targets with their corresponding source tables. Synchronization is faster when you replicate only the change data rather than all of the data.

The change data replication process consists of following high-level steps:

1. Change data capture. PowerExchange captures change data for the source tables. PowerExchange can read change data directly from the RDBMS log files or database. Optionally, you can use the PowerExchange Logger for Linux, UNIX, and Windows to capture change data to its log files.
2. Change data extraction. PowerExchange, in conjunction with PowerCenter, extracts captured change data for movement to the target.
3. Change data apply. PowerExchange, in conjunction with PowerCenter, transforms and applies the extracted change data to target tables or files.

## Change Data Capture

PowerExchange can capture change data directly from DB2 database logs, Microsoft SQL Server distribution databases, or Oracle redo logs. If you use the offloading feature in combination with the PowerExchange Logger for Linux, UNIX, and Windows, a PowerExchange Logger process can log change data from data sources on an i5/OS or z/OS system.

If you do not retain database log files long enough for CDC processing to complete, use the PowerExchange Logger for Linux, UNIX, and Windows. The PowerExchange Logger writes change data to its log files. PowerExchange can then extract change data from the PowerExchange Logger log files rather than from the database log files.

For each source table, you must define a capture registration in the PowerExchange Navigator. The capture registration provides metadata for the columns that are selected for change capture.

PowerExchange captures changes that result from successful SQL INSERT, DELETE, and UPDATE operations. Depending on the statement type, PowerExchange captures the following data images:

- For INSERTS, PowerExchange captures after images only. An *after image* reflects a row just after an INSERT operation. PowerExchange passes these changes as INSERTs to PowerCenter.
- For DELETES, PowerExchange captures before images only. A *before image* reflects a row just prior to the last DELETE operation. PowerExchange passes these changes as DELETES to PowerCenter.
- For UPDATES, PowerExchange captures the following image types:
  - Both before and after images if you select an image type of "BA" in the CDC application connection attributes for PowerCenter. PowerExchange passes an UPDATE to PowerCenter as a DELETE of the before-image data followed by an INSERT of the after-image data.
  - After images if you select an image type of "AI" in the CDC application connection attributes. PowerExchange passes only the after-image data for an updated row, unless you also request before-image data. PowerExchange passes an UPDATE to PowerCenter as an UPDATE or INSERT.

## Change Data Extraction and Apply

PowerExchange works with PowerCenter to extract change data and write it to one or more target tables or files. The targets can be on the same system as the source or on a different system.

When you create a capture registration for a source table, the PowerExchange Navigator generates a corresponding extraction map and application name for the extraction. The extraction map describes the columns for which to extract change data. You can edit the extraction map to remove columns from extraction processing. Also, you can create alternative extraction maps, each for a subset of the columns that are registered for capture. For DB2 for Linux, UNIX, and Windows data sources only, you can create a data map if you have user-defined or multi-field columns for which you want to manipulate data before loading it to the target.

From PowerCenter, you run a CDC workflow and session that extracts and applies change data. To define a data source in PowerCenter, you can import the extraction map or import the table definition from the source database through PowerExchange. For DB2 only, you can import a DB2 data map instead of the extraction map. In most situations, Informatica recommends that you import the extraction map.

Also, you must define a mapping, session, and workflow in PowerCenter. Optionally, you can include transformations in the mapping to manipulate the change data. When you define a CDC session, you must specify a connection type. The connection type determines the extraction mode and access method that PowerExchange uses to extract data.

To extract change data directly from source DB2 or Oracle log files or SQL Server distribution database, you must use the real-time extraction mode. To extract change data from PowerExchange Logger log files, you can use either the batch extraction mode or continuous extraction mode.

The following table describes these extraction modes:

Extraction Mode	Description
Real-time extraction mode	Reads change data directly from the database log files in near real time, on an ongoing basis. When the PowerExchange Listener receives an extraction request, it pulls the change data from the log files and transmits the data to PowerCenter for extraction and apply processing. This mode provides the lowest latency for change data extraction but potentially the highest impact on system resources.
Batch extraction mode	Reads change data from PowerExchange Logger log files that are in a closed state when an extraction request is made. After processing the log files, the extraction request ends. This mode provides the highest latency for change data extraction but minimizes the impact on system resources.
Continuous extraction mode	Reads change data continuously from open and closed PowerExchange Logger log files in near real time. This mode also minimizes database log accesses and the log retention period that is required for CDC.

To initiate change data extraction and apply processing, run a CDC workflow and session from PowerCenter.

During extraction processing, PowerExchange extracts changes from the change stream in chronological order based on the unit of work (UOW) end time. PowerExchange passes only the successfully committed changes to PowerCenter for processing. PowerExchange does not pass ABORTed or UNDO changes. If you are capturing changes from DB2 database logs or Oracle redo logs, changes that were contiguous in the change stream might not be contiguous in the reconstructed UOW that PowerExchange passes to PowerCenter.

To properly resume extraction processing, PowerExchange maintains restart tokens for each source table. Restart tokens are used for all extraction modes. To generate current restart tokens, you can use the PowerExchange Navigator, the special override statement in the restart token file, or the DTLUAPPL utility.

## RELATED TOPICS:

- [“Introduction to Change Data Extraction” on page 208](#)

# PowerExchange CDC Data Sources

PowerExchange can capture change data from DB2 and Oracle data sources on Linux, UNIX, or Windows systems. PowerExchange can also capture change data from Microsoft SQL Server data sources on Windows.

In the PowerExchange Navigator, you must create a capture registration for each source table. The PowerExchange Navigator generates a corresponding extraction map and application name. You can import the extraction map into PowerCenter to define the source for extraction and apply processing.

If you use the PowerExchange Logger for Linux, UNIX, and Windows in combination with the offloading feature, you can also process change data from data sources on i5/OS or z/OS.

## DB2 for Linux, UNIX, and Windows Data Sources

PowerExchange captures change data from DB2 recovery log files for the database that contains your source tables. For CDC to work, archive logging must be active for the database. Also, you must create a



PowerExchange capture catalog table in the source database. The capture catalog table stores information about the source tables and columns, including DB2 log positioning information.

If you have a source table with user-defined fields or multi-field columns, you can create a data map to manipulate these fields with expressions. For example, you might want to create data map to manipulate packed data in a CHAR column. If you create a data map, you must still create a capture registration and merge the data map with the extraction map that is generated for the capture registration.

#### RELATED TOPICS:

- [“DB2 for Linux, UNIX, and Windows CDC” on page 77](#)

## Microsoft SQL Server Data Sources

PowerExchange CDC uses Microsoft SQL Server transactional replication technology to access data in SQL Server distribution databases. For CDC to work, you must enable SQL Server Replication on the system from which change data is captured. Also, verify that each source table in the distribution database has a primary key. If your database has a high volume of change activity, use a distributed server as the host of the distribution database. When the extraction process runs, the Microsoft SQL Server Agent must also be running.

#### RELATED TOPICS:

- [“Microsoft SQL Server CDC” on page 94](#)

## Oracle Data Sources

PowerExchange for Oracle provides two alternative methods of capturing change data from Oracle sources: *PowerExchange Express CDC for Oracle* and *PowerExchange Oracle CDC with LogMiner*. Both methods are delivered as part of the PowerExchange for Oracle CDC option but you can use only one of them in a PowerExchange instance with the same configuration files.

With both methods, Informatica strongly recommends that you use the PowerExchange Logger for Linux, UNIX, and Windows.

#### **PowerExchange Express CDC for Oracle**

PowerExchange Express CDC reads change data directly from Oracle active and archived redo logs. It is more efficient and faster than PowerExchange Oracle CDC with LogMiner in many environments. Also, it avoids LogMiner reinitialization issues.

You can run PowerExchange Express CDC on the database server or on another supported 64-bit Linux, UNIX or Windows machine. At initialization, PowerExchange Express CDC stores the Oracle data dictionary in memory.

You must run Oracle in ARCHIVELOG mode with minimal global supplemental logging enabled.

PowerExchange Express CDC supports RAC and non-RAC environments, ASM, and Oracle Data Guard logical and physical standby databases.

Although PowerExchange Express CDC supports many types of Oracle environments, it has the following restrictions:

- PowerExchange Express CDC does not capture changes for the following types of objects:
  - Tables defined with sorted hash cluster columns
  - Virtual columns

- LOB columns that are defined with the STORE AS SECUREFILE storage parameter
- Columns that use Oracle Advanced Security Transparent Data Encryption (TDE)
- Tablespace that use any type of encryption other than TDE
- PowerExchange Express CDC does not capture the following types of operations:
  - Direct-path INSERTs
  - INSERT AS SELECT, MERGE, and PL/SQL FORALL operations that use the APPEND hint or parallel mode
  - SQL\*Loader direct-path loads
  - CREATE TABLE...AS SELECT operations

#### **PowerExchange Oracle CDC with LogMiner**

PowerExchange Oracle CDC with LogMiner uses Oracle LogMiner continuous mining to read change data from the Oracle archive logs that reside at the archive destination to which they were originally written.

You must run Oracle in ARCHIVELOG mode with minimal global supplemental logging enabled. You must also periodically copy the Oracle online data dictionary to the archive log destination so that PowerExchange can determine restart points for change data extraction processing.

PowerExchange Oracle CDC with LogMiner supports RAC and non-RAC environments, ASM, and Oracle Data Guard logical standby databases.

PowerExchange Oracle CDC with LogMiner does not capture change data from the following types of objects and environments:

- Oracle Data Guard physical databases
- Oracle 11g or later tables that use the system partitioning or reference partitioning scheme and virtual columns
- Oracle 12c multitenant environments

**Important:** You cannot use both PowerExchange Express CDC for Oracle and PowerExchange Oracle CDC with LogMiner in the same PowerExchange instance with the same dbmover.cfg and pwxocl.cfg configuration files.

#### **RELATED TOPICS:**

- [“Oracle CDC with LogMiner” on page 109](#)
- [“Express CDC for Oracle” on page 139](#)

## **i5/OS and z/OS Data Sources with Offload Processing**

You can use CDC offload processing in combination with the PowerExchange Logger for Linux, UNIX, and Windows to log change data from data sources on systems other than the system where the PowerExchange Logger runs.

With offload processing, a PowerExchange Logger process on Linux, UNIX, and Windows can log change data from i5/OS and z/OS systems as well as from other Linux, UNIX, or Windows systems. For example, a PowerExchange Logger process can log change data from a DB2 instance on z/OS.

# PowerExchange CDC Components

Several PowerExchange components are involved in change data capture (CDC).

These components are:

- **PowerExchange Listener.** Required, unless PowerExchange and the PowerCenter Integration Service are installed on the same physical machine.
- **PowerExchange Logger for Linux, UNIX, and Windows.** Optional.
- **PowerExchange Navigator.** Required.

## PowerExchange Listener

The PowerExchange Listener manages capture registrations and extraction maps for all CDC data sources. It also manages data maps if you create any for DB2 for Linux, UNIX, and Windows tables. The PowerExchange Listener maintains this information in the following files:

- CCT file for capture registrations
- CAMAPS directory for extraction maps
- DATAMAPS directory for DB2 data maps

The PowerExchange Listener also handles PowerCenter extraction requests for both change data replication and bulk data movement.

When you create, edit, or delete capture registrations or extraction maps in the PowerExchange Navigator, the PowerExchange Navigator uses the location value in the registration group and extraction group to contact the PowerExchange Listener. This location corresponds to a NODE statement in the dbmover.cfg file. For example, when you open a registration group for a RDBMS instance, the PowerExchange Navigator communicates with the PowerExchange Listener to get all capture registrations defined for that instance.

A PowerExchange Listener is not required if PowerExchange and the PowerCenter Integration Service run on the same physical machine.

### RELATED TOPICS:

- [“PowerExchange Listener” on page 26](#)

## PowerExchange Logger for Linux, UNIX, and Windows

The PowerExchange Logger for Linux, UNIX, and Windows captures change data from DB2 database logs, Oracle redo logs, or a SQL Server distribution database and writes that data to PowerExchange Logger log files.

Use of the PowerExchange Logger is optional. To use the PowerExchange Logger, run one PowerExchange Logger process for each database type and instance. The PowerExchange Logger writes all successful UOWs in chronological order based on end time to its log files. This practice maintains transactional integrity. You can extract the change data from the PowerExchange Logger log files in either batch or continuous mode.

Benefits of the PowerExchange Logger include:

- Source database overhead is reduced because PowerExchange makes fewer accesses to the source log files or database to read change data. For Oracle, this overhead reduction can be significant. The PowerExchange Logger can use only one Oracle LogMiner session to read change data for all extractions that process an Oracle instance.
- You do not need to retain the source RDBMS log files longer than normal for CDC.

- PowerExchange does not need to reposition its point in the DB2 or Oracle logs from which to resume reading data. This feature can significantly reduce restart times.

**Tip:** Informatica strongly recommends that you use the PowerExchange Logger rather than real-time extraction mode for both PowerExchange Oracle CDC with LogMiner sources and PowerExchange Express CDC for Oracle sources. For Oracle CDC with LogMiner, this configuration enables PowerExchange to use one Oracle LogMiner session for all extractions that process an Oracle instance.

## RELATED TOPICS:

- [“PowerExchange Logger for Linux, UNIX, and Windows” on page 35](#)

## PowerExchange Navigator

The PowerExchange Navigator is the graphical user interface from which you define and manage capture registrations, extraction maps, and data maps.

You must define a capture registration for each source table. The corresponding extraction map is automatically generated. For DB2 sources, you can also define data maps if you need to perform column-level processing, such as adding user-defined columns and building expressions to populate them. You can import the extraction maps into PowerCenter so that they can be used for moving change data to the target.

**Note:** If the PowerExchange Navigator is not installed on the same machine as a Microsoft SQL Server data source, you must install the SQL Server client software on the PowerExchange Navigator machine. The client software is required because PowerExchange uses SQL Server services when creating capture registrations. For the same situation with DB2 and Oracle data sources, you do not need the RDBMS client software. Instead, from the PowerExchange Navigator, you can point to the PowerExchange Listener on the machine that contains the source DB2 database or Oracle instance.

For more information about the PowerExchange Navigator, see the *PowerExchange Navigator User Guide*.

## PowerExchange CDC Architecture

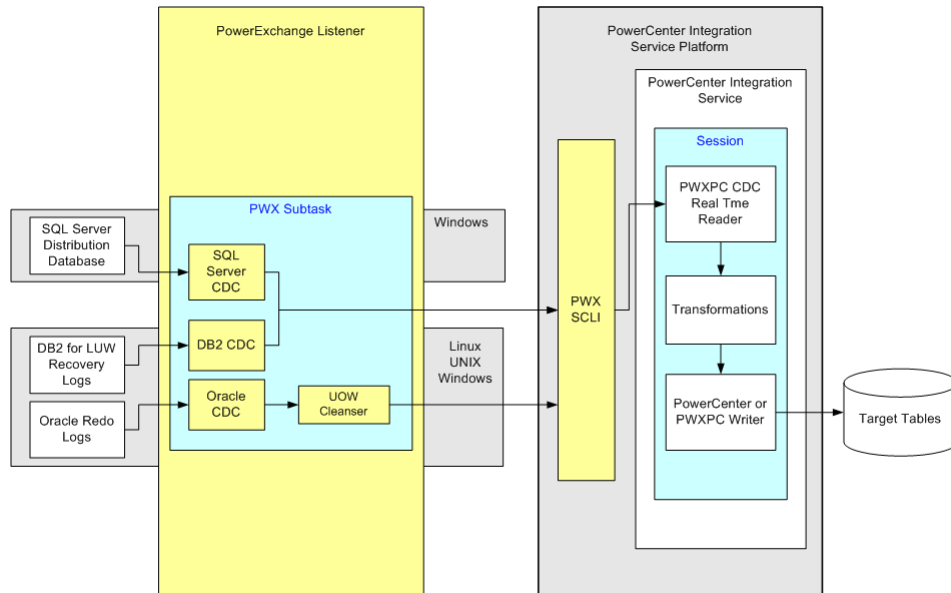
The PowerExchange CDC architecture is sufficiently flexible to handle many change data replication scenarios.

You can use PowerExchange in conjunction with PowerCenter to replicate change data from multiple sources of the same RDBMS type to multiple targets of different types in a single session.

The targets can be tables or files on the same system as the source or on other systems. The PowerCenter Integration Service can write data to tables in some RDBMSs as well as to flat files and XML files. If you installed PowerExchange or PowerExchange (PowerCenter Connect) products that provide connectivity to additional nonrelational or relational targets, you can also load data to those targets, for example, DB2 for z/OS tables, VSAM data sets, IMS segments, or WebSphere MQ.

You can run multiple instances of PowerExchange CDC components on a single system. For example, you might want to run a separate PowerExchange Logger for each source RDBMS to create separate sets of log files for each RDBMS type.

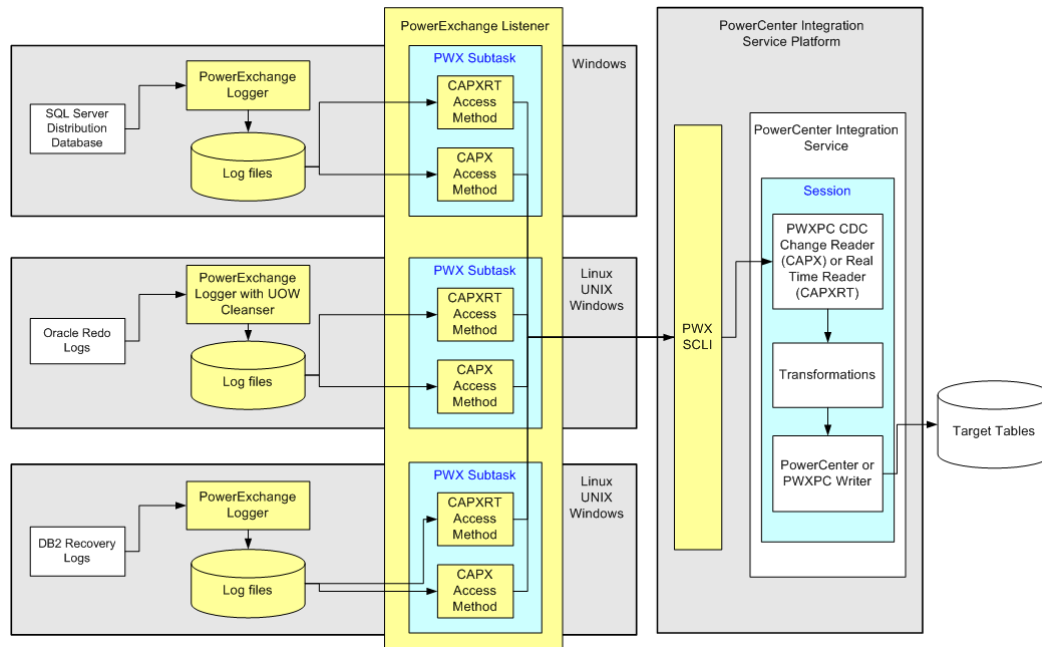
The following figure shows a simple CDC configuration that uses real-time extraction mode to access change data directly from the change stream without the PowerExchange Logger:



In this real-time configuration, PowerExchange CDC uses the CAPXRT access method to capture change data from a SQL Server distribution database, DB2 database logs, and Oracle redo logs. When an extraction request runs, PowerCenter connects to the PowerExchange Call Level Interface (SCLI) to contact the PowerExchange Listener. The change data is passed to the SCLI and then to the PWXPC CDC Real Time reader. In this manner, the PowerCenter extraction session pulls the change data that PowerExchange captured. After the PWXPC reader reads the change data, PowerCenter uses the mapping and workflow that you created to transform the data and load it to the target. With this configuration, you can replicate change data from multiple sources in the same database or instance to multiple target tables in a single extraction process.

**Note:** For PowerExchange Oracle CDC with LogMiner, the UOW Cleanser reconstructs UOWs from redo logs into complete and consecutive UOWs that are in chronological order by end time. For DB2 CDC and Express CDC for Oracle, PowerExchange incorporates the UOW Cleanser function into the consumer API (CAPI) for extracting changes from the data source.

The following figure shows a CDC configuration that uses the PowerExchange Logger in both batch extraction mode and continuous extraction mode:



In this configuration, the PowerExchange Logger captures change data from the change stream for SQL Server, Oracle, and DB2 tables and writes that data to its log files. After the data is in the PowerExchange log files, the source RDBMS log files can be deleted, if necessary. When an extraction session runs, PWXPC contacts the PowerExchange Listener. The PowerExchange Listener reads the PowerExchange Logger log files and calls the SCL on the PowerCenter Integration Service machine to transmit the change data to PowerCenter.

For some source tables, PWXPC extracts change data from the PowerExchange Logger log files in batch extraction mode with the CAPX access method. In this mode, the extraction session stops after it completes processing the log files. For other source tables, PWXPC extracts change data in continuous mode with the CAPXRT access method. In this mode, the extraction session extracts change data on an ongoing basis. In PowerCenter, you can create one source definition and one mapping that covers both extraction modes. However, batch and continuous extractions must run as separate sessions. For a batch extraction session, use a PWX CDC Change application connection. For a continuous extraction session, use a PWX CDC Real Time application connection. For example, you can run batch extractions to replicate change data to targets that need to be synchronized periodically, and run continuous extractions to replicate change data to targets that need to be synchronized in near real time. Batch and continuous extraction sessions can run concurrently.

## PowerExchange Integration with PowerCenter

PowerCenter works in conjunction with the PowerExchange Client for PowerCenter (PWXPC) to extract the change data that PowerExchange captured and apply it to one or more targets.

The primary function of PWXPC is to integrate PowerExchange with PowerCenter so that PowerCenter can access PowerExchange-controlled data and write it to various targets. With PWXPC, CDC sessions can extract change data directly from the change stream and from PowerExchange Logger for Linux, UNIX, and Windows log files.

PowerCenter provides transformation and data cleansing capabilities, which you can use in your CDC sessions.

For more information about PWXPC, see *PowerExchange Interfaces for PowerCenter*.

## Environment Variable Incompatibilities Between PowerExchange and PowerCenter

When PowerCenter and PowerExchange are installed on the same Linux, UNIX, or Windows machine, in certain cases, they have conflicting requirements for the PATH and LD\_LIBRARY\_PATH environment variables. To run correctly in these cases, PowerExchange and PowerCenter must run in separate environments.

This requirement applies when the PowerCenter Integration Service or PowerCenter Repository Service runs on the same machine as one of the following PowerExchange components:

- PowerExchange Listener
- PowerExchange Logger for Linux, UNIX, and Windows
- PowerExchange Navigator
- Any PowerExchange utility except the createdatamaps utility

The following table describes the restrictions that apply to the PATH and LD\_LIBRARY\_PATH variables in the PowerExchange and PowerCenter environments:

Environment	PATH	LD_LIBRARY_PATH
PowerExchange	\$INFA_HOME must not precede \$PWX_HOME. Otherwise, you cannot start the PowerExchange Listener or Logger from the command line.	LD_LIBRARY_PATH must not contain an entry for PowerCenter. This requirement ensures that PowerExchange utilities pick up their libraries from \$PWX_HOME only.
PowerCenter	The \$PWX_HOME entry must not precede the \$INFA_HOME entry.	The \$LD_LIBRARY_PATH variable definition must include both \$INFA_HOME and \$PWX_HOME, and \$INFA_HOME must precede \$PWX_HOME. For example:  <code>\$INFA_HOME/server/ bin:\$PWX_HOME:\$LD_LIBRARY_PATH</code>

To set the correct environment for PowerExchange or PowerCenter instances on the same machine, use one of the following strategies:

- Always start PowerExchange and PowerCenter using separate user accounts, and set the environment variables appropriately for each account.
- Run the pwxsettask.sh or pwxsettask.bat script each time you start a PowerExchange component.

# Summary of CDC Implementation Tasks

After you install PowerExchange, you can configure change data capture and extraction, materialize targets, and start extraction processing.

The following table identifies the tasks that you perform to implement change data capture and extraction processing for a Linux, UNIX, or Windows data source:

Step	Task	References
1	Configure parameters in the dbmover.cfg file for the PowerExchange Listener.	<a href="#">"Customizing the dbmover Configuration File for CDC" on page 26</a>
2	Start the PowerExchange Listener on the machine with the source database.	<a href="#">"Starting the PowerExchange Listener" on page 31</a>
3	Perform RDBMS-specific configuration tasks for CDC.	<ul style="list-style-type: none"><li>- <a href="#">Chapter 4, "DB2 for Linux, UNIX, and Windows CDC" on page 77</a></li><li>- <a href="#">Chapter 5, "Microsoft SQL Server CDC" on page 94</a></li><li>- <a href="#">Chapter 6, "Oracle CDC with LogMiner" on page 109</a></li><li>- <a href="#">Chapter 7, "Express CDC for Oracle" on page 139</a></li></ul>
4	(Optional) Configure the PowerExchange Logger.	<a href="#">"Configuring the PowerExchange Logger" on page 43</a>
5	(Optional) Start the PowerExchange Logger.	<a href="#">"Starting the PowerExchange Logger" on page 63</a>
6	From the PowerExchange Navigator, define and activate capture registrations and extraction maps for the data sources.	<i>PowerExchange Navigator User Guide</i>
7	For DB2 sources that have user-defined or multi-field columns that you want to manipulate, create DB2 data maps.	<i>PowerExchange Navigator User Guide</i>
8	Materialize the target from the source.	<i>PowerExchange Bulk Data Movement Guide</i>
9	Establish a start point for the extraction.	<a href="#">"Restart Tokens and the Restart Token File" on page 215</a>
10	From PowerCenter, configure mappings, workflows, connections, and sessions. Then run the workflow.	<ul style="list-style-type: none"><li>- <i>PowerExchange Interfaces for PowerCenter</i></li><li>- <i>PowerCenter Designer Guide</i></li><li>- <i>PowerCenter Workflow Basics Guide</i></li></ul>



# Part II: PowerExchange CDC Components

This part contains the following chapters:

- [PowerExchange Listener, 26](#)
- [PowerExchange Logger for Linux, UNIX, and Windows, 35](#)

## CHAPTER 2

# PowerExchange Listener

This chapter includes the following topics:

- [PowerExchange Listener Overview, 26](#)
- [Customizing the dbmover Configuration File for CDC, 26](#)
- [Starting the PowerExchange Listener, 31](#)
- [Stopping the PowerExchange Listener, 33](#)
- [Displaying Active PowerExchange Listener Tasks, 33](#)

## PowerExchange Listener Overview

In a change data capture (CDC) environment, a PowerExchange Listener can provide some or all of the following services:

- Store and manage capture registrations, extraction maps, and data maps for CDC data sources.
- Provide captured change data to PowerCenter when you run a PowerCenter CDC session.
- Provide captured change data or source table data to the PowerExchange Navigator when you perform a database row test of an extraction map or a data map.
- Interact with other PowerExchange Listeners on other nodes to facilitate communication among the PowerExchange Navigator, PowerCenter Integration Service, data sources, and any system to which PowerExchange processing is offloaded.

## Customizing the dbmover Configuration File for CDC

You must configure some statements in the dbmover configuration file for CDC processing.

The PowerExchange Listener uses the dbmover statements to perform the following functions:

- Connect to source relational databases and objects to capture change data.
- Determine the directory in which to store capture registrations, extraction maps, and PowerExchange Logger log files.
- Connect to the system where the PowerExchange Logger log files reside to extract change data.

The following key dbmover statements are required for CDC and pertain to all source RDBMSs that PowerExchange supports on Linux, UNIX, or Windows:

- CAPI\_CONNECTION statements
  - Source-specific CAPI\_CONNECTION statements, which are described for each source type
  - A CAPX CAPI\_CONNECTION if you use continuous extraction mode
- CAPI\_SRC\_DFLT statement
- CAPT\_PATH statement
- CAPT\_XTRA statement

Review the descriptions of each of these parameters. For more information about other dbmover.cfg statements, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“DB2 for Linux, UNIX, and Windows CDC” on page 77](#)
- [“Microsoft SQL Server CDC” on page 94](#)
- [“Oracle CDC with LogMiner” on page 109](#)

## CAPI\_CONNECTION Statements

PowerExchange requires that you define CAPI\_CONNECTION statements in the dbmover configuration file on any Linux, UNIX, or Windows system where PowerExchange captures or extracts change data.

PowerExchange uses the CAPI\_CONNECTION statements to connect to the change stream for a source to extract change data.

For each data source, you must define one of the following source-specific types of CAPI\_CONNECTION statements:

- For Microsoft SQL Server sources, an MSQL CAPI\_CONNECTION
- For Oracle CDC with LogMiner sources, an ORCL CAPI\_CONNECTION and a UOWC CAPI\_CONNECTION for the UOW Cleanser
- For Express CDC for Oracle sources, an ORAD CAPI\_CONNECTION
- For DB2 for Linux, UNIX, and Windows sources, a UDB CAPI\_CONNECTION

If you use continuous extraction mode to extract change data from PowerExchange Logger log files, you must also define a CAPX CAPI\_CONNECTION statement.

You can specify up to eight source-type CAPI\_CONNECTION statements in a dbmover configuration file, excluding CAPX CAPI\_CONNECTION statements. You can identify one of these statements as the overall default CAPI\_CONNECTION statement. If you define multiple CAPI\_CONNECTION statements for the same source type, you can also specify a source-specific default. In addition to or in lieu of the defaults, you can define specific CAPI\_CONNECTION overrides in multiple ways. The order of precedence that PowerExchange uses to determine which CAPI\_CONNECTION statement to use is described in the *PowerExchange Reference Manual*.

**Note:** To perform database row tests for data sources that are defined by capture registrations local to the PowerExchange Navigator, you must specify the appropriate CAPI\_CONNECTION statements on the PowerExchange Navigator machine. Otherwise, you do not need to specify CAPI\_CONNECTION statements to perform database row tests.

## RELATED TOPICS:

- [“CAPI\\_CONNECTION - CAPX Statement” on page 57](#)
- [“CAPI\\_CONNECTION - MSQL Statement” on page 102](#)
- [“CAPI\\_CONNECTION - UDB Statement” on page 85](#)
- [“CAPI\\_CONNECTION - ORCL Statement” on page 125](#)
- [“CAPI\\_CONNECTION - UOWC Statement” on page 132](#)
- [“CAPI\\_CONNECTION - ORAD Statement” on page 168](#)

## CAPI\_SRC\_DFLT Statement

The CAPI\_SRC\_DFLT statement specifies the CAPI\_CONNECTION statement that PowerExchange uses by default for a specific data source type when a CAPI connection override is not supplied.

Informatica recommends that you specify this statement for each source type to ensure that the appropriate type of CAPI\_CONNECTION statement is used for accessing a source. If you have multiple source types that run on different platforms and do not specify a CAPI\_SRC\_DFLT statement for each source type, source access problems might occur in certain situations during PowerCenter data previews or CDC sessions.

**Operating Systems:** All

**Data Sources:** All

**Related Statements:** CAPI\_CONN\_NAME and CAPI\_CONNECTION

**Required:** No

**Syntax:**

```
CAPI_SRC_DFLT=(source_type
               ,capi_connection_name)
```

**Parameters:**

**source\_type**

Required. The CDC source type.

The following table describes the possible values:

Option	Source Type
ADA	Adabas sources
AS4	DB2 for i5/OS sources
CAPX	Sources for which you are extracting data in continuous extraction mode and using the PowerExchange Logger for Linux, UNIX, or Windows or PowerExchange Condense
DB2	DB2 for z/OS sources
DCM	CA Datacom sources
IDL	CA IDMS/DB log-based CDC sources
IML	IMS log-based CDC sources
IMS	IMS synchronous CDC sources

Option	Source Type
MSS	Microsoft SQL Server sources
ORA	Oracle sources
UDB	DB2 for Linux, UNIX, and Windows sources
VSAM or VSM	VSAM sources

#### ***capi\_connection\_name***

Required. Unique name of the CAPI\_CONNECTION statement to use as the default for the specified source type. This name must match the NAME value in a CAPI\_CONNECTION statement that has a TYPE value that is compatible with the CAPI\_SRC\_DFLT *source\_type*.

The following table shows, for each CAPI\_SRC\_DFLT option, the compatible CAPI\_CONNECTION type:

CAPI_SRC_DFLT Option	CAPI_CONNECTION Statement Type
AS4	UOWC
CAPX	CAPX
ADA, DB2, DCM, IDL, IDM, IML, IMS, VSAM, or VSM	UOWC
ORA	UOWC for PowerExchange Oracle CDC with LogMiner ORAD for PowerExchange Express CDC for Oracle
MSS	MSQL
UDB	UDB

#### **Usage Notes:**

- If you define multiple CAPI\_CONNECTION statements for a source type, you can define a CAPI\_SRC\_DFLT statement to identify the default CAPI\_CONNECTION for that source type. The CAPI\_SRC\_DFLT statement must point to a CAPI\_CONNECTION statement of a compatible type.
- You can optionally define a CAPI\_CONN\_NAME statement that specifies an overall default statement, out of all of the CAPI\_CONNECTION statements in the DBMOVER file.
- Instead of or in addition to specifying defaults, you can use the following CAPI connection name overrides to point to a specific CAPI\_CONNECTION statement for CDC sessions or database row tests:
  - For CDC sessions, use the **CAPI Connection Name Override** attribute on the PWX CDC application connection.
  - For PowerExchange Condense, use the CONN\_OVR parameter in CAPTPARM configuration file.
  - For the PowerExchange Logger for Linux, UNIX, and Windows, use the CONN\_OVR parameter in pwxccl.cfg configuration file.
  - For DTLUAPPL utility operations that generate restart tokens, use the CONN\_OVR parameter in the DTLUAPPL control statement.

- For CAPXRT database row tests in the PowerExchange Navigator, use the **CAPI Connection Name** value in the **CAPXRT Advanced Parameters** dialog box. If you add an SQL statement for generating restart tokens, you can include the CONNAME parameter to point to the override CAPI\_CONNECTION.
- For PowerExchange ODBC connections, use the DTLCONN\_OVR parameter in the odbc.ini file or the SQL escape sequence override DTLCONN\_OVR.

## CAPT\_PATH Statement

The CAPT\_PATH statement specifies the path to a directory on a Linux, UNIX, or Windows system that contains the control files for CDC.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Microsoft SQL Server, Oracle, DB2 for Linux, UNIX, and Windows, and if you use offloading processing, other data sources on i5/OS or z/OS

**Related Statements:** CAPT\_XTRA

**Required:** Yes for CDC sources on Linux, UNIX, and Windows

**Syntax:**

```
CAPT_PATH=path
```

**Value:** For the *path* variable, enter the path to the local directory that contains the following control files for CDC:

- CCT file, which contains capture registrations.
- CDEP file, which contains application names for any PowerCenter extractions that use ODBC connections.
- CDCT file, which contains information about PowerExchange Logger for Linux, UNIX, and Windows log files.

This directory can be one that you created specifically for these files or another directory.

Informatica recommends that you use a unique directory name to separate these CDC objects from the PowerExchange code. This practice makes migrating to a another PowerExchange version easier.

Default is the PowerExchange installation directory.

**Usage Notes:**

- PowerExchange C-ISAM control files, such as the CCT, CDEP, and CDCT files, must be stored on local disk. Do not locate these files in SAN or NAS storage.
- To provide a path to the directory that contains extraction maps, use the CAPT\_XTRA statement.

## CAPT\_XTRA Statement

The CAPT\_XTRA statement specifies the path to the local directory on a Linux, UNIX, or Windows system that stores extraction maps for CDC.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Microsoft SQL Server, Oracle, and DB2 for Linux, UNIX, and Windows

**Related Statements:** CAPT\_PATH

**Required:** Yes for CDC sources on Linux, UNIX, and Windows

**Syntax:**

```
CAPT_XTRA=path
```

**Value:** For the *path* variable, enter the path to the local directory that stores extraction maps.

This directory can be a directory that you created specifically for extraction maps or another directory. Default is the PowerExchange installation directory.

**Usage Notes:** To specify the directory that contains the CCT and CDEP control files for CDC, use the CAPT\_PATH statement.

## Starting the PowerExchange Listener

To start the PowerExchange Listener, you can run the dtllst program or use other system-specific methods.

**Note:** You cannot use the pwxcmd or infacmd program to start the PowerExchange Listener.

### Starting the PowerExchange Listener on Linux or UNIX

To start the PowerExchange Listener on a Linux or UNIX system, use one of the following methods:

- Enter dtllst at the command line to run the PowerExchange Listener in foreground mode. Syntax is:

```
dtllst node1 [config=directory/myconfig_file] [license=directory/mylicense_key_file]
```

Include the optional config and license parameters if you want to specify configuration and license key files that override the original dbmover.cfg and license.key files.

Add an ampersand (&) at the end to run the PowerExchange Listener in background mode. Also, you can add the prefix nohup at the beginning to run the PowerExchange Listener persistently. Syntax is:

```
nohup dtllst node1 [config=directory/myconfig_file] [license=directory/  
mylicense_key_file] &
```

- Run the startlst script, which is supplied by PowerExchange installation. This script deletes the detail.log file and then starts the PowerExchange Listener.

**Caution:** If you run PowerExchange and PowerCenter on the same machine, using the same user account, you must create separate environments for PowerExchange and PowerCenter. To create the appropriate PowerExchange environment and start the PowerExchange Listener, run the pwxsettask.sh script.

Use the following syntax:

```
pwxsettask.sh dtllst node_name ["config=directory/config_file"] ["license=directory/  
license_key_file"]
```

The quotation marks are optional.

For more information, see [“Environment Variable Incompatibilities Between PowerExchange and PowerCenter” on page 23](#).

### Starting the PowerExchange Listener on Windows

To start the PowerExchange Listener on a Windows system, use one of the following methods:

- Run the PowerExchange Listener as a Windows service by completing one of the following actions:
  - From the Windows Start menu, click **Start > Programs > Informatica PowerExchange > Start PowerExchange Listener**.

- Use the dtllstsi program to enter the start command from a Windows command prompt:

```
dtllstsi start "service_name"
```

- Enter dtllst.

The syntax is the same as that for Linux and UNIX except that the & and nohup operands are not supported. You must have a product license that allows you to manually run dtllst.

If you use the PowerExchange Logger for Linux, UNIX, and Windows, the PowerExchange Listener must run under a user ID that has READ access to the PowerExchange Logger log files.

If you run the PowerExchange Listener as an application service in the Informatica domain, enable the PowerExchange Listener Service from the Informatica Administrator tool to start it. For more information, see the *Informatica Application Service Guide*.

**Caution:** If you run PowerExchange and PowerCenter on the same machine, using the same user account, you must create separate environments for PowerExchange and PowerCenter. To create the appropriate PowerExchange environment and start the PowerExchange Listener, run the pwxsettask.bat script.

Use the following syntax:

```
pwxsettask dtllst node_name ["config=directory/config_file"] ["license=directory/  
license_key_file"]
```

The quotation marks are required.

For more information, see ["Environment Variable Incompatibilities Between PowerExchange and PowerCenter" on page 23](#).



# Stopping the PowerExchange Listener

To stop the PowerExchange Listener, use the CLOSE or CLOSE FORCE command. To stop active PowerExchange Listener tasks, use the STOPTASK command.

The following table describes these commands and the syntax for issuing each command from the command line against a PowerExchange Listener task that is running in foreground mode:

Command	Description	Command Line Syntax
CLOSE	Stops the PowerExchange Listener after all of the following subtasks complete: <ul style="list-style-type: none"><li>- CDC subtasks, which stop at the next commit of a unit of work (UOW)</li><li>- Bulk data movement subtasks</li><li>- PowerExchange Listener subtasks</li></ul>	On Linux, UNIX, or Windows:  C
CLOSE FORCE	Forces the cancellation of all user subtasks and stops the PowerExchange Listener. PowerExchange waits 30 seconds for current user subtasks on the PowerExchange Listener to complete. Then PowerExchange cancels any remaining user subtasks and stops the PowerExchange Listener. This command is useful if you have long-running subtasks on the PowerExchange Listener.	On Linux or UNIX:  C F  On Windows:  CF
STOPTASK	Stops a PowerExchange Listener task for a specific extraction application process. PowerExchange waits to stop the PowerExchange Listener until either the end UOW or commit threshold is reached.	On Linux, UNIX, and Windows:  STOPTASK <i>app_name</i>  The <i>app_name</i> is the name of an active change data extraction process. You can get this name from a PWX-00712 message in the PowerExchange Listener DISPLAY ACTIVE command output.

Alternatively, you can use any of the following methods:

- On a Linux, UNIX, or Windows system, use the pwxcmd program to issue the close, closeforce, or stoptask command to a PowerExchange Listener running in foreground or background mode, on the local system or a remote system. You can issue these pwxcmd commands from the command line or include them in scripts or batch files.
- On a Linux or UNIX system, if the PowerExchange Listener is running in background mode, use the standard operating system commands to find the PowerExchange Listener process ID and then “kill” that process. A “kill” operation is similar to a CLOSE operation.
- On a Windows system, if the PowerExchange Listener does not respond to a CLOSE FORCE command, press Ctrl + C once to issue CLOSE or press Ctrl + C twice to issue CLOSE FORCE.

# Displaying Active PowerExchange Listener Tasks

You can use the DISPLAY ACTIVE command to display information about each active PowerExchange Listener task that is running in foreground mode on a Linux, UNIX, or Windows system. This information includes the TCP/IP address, port number, application name, access type, and status.

On a Linux, UNIX, or Windows system, enter the following command at the command line on the screen where the PowerExchange Listener task is running in foreground mode:

D

Alternatively, on a Linux, UNIX, or Windows system, you can issue the `pwxcmd listtask` command from a command line, script, or batch file to a PowerExchange Listener running on the local system or a remote system. The `pwxcmd listtask` command produces the same output as the `DISPLAY ACTIVE` command.

## CHAPTER 3

# PowerExchange Logger for Linux, UNIX, and Windows

This chapter includes the following topics:

- [PowerExchange Logger Overview, 35](#)
- [PowerExchange Logger Tasks, 37](#)
- [PowerExchange Logger Files, 37](#)
- [File Switches, 40](#)
- [PowerExchange Logger Operational Modes, 41](#)
- [PowerExchange Logger Usage Considerations, 42](#)
- [Logging Data from Remote z/OS or i5/OS Sources, 43](#)
- [Configuring the PowerExchange Logger, 43](#)
- [Starting the PowerExchange Logger, 63](#)
- [Managing the PowerExchange Logger, 68](#)

## PowerExchange Logger Overview

The PowerExchange Logger for Linux, UNIX, and Windows captures change data from PowerExchange data sources and writes that data to PowerExchange Logger log files. The PowerExchange Logger writes only the successful units of work (UOWs) to its log files, in chronological order based on end time.

When a PowerCenter CDC session runs, it extracts change data from the log files instead of from the change stream.

**Note:** The PowerExchange Logger for Linux, UNIX, and Windows is similar in function to PowerExchange Condense on i5/OS or z/OS systems.

The PowerExchange Logger can capture change data from DB2 database logs or Oracle redo logs on Linux, UNIX, or Windows, or from a Microsoft SQL Server distribution database on a Windows. If you use the offloading feature, a PowerExchange Logger process on Linux, UNIX, or Windows can also process data from data sources on i5/OS or z/OS systems.

Use the PowerExchange Logger to reduce the overhead of CDC processing. With the PowerExchange Logger, PowerExchange accesses the source database fewer times to read change data, which reduces database I/O. Also, because change data is extracted from the PowerExchange Logger log files, you often do not need to extend the retention period for source database log files to accommodate CDC processing.

You must run one PowerExchange Logger process for each source type and instance, as defined in a registration group. The PowerExchange Logger can run on the source database server, PowerCenter Integration Service machine, or another system in either continuous mode or batch mode.

Multiple PowerExchange Logger instances can run under the same PowerExchange Listener and dbmover.cfg configuration. However, because a single dbmover.cfg can contain a maximum of eight CAPI\_CONNECTION statements, the number of source instances and PowerExchange Logger instances that can run under a single Listener and dbmover.cfg is limited. For more information, see [“CAPI\\_CONNECTION Statements” on page 27](#).

When you create capture registrations for data sources, including i5/OS and z/OS data sources for which processing is offloaded, set the **Condense** option to **Part**. The PowerExchange Logger supports only partial condense processing. For i5/OS or z/OS data sources, if you set the **Condense** option to **Full** in capture registrations, the PowerExchange Logger ignores the registrations and does not process change data from those sources.

For each PowerExchange Logger process, you must define a configuration file. PowerExchange provides a sample configuration file named pwxcl.cfg. The configuration file contains parameters for controlling the PowerExchange Logger and for identifying the source instance. Use the COLL\_END\_LOG parameter to control whether the PowerExchange Logger runs in continuous mode or batch mode.

When PowerCenter CDC sessions run, the PowerExchange Log Reader process extracts change data from the PowerExchange Logger log files in continuous extraction mode or batch extraction mode.

**Tip:** For PowerExchange Oracle CDC with LogMiner, Informatica recommends that you use the PowerExchange Logger and continuous extraction mode. PowerExchange can then use one Oracle LogMiner session for all extractions that process an Oracle instance. If you use real-time extraction mode without the PowerExchange Logger, PowerExchange starts a separate LogMiner session for each extraction. Multiple, concurrent LogMiner sessions can significantly degrade the performance on the system where LogMiner runs. For PowerExchange Express CDC for Oracle, this LogMiner consideration does not apply.

You can secure sensitive data that is stored in PowerExchange Logger log files, such as Social Security numbers, by enabling AES encryption of the log files. In the PowerExchange Logger configuration file, you can select the AES encryption algorithm that you want to use. To enable encryption, you must also specify an encryption password either in the PowerExchange Logger configuration file or in the pwxcl command that you use to cold start the PowerExchange Logger from the command line. If you specify the encryption password in the pwxcl command for a cold start and need to restore the CDCT file later, you must enter the same encryption password for the restore operation.

**Tip:** To reduce the risk of unauthorized access to the encryption password, Informatica recommends that you specify the password in the pwxcl command for cold starting the Logger rather than specify the password in the configuration file.

## RELATED TOPICS:

- [“PowerExchange Logger Operational Modes” on page 41](#)
- [“PowerExchange Logger Tasks” on page 37](#)
- [“PowerExchange Logger Files” on page 37](#)
- [“File Switches” on page 40](#)
- [“Configuring the PowerExchange Logger” on page 43](#)
- [“Managing the PowerExchange Logger” on page 68](#)

# PowerExchange Logger Tasks

The PowerExchange Logger for Linux, UNIX, and Windows includes the following task and subtasks:

## **Controller task**

Loads parameter settings from the PowerExchange Logger pwxcl configuration file. Loads the capture registrations from the CCT file. After loading this information, the Controller starts the Command Handler subtask and then the Writer subtask.

## **Command Handler subtask**

Processes PowerExchange Logger commands from various sources, including user stdin and the pwxcmd program. If the PROMPT parameter is set to Y in the pwxcl.cfg file, the Command Handler waits for the Writer subtask to initialize before accepting a user command.

## **Writer subtask**

Performs most of the PowerExchange Logger work that uses CPU time. The Writer initializes the CAPI for the source database, determines the start or restart point in the change stream, reads change data from the change stream, and writes change data to PowerExchange Logger log files. The Writer also writes records to the CDCT file during a file switch, deletes expired CDCT records, and rolls back CDCT records when you warm start the PowerExchange Logger from an earlier point in time. If the PROMPT parameter is set to Y in the pwxcl configuration file, the Writer waits for you to respond to confirmation prompts before proceeding with a cold start or a rollback of CDCT records.

**Note:** The Log Reader, which extracts change data from PowerExchange Logger log files, runs as an independent process.

# PowerExchange Logger Files

A PowerExchange Logger process writes information to the CDCT file, PowerExchange Logger log files, and PowerExchange message logs.

The PowerExchange Logger also uses lock files during processing.

## **RELATED TOPICS:**

- [“CDCT File” on page 37](#)
- [“PowerExchange Logger Log Files” on page 38](#)
- [“Lock Files” on page 39](#)
- [“Message Log Files” on page 40](#)

## CDCT File

The PowerExchange Logger stores log file and restart information in the CDCT file.

When a PowerCenter CDC session runs in continuous extraction mode or batch extraction mode, the PowerExchange Listener reads the CDCT file to determine the PowerExchange Logger log files from which to extract change data.

The PowerExchange Logger creates the CDCT file in the directory that is specified by the CAPT\_PATH statement in the dbmover configuration file that is on the system where the PowerExchange Logger runs. The

CDCT file must reside on local disk. If the CAPT\_PATH statement is not specified, the CDCT file is in the local directory from which the PowerExchange Logger is invoked.

The generated CDCT file format is independent of the operating system type and system endianness. The file name has the format CDCT\_*dbid*, where *dbid* is the DBID value in the PowerExchange Logger configuration file.

The first time the PowerExchange Logger receives data of interest after startup or a file switch, the PowerExchange Logger opens a log file to which to write data and creates an entry for this log file in the CDCT file. After the PowerExchange Logger finishes writing data to the log file, it marks the file as closed in the CDCT file. The PowerExchange Logger also updates the restart information in the CDCT file during periods when no changes of CDC interest are received.

If a PowerExchange Logger failure occurs and leaves an open log file, the PowerExchange Logger marks the CDCT entry for that log file for deletion. The next time the PowerExchange Logger starts, it deletes that log file entry and creates a new log file. The log files contain records from which the CDCT entry for the log can be rebuilt.

PowerExchange automatically generates a backup of the CDCT file at PowerExchange Logger initialization and normal termination. These backups are located in the same directory as the CDCT file and have file names with the following formats: CDCT\_*dbid*\_INIT.bkp and CDCT\_*dbid*\_TERM.bkp.

**Tip:** You can use the PWXUCDCT utility to print information about CDCT records, back up the CDCT file, restore the CDCT file from a backup, re-create the CDCT file based on PowerExchange Logger log files, and delete expired CDCT records.

## RELATED TOPICS:

- [“PWXUCDCT Commands for Maintaining the PowerExchange Logger CDCT and Log Files” on page 72](#)

## PowerExchange Logger Log Files

The PowerExchange Logger creates log files for storing change data records when it first encounters changes for source tables and columns of interest. These source tables and columns must be defined in active capture registrations.

The PowerExchange Logger creates log files based on the EXT\_CAPT\_MASK parameter in the pwxctl configuration file. This parameter specifies a path to the directory where the log files are stored and a prefix for the log file names. Log file names have the following format:

*path/prefix.CND.CPyyymmdd.Thhmmnnn*

Where:

- *path/prefix* is the EXT\_CAPT\_MASK value.
- *yyymmdd* is the date when the file is created.
- *hhmm* is a 24-hour time when the file is created.
- *nnn* is a generated sequence number that makes each file name unique.

The log files remain open until a file switch occurs or the PowerExchange Logger shuts down.

When you run a PowerCenter CDC session in continuous extraction mode or batch extraction mode, PowerExchange extracts change data from the PowerExchange Logger log files.

## Lock Files

During initialization, a PowerExchange Logger process creates lock files to prevent other PowerExchange Logger processes from accessing the same CDCT file and log files concurrently.

As long as the PowerExchange Logger process holds a lock on the lock files, locking is in effect for the resources for which the lock files were created.

PowerExchange Logger locking works on local disks on Linux, UNIX, or Windows systems. It also works on the following shared file systems on Linux or UNIX systems:

- Veritas Storage Foundation™ Cluster File System by Symantec
- IBM General Parallel File System
- EMC Celerra network-attached storage (NAS) with Network File System (NFS) protocol version 3
- NetApp NAS with NFS version 3

The PowerExchange Logger creates lock files in the following order:

1. A lock file for the CDCT file for a source instance. The PowerExchange Logger generates the lock file name and location based on the directory that is specified in the CAPT\_PATH parameter of the dbmover configuration file.
2. One of the following lock files:
  - If you use a group definition file, a lock file for each set of the PowerExchange Logger log files that is defined by the GROUP statements in the group definition file. The PowerExchange Logger generates the lock file names and locations based on the external\_capture\_mask parameter in each GROUP statement. In this case, the PowerExchange Logger ignores the EXT\_CAPT\_MASK parameter in the pwxcl configuration file when creating lock files and processing log files.
  - If you do not use a group definition file, a lock file for PowerExchange Logger log files. The PowerExchange Logger generates the lock file name and location based the directory and file-name prefix that are specified in the EXT\_CAPT\_MASK parameter of the pwxcl configuration file.

Lock file names end with \_lockfile.lck. For example, a lock file for the CDCT file could have the name CDCT\_oracoll1\_lockfile.lck.

When the PowerExchange Logger process ends, it unlocks the lock files to enable other PowerExchange Logger processes to access the previously locked resources.

To identify a PowerExchange Logger process that holds a lock, look up the process ID (PID) in the Task Manager on a Windows system or issue the ps command on a UNIX or Linux system.

Also, the PowerExchange Logger writes messages to the PowerExchange message log that indicate the locking status. Look for the following key messages:

- To verify that lock files are created, look for PWX-25802 messages, such as:

```
PWX-25802 Process pwxcl.exe pid 5428 locked file C:\capture\captpath
\CDCT_instance_lockfile.lck
```

- To verify that lock files are unlocked, look for PWX-25803 messages, such as:

```
PWX-25803 Process pwxcl.exe pid 5428 unlocked file C:\capture\extcapt
\loggerfiles_lockfile.lck
```

- If the PowerExchange Logger process cannot find the lock file that it needs to access some resources, it writes message PWX-25800:

```
PWX-25800 Could not find lock file file_name
```

- If a lock file is locked by another process, the PowerExchange Logger process writes some or all of the following messages, depending on if it can acquire a lock before the maximum retry interval that is specified in PWX-25814 elapses:

```
PWX-25804 Error trying to lock PowerExchange Logger files
PWX-25811 File file_name is locked by process process_name pid process_id on host
host_name date date time time
PWX-25812 File file_name is locked by pid process_id start offset length bytes
PWX-25813 No information is available on process which locked file file_name
PWX-25814 Trying to lock file file_name until number seconds elapses
PWX-25815 File file_name is locked by another process and no more waiting is allowed.
```

If a PowerExchange Logger process ends abnormally with message PWX-25815 and return code 25815, try to determine the status of the other PowerExchange Logger process that is holding the lock. This other process is identified in message PWX-25811. For example, the other process might not have completely shut down, or both processes might be trying to use the same files because of an error in their pwxcl configuration files.

## Message Log Files

The PowerExchange Logger writes messages to the PowerExchange message log file.

By default, on Linux, UNIX, and Windows, this file is named `detail.log` and is located in the working directory where the PowerExchange Logger process runs. However, you can optionally specify another directory for PowerExchange message log files. You can also enable the use of alternative log files.

To specify a unique directory for PowerExchange message log files, include the `LOGPATH` parameter in the `dbmover` configuration file. You can then find the PowerExchange message log files more easily.

Also, you can implement alternative logging by specifying the `TRACING` statement in the `dbmover` configuration file. When alternative logging is enabled, PowerExchange creates a set of alternative log files for each PowerExchange process, including each PowerExchange Logger process, in a separate directory. When an alternative log file becomes full, PowerExchange switches to another alternative log file. This automatic rotation of message log files prevents out-of-space conditions. Also, PowerExchange buffers messages before writing them to the alternative log files on disk at a specific flush interval. This mode of writing messages can reduce I/O activity on the alternative log files.

## File Switches

When running in continuous mode, the PowerExchange Logger periodically closes its open log files if they contain data and then opens a new set of log files. This process is called a *file switch*.

The PowerExchange Logger automatically performs a file switch when the criteria in the following parameters of the pwxcl configuration file are met:

- `FILE_SWITCH_CRIT`
- `FILE_SWITCH_VAL`

If the open log files do not contain data when the file-switch criteria in these parameters are met, the file switch does not occur. The PowerExchange Logger waits until the next time the file-switch criteria are met. If the files still do not contain data, the PowerExchange Logger continues to check the log files at set intervals. Only when the log files contain data does the file switch occur.

Also, you can force a file switch by entering the `files witch` command from the command line. Alternatively, on Linux, UNIX, or Windows, you can send a `pwxcmd fileswitch` command to a PowerExchange Logger process running on the local system or a remote system.



If a file switch is initiated automatically or by the `fileswitch` command when the PowerExchange Logger is not on a commit boundary in the change stream, the PowerExchange Logger waits 10 seconds for the commit to occur and then forces the file switch.

## PowerExchange Logger Operational Modes

A PowerExchange Logger process can operate in continuous mode or batch mode.

To set the operational mode, use the `COLL_END_LOG` parameter in the `pwxccl` configuration file.

### Continuous Mode

In continuous mode, the PowerExchange Logger process runs continuously until you manually stop it.

Use continuous mode in the following situations:

- You have a database with a high level of change activity that occurs continuously.
- You have a database with intermittent activity that occurs at unpredictable intervals.
- You want to avoid the overhead of scheduling PowerExchange Logger runs.
- You cannot restart the PowerExchange Logger process often enough to keep up with the change volume.

To enable continuous mode, set the `COLL_END_LOG` parameter to 0.

In continuous mode, each time the Writer subtask completes a logging cycle, the PowerExchange Logger process is temporarily suspended. The next logging cycle is triggered by any of the following events:

- The wait interval that is defined in the `NO_DATA_WAIT` parameter of the `pwxccl` configuration file elapses.
- The `CONDENSE` command is manually entered at the command line or with the `pwxcmd` program.
- The `FILESWITCH` command is manually entered at the command line or with the `pwxcmd` program.

The PowerExchange Logger process continues to run until you enter the `SHUTDOWN` or `SHUTCOND` command. To prevent log files from becoming too large, the PowerExchange Logger process periodically performs a file switch. Log files that are very large can prolong restart times for CDC sessions that run in continuous extraction mode or batch extraction mode.

You can use the `NO_DATA_WAIT2` parameter in the `pwxccl` configuration file to prevent the PowerExchange Logger from consuming too much CPU time when PowerExchange is not receiving change data. For example, if you set the `NO_DATA_WAIT2` parameter to 30 seconds, the PowerExchange Logger sleeps for 30 seconds, provided that no updates are received, and then performs another processing cycle. However, a large `NO_DATA_WAIT2` value can delay processing of a `SHUTDOWN` command. If you need to reduce the amount of time that the PowerExchange Logger sleeps on a quiet system, you can adjust the `FILE_FLUSH_VAL` and `FILE_SWITCH_VAL` parameters.

Also, specify the `RSTRADV` time interval in the source-specific `MSQL`, `UDB`, or `UOWC` `CAPL_CONNECTION` statement or the `OPTIONS` statement of the PowerExchange Express CDC for Oracle configuration file to enable the PowerExchange Logger to advance its restart and sequence tokens even when UOWs do not contain any change data of interest for the data sources.

When you run the PowerExchange Logger in continuous mode, you can use either continuous or batch extraction mode for workflows that extract change data from the PowerExchange Logger log files.

**Tip:** Run the PowerExchange Logger in continuous mode unless you have a reason to use batch mode. On Linux or UNIX, you can run a continuous PowerExchange Logger process in background mode and use the `pwxcmd` program to send commands to that background PowerExchange Logger process.

## Batch Mode

In batch mode, the PowerExchange Logger process shuts down after it reaches the end-of-log (EOL) and waits for the interval that is specified in the NO\_DATA\_WAIT2 parameter of the pwxcl configuration file without receiving additional change data.

Use batch mode in the following situations:

- You run the PowerExchange Logger on a scheduled basis, after batch applications that update the database complete.
- You run the PowerExchange Logger manually for testing or other purposes.

To enable batch mode, set the COLL\_END\_LOG parameter to 1 in the pwxcl configuration file.

When you run the PowerExchange Logger in batch mode, use batch extraction mode for any workflows that extract change data from the PowerExchange Logger log files.

## PowerExchange Logger Usage Considerations

Before you run the PowerExchange Logger for Linux, UNIX, or Windows, review the following usage considerations:

- Informatica recommends that you run the PowerExchange Logger for Linux, UNIX, and Windows and the process that extracts data from the PowerExchange Logger log files on the same system. Configure a PowerExchange Listener on the PowerExchange Logger system and specify that node location in the **Location** attribute of the PowerCenter PWX CDC connection. If you run the PowerExchange Logger and extraction process on different systems and use NFS to access the log files, disable attribute and data caching for the NFS mount point. Otherwise, session failures might occur.
- On Linux and UNIX, the PowerExchange Logger requires sufficient amounts of main memory and virtual memory to process change data. If the memory is not sufficient, PowerExchange writes the error messages PWX-00271 and PWX-00904 to the PowerExchange message log file when you attempt to start the PowerExchange Logger on Linux or UNIX.  
To prevent this problem, use the Linux or UNIX ulimit command to set the size limits for maximum memory and virtual memory to unlimited. The specific ulimit syntax varies by platform and shell. For more information about this command, see the documentation for your Linux or UNIX operating system.
- On Linux and UNIX, you can run a PowerExchange Logger process in background mode. For background PowerExchange Logger processes, Informatica recommends that you set the COLL\_END\_LOG parameter to 0 in the pwxcl configuration file to run the PowerExchange Logger continuously. Also, set the PROMPT parameter to N. If you use PROMPT=Y, the PowerExchange Logger ignores this setting and issues an error message. To send commands to a PowerExchange Logger process that is running in the background, use the pwxcmd program. To enable pwxcmd use, define the CONDENSENAME statement in the pwxcl configuration file and define the SVCNODE statement in the dbmover configuration file.

# Logging Data from Remote z/OS or i5/OS Sources

You can log data for a z/OS or i5/OS data source to remote PowerExchange Logger log files on Linux, UNIX, or Windows log files.

CDC sessions that run in continuous extraction mode can then extract the change data from the PowerExchange Logger log files instead of from the source. This practice can reduce the amount of time for CDC processing on the z/OS or i5/OS system.

To provide the highest level of security for z/OS data sources, set the first parameter in the SECURITY statement to 2 in the z/OS DBMOVER configuration file. With this setting, the PowerExchange Logger for Linux, UNIX, and Windows can log data from z/OS systems only if its user credentials pass z/OS security checking. The PowerExchange Logger must use a valid z/OS user ID and password combination that has READ access to CAPX.REG.\* resource profiles in the FACILITY class, which are managed by your z/OS security product.

For more information, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“Remote Logging Overview” on page 188](#)

# Configuring the PowerExchange Logger

To configure the PowerExchange Logger, you must define a PowerExchange Logger configuration file for each source type and instance, as defined in a registration group.

Also, verify that the **Condense** option is set to **Part** in the capture registrations for all source tables that the PowerExchange Logger will process.

If you want the PowerExchange Logger to create separate sets of log files for different groups of tables, create a PowerExchange group definition file. The group definition file defines the group name, the path and file-name prefix for the log files, and the registrations in the group.

## RELATED TOPICS:

- [“Customizing the PowerExchange Logger Configuration File” on page 44](#)
- [“Customizing the dbmover Configuration File for the PowerExchange Logger” on page 56](#)
- [“Using PowerExchange Logger Group Definitions” on page 60](#)
- [“Enabling a Capture Registration for PowerExchange Logger Use” on page 43](#)

# Enabling a Capture Registration for PowerExchange Logger Use

For the PowerExchange Logger to use a capture registration, the registration must have a status of active and a **Condense** setting of **Part**.

If the PowerExchange Logger does not find any active capture registration, the PowerExchange Logger issues error message PWX-06427 and ends.

To enable a capture registration for PowerExchange Logger use:

1. In the PowerExchange Navigator, open the capture registration.
2. In the Resource Inspector, select **Active** in the **Status** list.

3. In the **Condense** list, select **Part**.

## Customizing the PowerExchange Logger Configuration File

Before you start the PowerExchange Logger, configure its parameters in the PowerExchange Logger configuration file.

PowerExchange provides an example configuration file, named `pwxccl`, in the PowerExchange installation directory that is specified in the `PWX_HOME` environment variable on Linux or UNIX or in the `PATH` environment variable on Windows. Use this example file as a starting point for creating a customized file. To preserve the example file in its original state, rename it and copy it to another directory. Then customize the copy. You must specify the `CS` parameter when you start the PowerExchange Logger to identify the custom configuration file path and file name.

When customizing the configuration file, if you enter a parameter value such as a Windows path that contains one or more spaces, enclose the value in double quotation marks (`"`). Make sure that you use straight quotation marks.

**Note:** If you used the similar PowerExchange Condense component in a PowerExchange release earlier than 8.6.1, you can copy its `dtlca.cfg` configuration file and then customize the copy. Rename the file to `pwxccl` or use the `CS` execution parameter. The PowerExchange Condense component is no longer supported on Linux, UNIX, and Windows.

### Parameter Descriptions

You can specify the following PowerExchange Logger parameters in the `pwxccl.cfg` configuration file.

The parameter descriptions include the following notational conventions:

- *Italics* indicate a variable.
- Curly braces { } enclose alternative options. Enter only one option. Do not type the braces when you enter the option.
- A vertical bar | indicates a mutually exclusive choice. When used with braces, you must enter one of the items.

#### **CAPT\_IMAGE={AI|BA}**

Type of data images that the PowerExchange Logger stores in its log files. Use this parameter to control whether the PowerExchange Logger stores after images only or both before and after images of the data in its log files.

This parameter affects the amount of storage that you use for PowerExchange Logger log files and whether before image data is available for use in extraction processing.

Enter one of the following options:

- **AI**. Stores only after images in the PowerExchange Logger log files.
- **BA**. Stores both before and after images in the PowerExchange Logger log files.

Default is **AI**. With **AI**, less storage is required for PowerExchange Logger log files. However, the following CDC limitations apply:

- You cannot use before images of the data in extraction processing. If you add before image (BI) fields to extraction maps, PowerCenter CDC sessions that reference the BI fields fail.
- If you add change indicator (CI) fields to extraction maps, PowerCenter CDC sessions that reference the CI fields fail.

Informatica recommends that you enter **BA** if you have sufficient storage for larger log files. The sample pwxcl configuration file that PowerExchange supplies specifies **BA**.

**Note:** If you use **BA** and add CI columns to the extraction maps, any Insert and Delete operations on the source result in Null values in the CI columns. Any Update operations on the source result in the Y or N indicator in the CI columns.

#### **CAPTURE\_NODE=node\_name**

Optional. The node name that the PowerExchange Logger uses to retrieve capture registrations and change data. Specify this parameter only if you use the PowerExchange Logger to capture change data from a source on a remote system.

Enter the node name of the remote source system, as specified in a NODE statement in the dbmover configuration file on the system where the PowerExchange Logger runs. The PowerExchange Logger uses the specified node name to connect to the PowerExchange Listener on the remote source node to read capture registrations and change data. The PowerExchange Logger then writes the change data to its local log files.

Default is local. Do not specify this parameter if the capture registrations and change data are on the local machine where the PowerExchange Logger runs.

You can also specify an optional user ID and password to control connection to the specified node. For more information, see the CAPTURE\_NODE\_UID parameter and the CAPTURE\_NODE\_EPWD or CAPTURE\_NODE\_PWD parameter.

#### **CAPTURE\_NODE\_EPWD=encrypted\_password**

An encrypted password that is associated with the user ID specified in the CAPTURE\_NODE\_UID parameter. This password, in conjunction with the CAPTURE\_NODE\_UID value, is used to control PowerExchange access to capture registrations and change data.

**Tip:** You can create an encrypted password in the PowerExchange Navigator by selecting **File > Encrypt Password**.

If you use remote logging of data from a data source on i5/OS or z/OS to a PowerExchange Logger for Linux, UNIX, and Windows instance, you can enter an encrypted PowerExchange passphrase instead of an encrypted password. Do not encrypt a passphrase that contains invalid characters, such as double-quotation marks, single quotation marks, or currency symbols.

**Note:** If you specify CAPTURE\_NODE\_EPWD, do not also specify CAPTURE\_NODE\_PWD.

#### **CAPTURE\_NODE\_PWD=password**

A clear text password that is associated with the user ID specified in the CAPTURE\_NODE\_UID parameter. This password, in conjunction with the CAPTURE\_NODE\_UID value, is used to control PowerExchange access to capture registrations and change data.

If you use remote logging of data from a data source on i5/OS or z/OS to a PowerExchange Logger for Linux, UNIX, and Windows instance, you can use a valid PowerExchange passphrase instead of a password. An i5/OS passphrase can be from 9 to 31 characters in length. A z/OS passphrase can be from 9 to 128 characters in length. A passphrase can contain the following characters:

- Uppercase and lowercase letters
- The numbers 0 to 9
- Spaces
- The following special characters:  
' - ; # \ , . / ! % & \* ( ) \_ + { } : @ | < > ?

**Note:** The first character is an apostrophe.

Passphrases cannot include single quotation marks ('), double quotation marks ("), or currency symbols.

When entering a passphrase, you must enclose it with double-quotation marks ("), for example:

```
CAPTURE_NODE_PWD="This is a passphrase!"
```

**Note:** On z/OS, a valid RACF passphrase can be up to 100 characters in length. PowerExchange truncates passphrases longer than 100 characters when passing them to RACF for validation.

To use passphrases, ensure that the PowerExchange Listener runs with a security setting of SECURITY=(1,N) or higher in the DBMOVER member. For more information, see "SECURITY Statement" in the *PowerExchange Reference Manual*.

**Note:** If you specify CAPTURE\_NODE\_PWD, do not also specify CAPTURE\_NODE\_EPWD.

#### **CAPTURE\_NODE\_UID=user\_id**

User ID that is used to control access to capture registrations and change data on the local machine or on the remote node that is specified in the CAPTURE\_NODE parameter.

Whether this parameter is required depends on the operating system of the local or remote node and the SECURITY setting in its DBMOVER configuration file on that node.

If CAPTURE\_NODE specifies a z/OS or i5/OS node that has a SECURITY setting of 0, do not specify this parameter. PowerExchange uses the user ID under which the PowerExchange Listener job runs to control access to capture registrations and change data.

If CAPTURE\_NODE specifies a z/OS or i5/OS node that has a SECURITY setting of 1, you must enter a valid operating system user ID for this parameter. Otherwise, error message PWX-00231 is issued, indicating a signon failure. However, PowerExchange uses the user ID under which the PowerExchange Listener job runs to control access to capture registrations and change data.

If CAPTURE\_NODE specifies a z/OS or i5/OS node that has a SECURITY setting of 2, you must enter a valid operating system user ID for this parameter. Otherwise, error message PWX-00231 is issued, indicating a signon failure. PowerExchange uses this user ID to control access to capture registrations and change data. If the specified user ID does not have the authority that is required to read capture registrations or change data, access fails.

For a Linux, UNIX, or Windows local or remote node, enter a user ID that is valid for your data source type:

- For DB2 for Linux, UNIX, or Windows sources, enter a valid operating system user ID that has DB2 DBADM or SYSADM authority.
- For Microsoft SQL Server instances that use SQL Server Authentication, enter a database user ID that permits access to the SQL Server distribution database. For SQL Server instances that use Windows Authentication, PowerExchange uses the user ID under which the PowerExchange Listener was started. In this case, do not specify this parameter unless you want to specify another user.
- For Oracle sources, if you use PowerExchange Oracle CDC with LogMiner, enter the ORACAPT user ID that you defined, which permits access to Oracle archive logs and to Oracle LogMiner. If you use PowerExchange Express CDC for Oracle, enter the ORACAPTL user ID that you defined, which permits access to the Oracle online and archive redo logs.

**Note:** If you specify CAPTURE\_NODE\_UID, you must enter a password or encrypted password in either the CAPTURE\_NODE\_PWD or CAPTURE\_NODE\_EPWD parameter, but not both.

#### **COLL\_END\_LOG={0|1}**

Required. The PowerExchange Logger operational mode.

Enter one of the following options:

- **0.** Runs the PowerExchange Logger continuously until you manually stop it. After the Writer subtask completes a processing cycle, it waits for the number of minutes specified in the NO\_DATA\_WAIT parameter before starting another processing cycle.
- **1.** Runs the PowerExchange Logger in batch mode. The PowerExchange Logger shuts down after the seconds specified in the NO\_DATA\_WAIT2 parameter elapse and no data has been received.

Default is **0** for continuous mode.

#### **COND\_CDCT\_RET\_P=days**

Recommended. Retention period, in days, for CDCT records and PowerExchange Logger log files. Log files that are older than this period and their corresponding CDCT records are deleted automatically during PowerExchange Logger cleanup processing. Cleanup processing occurs during startup, file switch, or shutdown processing.

Enter a number greater than 0. Default is 60.

When setting this parameter, try to minimize the size of the CDCT file while preserving the log files that contain the earliest change data that you might need to access. Use the following guidelines:

- If you set the retention period to a low value, ensure that PowerExchange extracts change data from the PowerExchange Logger log files during this period. Otherwise, the log files for which the retention period has elapsed are deleted and you can lose change data. For example, if the retention period is 5 and you plan not to run extractions during a 10-day holiday, increase the retention period to 15. This approach ensures that the log files with the change data you need are not deleted until the extractions run again.
- If you set the retention period to a high value, the CDCT can become very large, depending on the number active capture registrations. Also, the number of PowerExchange Logger log files might increase. For continuous extraction mode, you can use the PowerExchange Logger FILE\_SWITCH parameters to decrease the number of log files and increase their size instead.
- If you use continuous extraction mode, PowerExchange reads the CDCT file each time the interval specified in the FILEWAIT parameter of the CAPX CAPI\_CONNECTION statement elapses. If a CDCT file becomes large, this read activity can increase I/O, system resource use, and latency of change data extraction. If you use batch extraction mode, this high read activity is not a consideration.

#### **CONDENSENAME=service\_name**

Optional. A name for the command-handling service for a PowerExchange Logger for Linux, UNIX, and Windows process to which infacmd pwx or pwxcmd commands are issued.

Enter a service name up to 64 characters in length. No default is available.

This service name must match the service name that is specified in the associated SVCNODE statement in the dbmover configuration file. The SVCNODE statement specifies the TCP/IP port on which this service listens for infacmd pwx or pwxcmd commands.

**Tip:** If you run the PowerExchange Logger as a background process in continuous mode, specify this parameter so that you can use the pwxcmd program to issue commands to the PowerExchange Logger. Without the use of pwxcmd, you cannot shut down a PowerExchange Logger process that is running in the background or send status information to a computer that is remote from where the PowerExchange Logger runs.

#### **CONDENSE\_SHUTDOWN\_TIMEOUT=seconds**

Maximum amount of time, in seconds, that the PowerExchange Logger waits after receiving the SHUTDOWN or pwxcmd shutdown command before stopping.

Enter a number from 0 through 2147483647. Default is 600.

During a shutdown, the PowerExchange Logger updates the CDCT file for each capture registration that is used to capture change data. If you have a large number of capture registrations, you might need to increase this timeout period.

**CONN\_OVR=capi\_connection\_name**

Recommended. The name of the override CAPI\_CONNECTION statement to use for the PowerExchange Logger. If you do not enter CONN\_OVR, the PowerExchange Logger uses the default CAPI\_CONNECTION in the dbmover configuration file, if specified.

Enter a valid CAPI\_CONNECTION name for the source type.

Informatica recommends that you specify CONN\_OVR because it is the only type of override that the PowerExchange Logger can use.

**DBID=instance\_name**

Required. A source identifier, sometimes called the *instance* name, that is defined in capture registrations. When used with DB\_TYPE, it defines selection criteria for capture registrations in the CCT file.

This value matches the instance or database name that is displayed in the Resource Inspector of the PowerExchange Navigator for the registration group that contains the capture registrations.

For sources on Linux, UNIX, and Windows, enter one of the following options:

- For DB2 for Linux, UNIX, and Windows, enter the **Database** name that is displayed for the registration group in the **Resource Inspector**.
- For Microsoft SQL Server, this value depends on whether you also specify the optional DISTSRV and DISTDB parameters in the PowerExchange Logger configuration file and whether you entered the **Instance** identifier when creating the registration group:
  - If you specify the optional DISTSRV and DISTDB parameters to have the single PowerExchange Logger instance log data for all of the registered articles that are published to the distribution database, enter a name that serves as the collection identifier for all of the registrations. This name must be one to eight characters in length and start with a letter. This name overrides the instance name that is associated with the individual registrations.  
**Note:** When you perform a CAPXRT database row test, you must enter this name in the **MSS LUW DBId** field in the CAPXRT Advanced Parameters dialog box. When you define a PWXPC connection for PowerCenter CDC sessions that extract data from PowerExchange Logger log files, enter this value for the **Logger DBID** attribute on the PowerCenter PWX MSSQL CDC Real Time connection.
  - If you do not specify the DISTSRV and DISTDB parameters, enter the value from the **Instance** field that is displayed for the registration group in the PowerExchange Navigator Resource Inspector. The instance identifier is either the unique user-defined identifier that was optionally entered for the database name and database server combination during registration group creation in the PowerExchange Navigator or the instance identifier that PowerExchange generated if you did not specify an instance identifier.
- For Oracle, enter the **Instance** name that is displayed for the registration group in the **Resource Inspector**. This value also should match the first positional parameter of the ORACLEID statement in the dbmover configuration file.

If you use the PowerExchange Logger to log data from remote data sources on z/OS or i5/OS, enter one of the following options:

- For Adabas, enter the **Instance** name that is displayed for the registration group.
- For Datacom, enter the **MUF Name** value that is displayed for the registration group.



Alternatively, if you use Datacom synchronous CDC, enter the value of the MUF parameter in the DTLINPUT data set specified in the MUF JCL. Or, if you use Datacom table-based CDC, enter the value of REG\_MUF parameter in the ECCRDCMP member of the RUNLIB library.

- For DB2 for i5/OS, enter the **Instance** name that is displayed for the registration group. This name should match the INST parameter value in the AS4J CAPI\_CONNECTION statement in the DBMOVER member of the CFG file.
- For DB2 for z/OS, enter the **Instance** name that is displayed for the registration group. This name should match the RN parameter value in the DB2 statement in the RUNLIB(REPDB2OP) member.
- For IDMS Log-based CDC, enter the **Logsid** value that is displayed for the registration group. This value should match the LOGSID parameter value in the RUNLIB(ECCRIDLP) member.
- For IMS, enter the **IMSID** value that is displayed for the registration group. For IMS log-based CDC, this value should match the first parameter value in the IMSID statement in the RUNLIB(CAPTIMS) member.
- For VSAM, enter the **Instance** name that is displayed for the registration group.

**DB\_TYPE={ADA|AS4|DB2|DCM|IDL|IMS|MSS|ORA|UDB|VSM}**

Required. Source database type.

For sources on Linux, UNIX, and Windows, enter one of the following options:

- **MSS**. For Microsoft SQL Server sources.
- **ORA**. For Oracle sources.
- **UDB**. For DB2 for Linux, UNIX, and Windows sources.

If you use the PowerExchange Logger to log data from remote data sources on z/OS or i5/OS, enter one of the following options:

- **ADA**. For Adabas sources.
- **AS4**. For DB2 for i5/OS sources.
- **DB2**. For DB2 for z/OS sources.
- **DCM**. For Datacom sources.
- **IDL**. For IDMS log-based CDC sources.
- **IMS**. For IMS sources.
- **VSM**. For VSAM sources.

**DISTDB=distribution\_database\_name**

Optional for Microsoft SQL Server sources. The name of the distribution database. Enter this parameter if you want the single PowerExchange Logger instance to read data for all of the registered articles that are published to the distribution database, regardless of their source publication databases. Otherwise, you must run a separate PowerExchange Logger instance for each source database.

If you specify this parameter, you must also specify the DISTSRV parameter.

**Note:** For the PowerExchange Logger to extract change data from the distribution database for articles in multiple publication databases in one pass, you must also set the MULTIPUB parameter in the MSQ CAPI\_CONNECTION statement to Y. Otherwise, the extraction fails with message PWX-15757. Also, you must specify the DBID parameter value for the **Logger DBID** attribute on the PowerCenter PWX MSSQL CDC Real Time connection.

**DISTSRV=distribution\_database\_server**

Optional for Microsoft SQL Server sources. The network name of the server that hosts the distribution database. This name might be different from the network name of the SQL Server instance if the

distribution database resides on a different server. Enter this parameter if you want the single PowerExchange Logger instance to read data for all of the registered articles that are published to the distribution database, regardless of their source publication databases. Otherwise, you must run a separate PowerExchange Logger instance for each source database.

If you specify this parameter, you must also specify the DISTDB parameter.

**Note:** For the PowerExchange Logger to extract change data from the distribution database for articles in multiple publication databases in one pass, you must also set the MULTIPUB parameter in the MSQL CAPI\_CONNECTION statement to Y. Otherwise, the extraction fails with message PWX-15757. Also, you must specify the DBID parameter value for the **Logger DBID** attribute on the PowerCenter PWX MSSQL CDC Real Time connection.

#### **ENCRYPTPWD=encrypted\_encryption\_password**

A password in encrypted format for enabling the encryption of PowerExchange Logger log files. With this password, the PowerExchange Logger can generate a unique encryption key for each Logger log file. The password is stored in the CDCT file in encrypted format. For security purposes, the password is not stored in CDCT backup files and is not displayed in the CDCT reports that you can generate with the PWXUCDCT utility.

You can set the AES algorithm to use for log file encryption in the ENCRYPTOPT parameter. The default is AES128.

If you specify this parameter, do not also specify the ENCRYPTPWD parameter in the same pwxcl.cfg file.

If you specify this parameter and cold start the PowerExchange Logger with a pwxcl command that includes the encryptpwd parameter, the ENCRYPTPWD parameter in the configuration file takes precedence.

If you change this ENCRYPTPWD password after log files have been encrypted, you must cold start the PowerExchange Logger. Otherwise, the change is ignored.

**Tip:** For optimal security, Informatica recommends that you specify the encryption password in a pwxcl command for cold starting the PowerExchange Logger rather than in the pwxcl.cfg configuration file. This practice can reduce the risk of malicious access to the encryption password for the following reasons: 1) The encryption password is not stored in the pwxcl.cfg file, and 2) You can remove the password from the command line after a successful cold start. If you specify the encryption password in a pwxcl command for a cold start and then need to restore the CDCT file later, you must enter the same encryption password in the RESTORE\_CDCT command of the PWXUCDCT utility.

To *not* encrypt PowerExchange Logger log files, do not enter an encryption password in the pwxcl.cfg configuration file or in the pwxcl command for a cold start.

#### **ENCRYPTOPT={AES128|AES192|AES256}**

The AES encryption algorithm that you want to use for encrypting PowerExchange log files. To enable encryption, you must also specify an encryption password in the ENCRYPTPWD or ENCRYPTPWD parameter in the pwxcl.cfg configuration file or specify the encryptpwd parameter in a pwxcl command that you use to cold start the PowerExchange Logger.

Default is AES128.

#### **ENCRYPTPWD=clear\_text\_encryption\_password**

A password in clear text format for enabling the encryption of PowerExchange Logger log files. With this password, the PowerExchange Logger can generate a unique encryption key for each Logger log file. The password is stored in the CDCT file in encrypted format. For security purposes, the password is not included in CDCT backup files and is not displayed in the CDCT reports that you can generate with the PWXUCDCT utility.

You can set the AES algorithm to use for log file encryption in the ENCRYPTOPT parameter. The default is AES128.

If you specify this parameter, do not also specify the ENCRYPTPWD parameter in the same pwxcl.cfg file.

If you specify this parameter and cold start the PowerExchange Logger with a pwxcl command that includes the encryptpwd parameter, an error occurs. Do not specify the ENCRYPTPWD parameter in the configuration file and also specify the encryptpwd parameter in the PWXCCL command.

If you change this ENCRYPTPWD password after log files have been encrypted, you must cold start the PowerExchange Logger. Otherwise, the change is ignored.

**Tip:** For optimal security, Informatica recommends that you specify the encryption password in a pwxcl command for cold starting the PowerExchange Logger rather than in the pwxcl.cfg configuration file. This practice can reduce the risk of malicious access to the encryption password for the following reasons: 1) The encryption password is not stored in the pwxcl.cfg file, and 2) You can remove the password from the command line after a successful cold start. If you specify the encryption password in a pwxcl command for a cold start and then need to restore the CDCT file later, you must enter the same encryption password in the RESTORE\_CDCT command of the PWXUCDCT utility.

To *not* encrypt PowerExchange Logger log files, do not enter an encryption password in the pwxcl.cfg configuration file or in the pwxcl command for a cold start.

#### EPWD

A deprecated parameter. Use CAPTURE\_NODE\_EPWD instead. If both CAPTURE\_NODE\_EPWD and EPWD are specified, CAPTURE\_NODE\_EPWD takes precedence.

#### EXT\_CAPT\_MASK=path/prefix

Required. An existing directory path and a unique prefix to be used for generating the PowerExchange Logger log files.

Maximum length is 256 characters.

For example:

```
/capture/pwxlog
```

**Note:** If you enter a value that include spaces, you must enclose the value in double quotation marks (").

Verify that no existing files match this path and prefix. PowerExchange considers any file that matches this path and prefix to be a PowerExchange Logger log file, even if it is unrelated to PowerExchange Logger processing.

To create the log files, the PowerExchange Logger appends the following information:

```
.CND.CPyymmdd.Thhmmssnnn
```

Where:

- *yymmdd* is a date composed of a two-digit year, a month, and a day.
- *hhmmss* is 24-hour time value, including hours, minutes, seconds.
- *nnn* is a generated sequence number that ensures uniqueness of the file name.

For example:

```
/capture/pwxlog.CND.CP080718.T1545001
```

#### FILE\_FLUSH\_VAL=seconds

Recommended. File flush interval in seconds. When this interval elapses, the PowerExchange Logger writes any outstanding change data that it read from the source to log files on disk. After the change

data is flushed to disk, CDC sessions that use continuous extraction mode can read the change data. This parameter affects the latency of continuous change data extractions.

Valid values are:

- -1. The PowerExchange Logger does *not* flush outstanding change data to the current log file on disk based on this parameter. Enter this value only if you use batch extraction mode. If you use continuous extraction mode, this value can increase the latency of your continuous extraction sessions.
- 0. The PowerExchange Logger flushes outstanding change data after every record.
- Any value from 1 through 86400. The PowerExchange Logger flushes outstanding change data at the specified interval.

Default is -1.

**Warning:** A value of 0 can degrade the performance of the PowerExchange Logger and file system.

Set this value as appropriate for your CDC environment. Values that are too high can increase change extraction latency, and values that are too low can degrade PowerExchange Logger and system performance. Informatica recommends that you set this parameter to a value that is equal to or greater than the NO\_DATA\_WAIT2 value because file flushes cannot occur until the NO\_DATA\_WAIT2 period expires.

#### **FILE\_SWITCH\_CRIT={M|R}**

Type of units to use for the FILE\_SWITCH\_VAL parameter, which determines when to do an automatic file switch.

Enter one of the following options:

- **M** for minutes.
- **R** for records.

Default is **M**.

#### **FILE\_SWITCH\_VAL**

Number of minutes or change records, as determined by FILE\_SWITCH\_CRIT, that must elapse before PowerExchange performs a file switch.

Enter a number greater than 0. Default is 30.

For example, if you use 30 and enter FILE\_SWITCH\_CRIT=R, the PowerExchange Logger performs a file switch every 30 records. If FILE\_SWITCH\_CRIT=M, the PowerExchange Logger performs a file switch every 30 minutes.

If the PowerExchange Logger log files contain no data when the FILE\_SWITCH\_VAL threshold is reached, the file switch does not occur.

This value affects the size of the PowerExchange Logger log files. Specify a value that results in log files of the appropriate size for your environment.

**Tip:** When using continuous extraction mode, set this parameter such that you have larger log files and a smaller CDCT file. When using batch extraction mode, set this parameter to a value that causes file switches to occur within the time frame that meets your change extraction latency requirements.

#### **GROUPDEFS=path/file\_name**

Path and file name of the optional PowerExchange Logger group definition file. This file defines groups of capture registrations that the PowerExchange Logger uses to capture change data to separate sets of log files. It also defines the path that the PowerExchange Logger uses to create the log files that contain the change data for each group.

This parameter is optional and has no default.

Maximum length is 255 characters.

#### **LOGGER\_DELETES\_EXPIRED\_CDCT\_RECORDS={Y|N}**

Controls whether the PowerExchange Logger deletes CDCT records for log files for which the retention period has expired.

Enter one of the following options:

- **Y.** The PowerExchange Logger deletes expired CDCT records during file switches. You cannot use the PWXUCDCT utility DELETE\_EXPIRED\_FILES command to manually delete expired log files and their related CDCT records.
- **N.** The PowerExchange Logger does not delete expired CDCT records. However, you can use the PWXUCDCT utility DELETE\_EXPIRED\_FILES command to manually delete expired log files and their related CDCT records.

**Note:** This parameter does not affect PowerExchange Logger deletions of CDCT records that are rolled back because of a cold start or a warm start to a prior point in time.

Default is **Y**.

#### **NO\_DATA\_WAIT=minutes**

If you run the PowerExchange Logger in continuous mode, specify the number of minutes that the PowerExchange Logger must wait before starting the next logging cycle.

Enter 0 or a number greater than 0. Default is 1.

A value of 0 causes no waiting to occur between PowerExchange Logger processing cycles. If source data is not available, the CAPI sleeps.

For continuous extraction mode, enter 0 for no waiting, or enter a low value so that the next logging cycle starts shortly after the current one completes.

If the value of FILE\_SWITCH\_CRIT is M and the value of FILE\_SWITCH\_VAL is less than the value of NO\_DATA\_WAIT, the PowerExchange Logger uses the FILE\_SWITCH\_VAL value instead.

#### **NO\_DATA\_WAIT2=seconds**

The number of seconds that PowerExchange waits at the end-of-log for more change data before returning control to the PowerExchange Logger. If this wait period elapses and no new change data has been received, PowerExchange returns control to the PowerExchange Logger, and the PowerExchange Logger then stops the current logging cycle.

Enter a number greater than 0. Default is 600.

Informatica recommends a value of 10. If you enter a greater value, execution of commands for the PowerExchange Logger might be delayed.

Use the same value for the FILE\_FLUSH\_VAL parameter.

#### **PROMPT={Y|N}**

When you run the PowerExchange Logger in foreground mode, controls whether PowerExchange displays a user confirmation prompt and waits for a response when you perform one of the following actions:

- Cold start the PowerExchange Logger.
- Warm start the PowerExchange Logger from a previous position in the change stream. This situation occurs only if the CDCT file still contains records related to the deleted files.

Enter one of the following options:

- **Y.** Displays the confirmation message PWX-33236 for a cold start or PWX-33242 for a warm start. You must respond to the message for startup processing to continue.
- **N.** Does not display the confirmation messages. PowerExchange attempts to start without first prompting for user confirmation.

If you run the PowerExchange Logger in foreground mode, the default is Y.

If you run the PowerExchange Logger in background mode or as a PowerExchange Logger Service in the Informatica domain, the default is N. In this case, if you enter PROMPT=Y in the pwxocl.cfg file, the PowerExchange Logger ignores this setting, issues error message PWX-33253, and continues processing.

#### **PWD**

A deprecated parameter. Use CAPTURE\_NODE\_PWD instead. If both CAPTURE\_NODE\_PWD and PWD are specified, CAPTURE\_NODE\_PWD takes precedence.

#### **RESTART\_TOKEN and SEQUENCE\_TOKEN**

A pair of token values that define a restart point for starting change data processing when a PowerExchange Logger is cold started.

Depending on how you set these parameters, PowerExchange Logger processing starts from one of the following restart points during a cold start:

- If you do not specify these parameters, processing starts from the current end-of-log position.
- If you enter 0 for both parameters, processing starts from the default start location:
  - For DB2, the default location is the current log position at the time the PowerExchange capture catalog was created.
  - For Oracle, if you use PowerExchange Oracle CDC with LogMiner, the default location is the most current Oracle catalog dump. If you use PowerExchange Express CDC for Oracle, the default location is the beginning of the most recent archive log.
  - For Microsoft SQL Server, the default location is the oldest data available in the publication database.
- If you enter restart token and sequence token values other than 0, processing resumes from the specific restart point defined by these token values. To perform a special start of the PowerExchange Logger, you must specify the RESTART\_TOKEN and SEQUENCE\_TOKEN parameters with specific values, and the SEQUENCE\_TOKEN value must be greater than the sequence token in the CDCT file.

If you use remote logging of change data from z/OS or i5/OS data sources, see the PowerExchange Condense chapter in the *PowerExchange CDC Guide for i5/OS* or *PowerExchange CDC Guide for z/OS* for information about what to enter for these parameters.

#### **SIGNALLING={N|Y}**

Indicates whether the PowerExchange Logger attempts to take automatic action in the event of certain errors.

Enter one of the following options:

- **N.** The PowerExchange Logger does not automatically trap and handle system errors. Instead, the operating system uses default error handling. Usually, the default handling is to report the program line in error and dump memory.
- **Y.** The PowerExchange Logger automatically handles certain errors such as memory corruption. After the PowerExchange Logger handles the error, it attempts to shut down in a controlled manner.

Default is **N**.

**STATS=(MONITOR[,interval|0])**

Enables PowerExchange Logger collection of the following monitoring statistics:

- PowerExchange Logger process ID (PID)
- Status of the PowerExchange Logger Writer task
- CPU time used by the PowerExchange Logger
- Memory use (current/total/maximum) in kilobytes, total and for the Controller, Command Handler, and Writer tasks
- The number of inserts, updates, deletes, and commits that the PowerExchange Logger processed, total since the Logger started and for the current active log file and the active logging cycle

Also enables collection of the following statistics for PowerExchange Logger group definitions, if defined:

- The number of DML operations and commits processed for each group
- Then number of change records that have not yet been flushed to a Logger log file on disk
- The name of the open Logger log file for each group and the file open timestamp

The statistics are printed on demand to the PowerExchange message log and on screen when you enter the DL or DG command from the command line or the `pwxcmd displaystats -tp {logger|groups}` command from a remote system.

For more information about the monitoring commands and related reports, see the *PowerExchange Command Reference*.

Optionally, include the *interval* subparameter in the STATS statement to publish the Logger statistics at a regular interval.

**{interval|0}**

Optional. The interval, in minutes, after which PowerExchange publishes monitoring statistics for the PowerExchange Logger. The interval-based statistics that are written to the PowerExchange message log file are the same as those published by the DL (or DS) command and `pwxcmd displaystats -tp logger` command. However, a subset of the message output is displayed on screen to prevent flooding the screen with messages over time.

Valid values are 0 through 120. Default is 0, which disables interval-based reporting of PowerExchange Logger monitoring statistics. With the default value, PowerExchange writes these statistics only when a DL or `pwxcmd displaystats -tp logger` command is issued.

**UID**

A deprecated parameter. Use `CAPTURE_NODE_UID` instead. If both `CAPTURE_NODE_UID` and `UID` are specified, `CAPTURE_NODE_UID` takes precedence.

**VERBOSE={Y|N}**

Indicates whether the PowerExchange Logger writes verbose or terse messages to the PowerExchange message log file for activities that it performs frequently, such as cleanup, condense, and file-switch processing.

Enter one of the following options:

- **Y**. Verbose messaging. The PowerExchange Logger logs multiple messages at various processing points, such as when starting or ending a cycle of reading source data or doing a file switch. Verbose messaging often includes processing statistics such as records processed and elapsed time.

- **N.** Terse messaging. The PowerExchange Logger logs a single terse message for each file switch.

Default is **Y**.

## Example pwxcl Configuration File

PowerExchange provides an example PowerExchange Logger configuration file, pwxcl, in the PowerExchange installation directory. You can use this file to create a custom configuration file.

The following example shows basic configuration statements:

```
/* Name for PWXCMD control
/*CONDENSENAME=PWXCCL1

DBID=ORACOLL1
DB_TYPE=ORA
CAPTURE_NODE_UID=user_id
CAPTURE_NODE_EPWD=encrypted_password
/* CAPTURE_NODE_PWD=plain_text_password

PROMPT=Y

EXT_CAPT_MASK=/capture/condense0
COND_CDCT_RET_P=50
LOGGER_DELETES_EXPIRED_CDCT_RECORDS=Y

/* 0 = continuous, 1 = Stop at end-of-log (batch)
COLL_END_LOG=0

/* Number of minutes to wait between CAPI read cycles
NO_DATA_WAIT=0
/* Number of seconds to wait at the end-of-log for more change data
NO_DATA_WAIT2=60

/* Number of seconds before flushing, or writing, data to the current log file on disk
/* -1 = No flush, 0 = flush every record, 1 to N flush every N seconds
/*FILE_FLUSH_VAL=60
/* Minimum number of FILE_SWITCH_CRIT units after new CDCT source entry
(normal,coldstart)
FILE_SWITCH_CRIT=M
FILE_SWITCH_VAL=20

CAPT_IMAGE=BA
SEQUENCE_TOKEN=00
RESTART_TOKEN=00
```

## Customizing the dbmover Configuration File for the PowerExchange Logger

To use the PowerExchange Logger, you must define the CAPT\_PATH statement and certain source-specific statements in the dbmover configuration file.

Also, you can include some optional parameters to help make finding messages for the PowerExchange Logger easier or to send commands to a PowerExchange Logger process that is running in background mode.

Use the following key parameters:

### **CAPT\_PATH**

Required. Path to the local directory on a Linux, UNIX, and Windows system that contains the control files for CDC, including the CCT and CDCT files. The CCT file contains information about capture registrations. The CDCT file contains information about the PowerExchange Logger log files, such as file names and number of records.



## CAPX CAPI\_CONNECTION

Required for continuous extraction mode. If you want the CAPI to use continuous extraction mode for the extraction of change data from PowerExchange Logger log files, you must define a CAPX CAPI\_CONNECTION statement.

## LOGPATH

Optional. A unique path and directory for PowerExchange message log files on a Linux, UNIX, or Windows system. Use this parameter to create message log files in a directory that is separate from your current working directory so that you can find the message log files more easily.

## SVCNODE

Optional. The TCP/IP port on which a command-handling service for a PowerExchange Logger process listens for commands that you issue with the pwxcmd program. You must define this parameter if you run the PowerExchange Logger process in background mode on a Linux or UNIX system. For more information about pwxcmd commands, see the *PowerExchange Command Reference*.

## TRACING

Optional. Enables alternative logging. PowerExchange creates a set of alternative log files for each PowerExchange process in a separate directory. You can specify the directory location, the number of log files, and the log file size in MB. When a log file reaches the specified size, PowerExchange switches to the next log file and begins overwriting any data in that file. Alternative logging is faster and enables you to customize the amount of data logged for long-running jobs, such as a PowerExchange Logger process that runs in continuous mode. If you specify this statement, also specify the LOGPATH statement.

In addition to these parameters, the PowerExchange Logger requires source-specific statements. For example, for PowerExchange Oracle CDC with LogMiner sources, you must define ORCL CAPI\_CONNECTION, UOWC CAPI\_CONNECTION, and ORACLEID statements.

For more information about all DBMOVER configuration parameters, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“DB2 for Linux, UNIX, and Windows CDC” on page 77](#)
- [“Microsoft SQL Server CDC” on page 94](#)
- [“Oracle CDC with LogMiner” on page 109](#)
- [“Express CDC for Oracle” on page 139](#)

## CAPX CAPI\_CONNECTION - CAPX Statement

The CAPX CAPI\_CONNECTION statement specifies a named set of parameters that the Consumer API (CAPI) uses for continuous extraction of change data from PowerExchange Logger for Linux, UNIX, and Windows log files.

**Operating Systems:** Linux, UNIX, and Windows

**Required:** Yes for continuous extraction mode

**Syntax:**

```
CAPI_CONNECTION= ([DLLTRACE=trace_id]
                  ,NAME=capi_connection_name
                  [,TRACE=trace_name]
                  ,TYPE=(CAPX
                        ,DFLTINST=instance_name
                        [,FILEWAIT={seconds|1}])
```

```

        [, NOSEQVAL={N|Y}]
        [, RSTRADV=seconds]
        [, RSTRANMODE={N|Y}]
    )

```

**Parameters:**

**DLLTRACE=*trace\_id***

Optional. User-defined name of the TRACE statement that activates internal DLL tracing for this CAPI. Specify this parameter only at the direction of Informatica Global Customer Support.

**NAME=*capi\_connection\_name***

Required. Unique user-defined name for this CAPI\_CONNECTION statement.

Maximum length is eight alphanumeric characters.

**TRACE=*trace\_name***

Optional. User-defined name of the TRACE statement that activates the common CAPI tracing. Specify this parameter only at the direction of Informatica Global Customer Support.

**TYPE=(CAPX, ... )**

Required. Type of CAPI\_CONNECTION statement. For continuous extraction mode, this value must be CAPX.

**DFLTINST=*instance\_name***

Required. A source instance identifier that is specified for the registration group. This value must match the DBID value that is specified in the PowerExchange Logger configuration file.

To determine this value, view the registration group properties in the PowerExchange Navigator. Depending on the source type, enter one of the following values:

- For Adabas, DB2 for i5/OS, DB2 for z/OS, Oracle, and VSAM, the name that is displayed in the **Instance** field for the registration group.
- For Microsoft SQL Server, this value depends on whether you specify the optional DISTSRV and DISTDB parameters in the PowerExchange Logger configuration file:
  - If you specify the DISTSRV and DISTDB parameters, enter the DBID name that you use as the collection identifier for all of the registrations. This name overrides the instance name that is associated with the individual registrations.
  - If you do not specify the DISTSRV and DISTDB parameters, enter the value that the PowerExchange Navigator generates and displays in the **Instance** field of the **Resource Inspector** for the registration group. The generated value is composed of the first four characters of the database name followed by a generated number, which starts at 000.
- For Datacom, the name of the Multi-User Facility (MUF) in the **MUF Name** field.
- For a DB2 for Linux, UNIX, and Windows source, the name of the database in the **Database** field.
- For an IDMS log-based source, the name of the database in the **DB Name** field.
- For an IMS source, the recon identifier for the database in the **RECON ID** field.

Maximum length is eight alphanumeric characters.

**FILEWAIT={*seconds*|1}**

Optional. The number of seconds that PowerExchange waits before checking for new PowerExchange Logger log files.

For the *seconds* variable, enter a number from 1 through 86400. Default is 1.

#### **NOSEQVAL={N|Y}**

If you receive error message PWX-36944 after starting a CDC session, the sequence token that PWXPC passed to PowerExchange is earlier than the sequence token that is recorded in the PowerExchange Logger CDCT file. If you want the session to continue and you can tolerate some data loss, you can set this parameter to Y. The Log Reader then begins extracting the earliest available data in the log files. With the default value of N, the session fails.

#### **RSTRADV=seconds**

Time interval, in seconds, that PowerExchange waits before advancing the restart and sequence tokens for a registered data source during periods when UOWs contain no changes of CDC interest for a data source. When the wait interval expires, PowerExchange returns the next committed "empty UOW," which includes only updated restart information.

For the *seconds* variable, enter a number from 0 through 86400. No default value is provided. A value of 0 causes PowerExchange to return an empty UOW after each UOW processed. Consequently, a value of 0 can degrade performance.

PowerExchange resets the wait interval to 0 when one of the following events occur:

- PowerExchange completes processing a UOW that includes changes of CDC interest.
- PowerExchange returns an "empty UOW" because the RSTRADV wait interval expired without any new changes of CDC interest having been received.

For example, if you specify 5, PowerExchange waits 5 seconds after it completes processing the last UOW or after the previous RSTRADV interval expires. Then PowerExchange returns the next committed "empty UOW" that includes the updated restart information and resets the wait interval to 0.

If you specify the RSTRADV parameter with any valid value, PowerExchange always advances the restart and sequence tokens when the Log Reader reaches the end of a Logger log file, even if the RSTRADV interval has not expired. This behavior ensures that restart and sequence tokens are advanced even when the CDC session run time is shorter than the RSTRADV interval. This situation is most likely to occur if you have source tables that have a low level of update activity.

If you do not specify RSTRADV, PowerExchange does not advance restart and sequence tokens for a registered source during periods when PowerExchange receives no changes of interest. In this case, when PowerExchange warm starts, it reads all changes, even those not of CDC interest, from the restart point.

For DB2 for i5/OS sources, Informatica recommends that you use this parameter if the change records that PowerExchange reads from i5/OS journal receivers are created under commitment control. If the change records are created without commitment control, do not specify this parameter.

#### **RSTRANMODE={N|Y}**

If you are migrating from real-time extraction mode to continuous extraction mode for CDC sessions that have not previously used the PowerExchange Logger for Linux, UNIX, and Windows, set this parameter to Y. This setting enables PowerExchange to convert restart token information to the format that PowerExchange Logger requires, when CDC sessions run. Retain the Y setting until PowerExchange completes converting the restart tokens for all registered source tables and all CDC sessions that use this CAPX CAPI\_CONNECTION. Then, set this parameter to N or delete it. Default is N.

## Using PowerExchange Logger Group Definitions

To create separate sets of PowerExchange Logger log files for groups of tables, create a PowerExchange Logger group definition file. Then, specify its path and file name in the GROUPDEFS parameter of the pwxccf configuration file.

When the PowerExchange Logger process starts, it reads the group definition file and creates a separate set of log files for each defined group.

Group definitions can help improve the efficiency of extraction sessions because the extractions target a more specific set of PowerExchange Logger log files.

By default, the PowerExchange Logger processes change data for all tables that reside on the instance specified by the DBID parameter and that have active capture registrations with the **Condense** option set to **Part**. Changes for all of these tables are written to a single set of log files. For a table with a low level of change activity, the extraction process might need to read many change records in the PowerExchange Logger log files before finding the changes of interest.

With group definitions, you can define a group that includes a subset of capture registrations. The PowerExchange Logger then writes change data to a separate set of log files for the tables that are associated with these registrations. When an extraction process runs, it is more likely to find the change data for a table in the group faster because it reads only the log files for that group.

For example, if you have five source tables with a low level of change activity and one table with a high level of change activity, you can define a group that includes the low-activity tables and another group that includes only the high-activity table only. Then, in PowerCenter, define a CDC session that extracts change data from the PowerExchange Logger log files for the low-activity group, and define another CDC session that extracts change from the log files for the high-activity group. This configuration enables the CDC session for the low-activity tables to find and extract the few change records for these tables much more quickly.

If you have multiple tables with the same table name but different schemas, you can define a single capture registration for the table and specify it once, under a single group, in the group definition file. For any other group that includes the same table with a different schema, you can override the schema name in the group definition by using a SCHEMA statement. By using the SCHEMA statement, you can avoid creating multiple capture registrations and specifying each one in the group definition file. For example, if you have an EMPLOYEE table with different schemas for the north, south, east, and west regions, you can register the north EMPLOYEE table only and specify the capture registration name in the NORTH group. Then specify only the override schemas in the EAST, WEST, and SOUTH groups.

**Note:** SCHEMA statements are optional for DB2 for i5/OS sources and for DB2 and Oracle sources on Linux, UNIX, and Windows. SCHEMA statements are not supported for SQL Server sources on Windows or any data source on z/OS.

On Linux, UNIX, and Windows, PowerExchange requirements for unregistered versions of tables, those for which you specify a SCHEMA statement instead of a REG statement in the group definition file, vary by source type:

- For DB2 for Linux, UNIX, and Windows, you must define any unregistered version of a table with the DB2 DATA CAPTURE CHANGES clause.
- For Oracle, you must create an Oracle supplemental log group for the unregistered table, which is similar to the supplemental log group that was created for the registered copy of the table at registration completion.
- For Microsoft SQL Server, you must register all versions of a table in PowerExchange and specify a REG statement in the group definition file.

**Tip:** When using group definitions, you can optimize extraction efficiency by defining a CDC session in PowerCenter for each group of tables defined in the group definition file.

## RELATED TOPICS:

- [“PowerExchange Logger Group Definition File” on page 61](#)
- [“Example Group Definition File” on page 62](#)

## PowerExchange Logger Group Definition File

A PowerExchange Logger group definition file contains one or more GROUP statements. Each GROUP statement contains REG or SCHEMA parameters that directly or indirectly identify a group of capture registrations and tables for which you want to create separate sets of PowerExchange Logger log files.

For the PowerExchange Logger to use the group definition file, you must specify the path and file name of the file in the GROUPDEFS parameter of the pwxcl.cfg file.

**Note:** If you specify the GROUPDEFS parameter, the PowerExchange Logger ignores the EXT\_CAPT\_MASK parameter in the pwxcl.cfg file when creating log files.

The following table describes the statements and parameters in the group definition file:

Statement	Positional Parameter	Description	Data Type and Length
GROUP	<i>group_name</i>	A unique user-defined name for the group. This parameter is required.	VARCHAR(255)
	<i>external_capture_mask</i>	A unique path and file-name prefix for the PowerExchange Logger log files that are created for tables in the group. This parameter is required. <b>Note:</b> This path and prefix is used for the group instead of the path and prefix that are specified in the EXT_CAPT_MASK parameter of the pwxcl.cfg file.	VARCHAR(255)
REG	<i>registration_name</i>	Optional. Registration name that is specified in the <b>Name</b> field of a capture registration. This lowercase name can be the full registration name or the first part of the name followed by an asterisk (*) wildcard. This parameter is optional. If omitted, the PowerExchange Logger assumes REG=.*.	VARCHAR(8)
SCHEMA	<i>schema_name</i>	Optional. Name of the override schema. You can optionally use this parameter for DB2 for i5/OS sources and for DB2 and Oracle sources on Linux, UNIX, and Windows. <b>Note:</b> This parameter is not supported for SQL Server sources on Windows. If you use the offloading feature to have the PowerExchange Logger process data from z/OS sources, this parameter is also not supported for the z/OS sources.	VARCHAR(255)

Use the following rules and guidelines when you create a PowerExchange Logger group definition file:

- Each *group\_name* must be unique within the group definition file.
- Each *external\_capture\_mask* must be unique on the system.

- SCHEMA statements are optional for DB2 for i5/OS sources and for DB2 and Oracle sources on Linux, UNIX, and Windows. SCHEMA statements are not supported for SQL Server sources on Windows or any data source on z/OS.
- If you use a SCHEMA statement, you must define a capture registration in the group. You can specify multiple SCHEMA statements under a GROUP if you want the tables with those schemas to be included in the group.
- REG statements apply to the preceding SCHEMA statement. If a SCHEMA statement is not present, the REG statements apply to the preceding GROUP statement.
- If the file contains a SCHEMA or REG statement without a preceding GROUP statement, the PowerExchange Logger issues a syntax error.
- Do not include the same *schema.table* value in more than one group. If a table is included in multiple groups, only the first group that includes the table logs changes for it.
- If you do not define at least one REG statement for a GROUP, the PowerExchange Logger includes all of the active capture registrations that are defined for the specified DBID instance and for which the **Condense** option is set to **Part**.
- If a registration belongs to multiple groups, the PowerExchange Logger logs changes for that registration only under the first group in the group definition file that includes the registration.

## Example Group Definition File

PowerExchange provides an example group definition file, `pwxcclgrp.cfg`, in the PowerExchange installation directory. Use this example as a starting point when creating your group definition file.

The example file contains the following statements:

```
GROUP=(Company1People,"/user/logger_files/people/company1/condense")
REG=Emp*
REG=Manager
GROUP=(UK_People,"/user/logger_files/people/UK/condense")
SCHEMA=Company2
REG=Manager
REG=Emp*
REG=Em*
SCHEMA=Company3
REG=Manager
REG=Emp*
GROUP=(All_Managers,"/user/logger_files/people/managers/condense")
SCHEMA=Company1
REG=Manager
SCHEMA=Company2
REG=Manager
SCHEMA=Company3
REG=Manager
GROUP=(AllCompany3_Locations,"/user/logger_files/locations/company3/condense")
REG=loc*
GROUP=(Company2Jobs,"/user/logger_files/jobs/company2/condense")
REG=Job*
```

**Note:** Because this example is for a group definition file on a Linux or UNIX system, the paths include forward slashes. A group definition file on Windows system would be similar but have back slashes.

This example file defines the following groups:

- **Company1People group.** Groups all tables associated with capture registrations that have names beginning with “Emp” or the name “Manager.” Changes for these tables are logged to log files that have file names beginning with “condense” and that are located at “/user/logger\_files/people/company1/.”
- **UK\_People group.** Groups tables that have the schema Company2 and that are associated with capture registrations that have names beginning with “Emp” or “Em” or the name “Manager.” Changes for these

tables are logged to log files that have names beginning with "condense" and that are located at "/user/logger\_files/people/UK/."

- **All\_Managers group.** Groups tables that have the schema Company1, Company2, or Company3 and that are associated with the capture registration with the name "Manager." Changes for these tables are logged to log files that have names beginning with "condense" and that are located at "/user/logger\_files/people/managers/."
- **AllCompany3\_Locations group.** Groups all tables that are associated with capture registrations that have names beginning with "loc." Changes for these tables are logged to log files that have names beginning with "condense" and that are located at "/user/logger\_files/locations/company3/."
- **Company2Jobs group.** Groups all tables that are associated with capture registrations that have names beginning with "Job." Changes for these tables are logged to log files that have names beginning with "condense" and that are located at "/user/logger\_files/jobs/company2/."

Some tables might be included in more than one group. For example, the table COMPANY2.MANAGERS is in the Company1People, UK\_People, and All\_Managers groups. However, changes for this table are logged only under the Company1People group because it is the first group in the file that includes this table.

## Starting the PowerExchange Logger

You can cold start, warm start, or special start a PowerExchange Logger process.

You must start the PowerExchange Logger under a user ID that has READ and WRITE access to PowerExchange Logger log files. Also, the PowerExchange Listener must be running under a user ID that READ access to the PowerExchange Logger log files.

Choose the start method that is appropriate for your current situation:

- A *cold start* uses the restart and sequence tokens, if present, in the pwxcl configuration file to determine the point in the change stream from which the PowerExchange Logger starts reading changes. If you are starting the PowerExchange Logger for the first time, you must perform a cold start. When you enter the pwxcl command to start the PowerExchange Logger, set the coldstart parameter to Y.
- A *warm start* uses the restart and sequence tokens in the CDCT file to resume CDC processing. You can perform a warm start only if you have run the PowerExchange Logger previously and have a recent CDCT file for the database instance. When you enter the pwxcl command for starting the PowerExchange Logger, set the coldstart parameter to N or omit the parameter.
- A *special start* uses the restart and sequence tokens in the pwxcl.cfg file to override the token values from the CDCT file for the PowerExchange Logger run. None of the data that was captured prior to the special start is lost. You must enter the SEQUENCE\_TOKEN and RESTART\_TOKEN parameters in the pwxcl.cfg. The SEQUENCE\_TOKEN value must be greater than or equal to the sequence token in the CDCT file.

Use a special start to avoid capturing changes from problematic portions of the logs. For example, perform a special start in the following situations:

- You do not want the PowerExchange Logger to capture an upgrade of an Oracle catalog. In this case, stop the PowerExchange Logger before the upgrade. After the upgrade is complete, generate new sequence and restart tokens for the PowerExchange Logger based on the post-upgrade SCN. To perform a special start, you must specify the SEQUENCE\_TOKEN and RESTART\_TOKEN parameters in the pwxcl.cfg file. Then special start the PowerExchange Logger.

- You do not want the PowerExchange Logger to reprocess old, unavailable logs that were caused by outstanding UOWs that are not of CDC interest. In this case, stop the PowerExchange Logger. Edit the RESTART\_TOKEN value to reflect the SCN of the earliest available log, and then perform a special start. If any of the outstanding UOWs are of CDC interest, data might be lost.

**Note:** You cannot use the pwxcmd or infacmd program to start the PowerExchange Logger.

## RELATED TOPICS:

- [“How the PowerExchange Logger Determines the Start Point for a Cold Start” on page 67](#)
- [“Cold Starting the PowerExchange Logger ” on page 67](#)
- [“PWXCL Command Parameters” on page 65](#)

## PWXCL Syntax and Parameters

To start the PowerExchange Logger process, run the pwxcl program, which is located in the PowerExchange installation directory by default.

### PWXCL Command Syntax

The pwxcl command that is used to start the PowerExchange Logger has the following syntax:

```
pwxcl
  [coldstart={Y|N}]
  [specialstart={Y|N}]
  [config=path/pwx_config_file]
  [cs=path/pwxlogger_config_file]
  [license=path/license_file]
  [encryptpwd=encrypted_password]
```

Use the following rules and guidelines when you enter the pwxcl command:

- To cold start the PowerExchange Logger, set the coldstart parameter to Y. The default is N.
- To special start the PowerExchange Logger from a specific point in the change stream, set the specialstart parameter to Y. The default is N. You must also specify the SEQUENCE\_TOKEN and RESTART\_TOKEN parameters in the pwxcl.cfg file.
- All parameters on the pwxcl command are optional. However, if you specify the config or license parameter, the cs parameter is required.
- In the config, cs, and license parameters, the full path is required only if the file is not in the default location.
- On Linux and UNIX, append an ampersand (&) at the end of the statement to run the PowerExchange Logger in background mode. For example:

```
pwxcl [coldstart=Y|N] [specialstart={Y|N}] [config=directory/myconfig_file]
[cs=directory/mycondense_config_file]
[license=directory/mylicense_key_file] &
```

For more information about pwxcl syntax, see the *PowerExchange Command Reference*.

**Caution:** If you run PowerExchange and PowerCenter on the same machine and with the same user account, you must create separate environments for PowerExchange and PowerCenter. To create the appropriate PowerExchange environment and start the PowerExchange Logger, run the pwxsettask.bat script on Windows or the pwxsettask.sh script on Linux or UNIX.

On Windows, use the following syntax:

```
pwxsettask pwxcl
  ["coldstart={Y|N}"]
```



```
["config=path/pwx_config_file"]
["cs=path/pwxlogger_config_file"]
["license=path/license_file"]
["encryptepwd=encrypted_password"]
```

The quotation marks are required on Windows.

On Linux and UNIX, use the following syntax:

```
pwxsettask.sh pwxcccl
[coldstart={Y|N}]
[config=path/pwx_config_file]
[cs=path/pwxlogger_config_file]
[license=path/license_file]
[encryptepwd=encrypted_password]
```

The quotation marks are optional on Linux and UNIX.

For more information, see [“Environment Variable Incompatibilities Between PowerExchange and PowerCenter” on page 23](#).

## PWXCCCL Command Parameters

You can specify several optional parameters in the pwxcccl command that starts the PowerExchange Logger.

The following table describes each parameter:

Parameter	Description
coldstart	<p>Indicates whether to cold start or warm start the PowerExchange Logger.</p> <p>Enter one of the following values:</p> <ul style="list-style-type: none"> <li>- <b>Y</b>. Cold start the PowerExchange Logger. You must specify coldstart=Y to perform a cold start. If the CDCT file contains log entries, the PowerExchange Logger deletes these entries.</li> <li>- <b>N</b>. Warm start the PowerExchange Logger from the restart point that is indicated in the CDCT file. If no restart information exists in the CDCT file, the PowerExchange Logger ends with error message PWX-33239.</li> </ul> <p>Default is <b>N</b>.</p>
config	<p>The full path and file name for a dbmover configuration file that overrides the default dbmover configuration file in the installation directory. The override file must have a path or file name that is different from that of the default file.</p> <p>This override file takes precedence over any other override configuration file that you optionally specify with the PWX_CONFIG environment variable.</p>
cs	<p>The full path and file name for the PowerExchange Logger configuration file. Use this parameter to specify a PowerExchange Logger configuration file that overrides the default pwxcccl configuration file in the installation directory. The override file must have a path or file name that is different from that of the default file.</p>

Parameter	Description
encryptepwd	<p>A password in encrypted format for enabling the encryption of PowerExchange Logger log files. With this password, the PowerExchange Logger can generate a unique encryption key for each Logger log file. The password is stored in the CDCT file in encrypted format. For security purposes, the password is not stored in CDCT backup files and is not displayed in the CDCT reports that you can generate with the PWXUCDCT utility. You can generate an encrypted password from the PowerExchange Navigator.</p> <p>If you specify this parameter, you must also specify coldstart=Y in the same pwxcl command.</p> <p>If you specify this command-line parameter and the ENCRYPTPWD parameter in the PowerExchange Logger configuration file, the parameter in the configuration file takes precedence. If you specify this command-line parameter and the ENCRYPTPWD parameter in the PowerExchange Logger configuration file, an error occurs.</p> <p>You can set the AES algorithm to use for log file encryption in the ENCRYPTOPT parameter of the pwxcl.cfg file. The default is AES128.</p> <p><b>Tip:</b> For optimal security, Informatica recommends that you specify the encryption password in a pwxcl command for cold starting the PowerExchange Logger rather than in the pwxcl.cfg configuration file. This practice can reduce the risk of malicious access to the encryption password for the following reasons: 1) The encryption password is not stored in the pwxcl.cfg file, and 2) You can remove the password from the command line after a successful cold start. If you specify the encryption password in a pwxcl command for a cold start and then need to restore the CDCT file later, you must enter the same encryption password in the RESTORE_CDCT command of the PWXUCDCT utility.</p> <p>To <i>not</i> encrypt PowerExchange Logger log files, do not enter an encryption password in the pwxcl command for a cold start or in the pwxcl.cfg configuration file.</p>
license	<p>The full path and file name for a license key file that overrides the default license.key file in the installation directory. The override file must have a file name or path that is different from that of the default file.</p> <p>This override file takes precedence over any other override license key file that you optionally specify with the PWX_LICENSE environment variable.</p>
specialstart	<p>Indicates whether to perform a special start of the PowerExchange Logger. A special start begins PowerExchange capture processing from a point in the change stream that you specify. This start point overrides the restart point based on information in the CDCT file. A special start does not delete any content from the CDCT file.</p> <p>Use this parameter to progress the restart point beyond problematic parts in the source logs. None of the data that was captured prior to the special start is lost.</p> <p>Enter one of the following values:</p> <ul style="list-style-type: none"> <li>- <b>Y.</b> Perform a special start of the PowerExchange Logger from the point in the change stream that is defined by the SEQUENCE_TOKEN and RESTART_TOKEN parameter values in the pwxcl.cfg configuration file. You must specify valid token values in the pwxcl.cfg file to perform a special start. These token values override the restart point based on information in the CDCT file. To generate valid token values, contact Informatica Global Customer Support. Ensure that the SEQUENCE_TOKEN value in the pwxcl.cfg file is greater than or equal to the current sequence token from the CDCT file.</li> </ul> <p><b>Note:</b> Because assistance from Customer Support is required to generate valid sequence and restart tokens, Informatica recommends that you use the specialstart parameter only at the direction of Customer Support.</p> <ul style="list-style-type: none"> <li>- <b>N.</b> Do not perform a special start. Perform a cold start or warm start as indicated by the coldstart parameter.</li> </ul> <p>Default is <b>N</b>.</p> <p>Do not specify both coldstart=Y and specialstart=Y. If you do, the coldstart=Y parameter takes precedence.</p>

**Note:** In these parameters, the full path is required only if the file is not in the default location.

## How the PowerExchange Logger Determines the Start Point for a Cold Start

When you cold start a PowerExchange Logger for Linux, UNIX, and Windows process, it uses the `RESTART_TOKEN` and `SEQUENCE_TOKEN` parameters, if present, in the `pwxccl` configuration file to determine the point in the change stream at which to start reading changes.

Based on how you set these parameters, the PowerExchange Logger starts from one of the following points in the change stream:

- If you enter valid token values in the `RESTART_TOKEN` and `SEQUENCE_TOKEN` parameters, the PowerExchange Logger starts from the point in the change stream that the token values specify. Use this method to start the PowerExchange Logger from a specific point.
- If you do not define the `RESTART_TOKEN` and `SEQUENCE_TOKEN` parameters, the PowerExchange Logger starts from the current end-of-log (EOL), or current point in time in the change stream.

**Tip:** You can generate restart and sequence tokens for the current EOL by running the `DTLUAPPL` utility with the `RSTTKN GENERATE` parameter or by performing a database row test with the `SELECT CURRENT_RESTART` SQL statement in PowerExchange Navigator.

- If you enter only zeroes (a single 0, or an even number of 0s) in the `RESTART_TOKEN` and `SEQUENCE_TOKEN` parameters, the PowerExchange Logger processing starts from one of the following start positions, depending on the data source type:
  - For DB2 for Linux, UNIX, and Windows sources, processing starts from the position at which the `DTLUCUDB` utility created the DB2 catalog snapshot to initialize the PowerExchange capture catalog table. However, you can change this default restart position with the `DTLUCUDB UPDTRDP` command.
  - For Microsoft SQL Server sources, processing starts from the position of the oldest available data in the distribution database.
  - For PowerExchange Express CDC for Oracle sources, processing starts from the beginning of the most recent Oracle archive log.
  - For PowerExchange Oracle CDC with LogMiner sources, processing starts from the position that corresponds to the most recent Oracle LogMiner dictionary dump.
  - For remote DB2 for i5/OS sources, processing starts from the beginning of the oldest receiver in the current chain of receivers.
  - For remote z/OS data sources, processing starts from the beginning of the PowerExchange Logger for MVS active log files.

## Cold Starting the PowerExchange Logger

Use this procedure to cold start the PowerExchange Logger. In the start statement, you must include the parameter `COLDSTART=Y`.

During a cold start, the PowerExchange Logger deletes the records in the CDCT file.

1. If you previously ran the PowerExchange Logger and have existing CDCT and log files, retain these files for historical purposes.

You can move or rename the files, as long as another PowerExchange Logger process is not using them. Do not delete them if you want to retain change processing history.

**Warning:** If you delete, move, or rename the CCT file, the capture registrations are no longer available.
2. In the `pwxccl` configuration file, set the `RESTART_TOKEN` and `SEQUENCE_TOKEN` parameters in a manner that causes the PowerExchange Logger to start from the appropriate point in the change stream.

3. To cold start the PowerExchange Logger, enter the following command at the command line:

```
pwxccl coldstart=y
```

The coldstart parameter must be set to y.

Include the optional config, cs, and license parameters if you want to override the default dbmover.cfg, pwxccl.cfg, and license.key files. On Linux and UNIX systems, you can add an ampersand (&) at the end of the statement to run the PowerExchange Logger in background mode. For more information about PowerExchange Logger start syntax, see the *PowerExchange Command Reference*.

## Managing the PowerExchange Logger

To assess the status of the PowerExchange Logger for Linux, UNIX, and Windows, you can display messages about PowerExchange Logger processing, memory use, and CPU use.

Occasionally, you might need to stop the PowerExchange Logger.

### RELATED TOPICS:

- [“Monitoring the PowerExchange Logger” on page 69](#)
- [“PowerExchange Logger Verbose Messages” on page 72](#)
- [“Controlling and Stopping PowerExchange Logger Processing” on page 68](#)
- [“PWXUCDCT Commands for Maintaining the PowerExchange Logger CDCT and Log Files” on page 72](#)
- [“Backing Up PowerExchange Logger Files” on page 74](#)
- [“Re-creating the CDCT File After a Failure” on page 75](#)

## Controlling and Stopping PowerExchange Logger Processing

PowerExchange provides commands for stopping the PowerExchange Logger, manually initiating a file switch, or starting another logging cycle.

You can enter these commands from the command line or by using the pwxcmd program on a remote Linux, UNIX, or Windows system. The output is displayed on screen and written to the PowerExchange message log.

**Note:** To use the pwxcmd program, you must specify the CONDENSENAME parameter in the pwxccl.cfg file and the SVCNODE statement in the dbmover.cfg file.

The following table describes each of these commands:

Command-line Command	pxcmd Command	Description
CONDENSE	condense	When the PowerExchange Logger is running in continuous mode, manually starts a new PowerExchange Logger logging cycle before the wait period for starting another cycle has elapsed. The wait period is defined by the NO_DATA_WAIT parameter in pwxcl.cfg.
FILESWITCH	fileswitch	Closes open PowerExchange Logger log files if they contain data and then switches to a new set of log files. If the log files do not contain data, the file switch does not occur.  If you use batch extraction mode, you can use this command to make change data in the current log files available for extraction processing before the next file switch is due to occur. To issue the fileswitch command from a script or batch file, you must use the pwxcmd program. Usually, you do not need to perform manual file switches if you use continuous extraction mode.
SHUTCOND	shutcond	Stops the PowerExchange Logger in a controlled manner after initiating and completing a final logging cycle. The final logging cycle enables the PowerExchange Logger to capture all of the changes up to point when the command is issued. After the logging cycle completes, the PowerExchange Logger closes open log files, updates the CDCT file, closes the CAPI, stops the Writer and Command Handler subtasks, and then ends the pwxcl program. Use this command if a logging cycle has not run recently.
SHUTDOWN	shutdown	Stops the PowerExchange Logger in a controlled manner after closing any open PowerExchange Logger log files and writing the latest restart position to the CDCT file. During shutdown processing, the PowerExchange Logger closes open log files, updates the CDCT file, closes the CAPI, stops the Writer and Command Handler subtasks, and then ends the pwxcl program. Use this command to stop a PowerExchange Logger process that is running in continuous mode.

For more information about command syntax, example output, and pwxcmd use, see the *PowerExchange Command Reference*.

## Monitoring the PowerExchange Logger

PowerExchange provides several commands that you can use for monitoring PowerExchange Logger processing and performance.

You can enter these commands from the command line or by using the pwxcmd program on a remote Linux, UNIX, or Windows system. The output is displayed on screen and written to the PowerExchange message log.

**Note:** To use the pwxcmd program, you must specify the CONDENSENAME parameter in the pwxcl.cfg file and the SVCNODE statement in the dbmover.cfg file.

The following table summarizes these commands:

Command-line Command	pwxcmd Command	Description
DG	pwxcmd displaystats -tp groups	<p>Displays monitoring statistics for each PowerExchange Logger group that is defined, if any. The statistics include:</p> <ul style="list-style-type: none"> <li>- The group name and the number of capture registrations in the group</li> <li>- The total number of insert, update, and delete records that the PowerExchange Logger processed for the group</li> <li>- The number of commits that the PowerExchange Logger processed for the group</li> <li>- The number of change records that the PowerExchange Logger has not yet flushed from memory to its log files on disk</li> <li>- The file name of the open Logger log file and the timestamp for when the file was opened</li> </ul> <p>To use this command, you must specify the STATS=(MONITOR) parameter in the pwxccl.cfg configuration file.</p>
DISPLAY CPU	pwxcmd displaycpu	<p>Displays the CPU time spent, in microseconds, for PowerExchange Logger processing during the current logging cycle, by processing phase. Processing phases include:</p> <ul style="list-style-type: none"> <li>- Reading source data</li> <li>- Writing data to PowerExchange Logger log files</li> <li>- Performing file switches</li> <li>- Performing "other processing," such as initialization and Command Handler processing of commands</li> </ul> <p>Also includes the total CPU time for all PowerExchange Logger processing.</p>
DISPLAY EVENTS	pwxcmd displayevents	Displays events that the PowerExchange Logger Controller, Command Handler, and Writer tasks are waiting on. Also indicates if the Writer is processing data or is in a sleep state while waiting for an event or timeout to occur.
DISPLAY MEMORY	pwxcmd displaymemory	Displays PowerExchange Logger memory use, in bytes, for each PowerExchange Logger task and subtask, with totals for the entire PowerExchange Logger process.
DISPLAY RECORDS	pwxcmd displayrecords	Displays counts of inserts, updates, deletes, and commits that the PowerExchange Logger processed during the current processing cycle. If the PowerExchange Logger did not receive changes during the current cycle, displays counts of change records for the current set of PowerExchange Logger log files.

Command-line Command	pwxcmd Command	Description
DISPLAY STATUS	pwxcmd displaystatus	Displays the status of the PowerExchange Logger Writer subtask.
DL	pwxcmd displaystats -tp logger	<p>Displays monitoring statistics for a PowerExchange Logger for Linux, UNIX, and Windows process and its tasks. The statistics include:</p> <ul style="list-style-type: none"> <li>- The PowerExchange Logger process ID</li> <li>- The status of the PowerExchange Logger Writer subtask at the time the command is issued</li> <li>- The CPU time used by the PowerExchange Logger since it started</li> <li>- PowerExchange Logger memory use by the Controller, Command Handler, and Writer tasks. For tasks, memory use is reported in the following categories: Current, Total, and Maximum.</li> <li>- Counts of inserts, updates, deletes, and commits that the PowerExchange Logger has processed, total and for the open Logger log file and the active logging cycle</li> </ul> <p>To use this command, you must specify the STATS=(MONITOR) parameter in the pwxccl.cfg configuration file.</p>

For more information about these commands and sample output, see the *PowerExchange Command Reference*.

Also, if you specify the optional *interval* subparameter in the STATS parameter in the pwxccl.cfg configuration file, you can publish the same monitoring statistics that are reported by the DL command at a specific interval:

```
STATS=(MONITOR,interval)
```

Fewer interval-based statistics messages are displayed on screen than are written to the message log to prevent flooding the screen with messages over time. For example, the following messages are displayed on the screen:

```
PWX-37132    Controller: (981/983/1849) KB    Command Handler: (0/0/34) KB    Writer: (5127/5147/5181)
KB
PWX-37135    Status 7144                      Totals  I=000000024344 U=000000000000 D=000000024336
C=000000004004 Total=000000052684
PWX-37136    CurrFileOpened : 2015-08-11 13:20:39 I=000000024344 U=000000000000 D=000000024336
C=000000004004 Total=000000052684
PWX-37137    Active Cycle : 2015-08-11 13:21:01 I=000000024344 U=000000000000 D=000000024336
C=000000004004 Total=000000052684
```

The following interval-based statistics are written to the PowerExchange message log:

```
PWX-37130    PWXCCL pid = 7144                Writer status = Reading or waiting for source data
PWX-37134    CPU Time = 0:00:02.589616
PWX-37131    Memory (Current/Total/Maximum)
PWX-37132    Controller: (981/983/1849) KB    Command Handler: (0/0/34) KB    Writer: (5127/5147/5181)
KB
PWX-37135    Status 7144                      Totals  I=000000024344 U=000000000000 D=000000024336
C=000000004004 Total=000000052684
PWX-37136    CurrFileOpened : 2015-08-11 13:20:39 I=000000024344 U=000000000000 D=000000024336
C=000000004004 Total=000000052684
PWX-37137    Active Cycle : 2015-08-11 13:21:01 I=000000024344 U=000000000000 D=000000024336
C=000000004004 Total=000000052684
```

## Determining If the PowerExchange Logger Captured Changes

To determine if the PowerExchange Logger for Linux, UNIX, and Windows captured committed changes for registered tables of interest, look for the following message in the PowerExchange message log:

```
PWX-09967 CAPI i/f: End of log for time 10/05/20 14:18:18 reached
```

This message indicates that the PowerExchange Logger read all of the changes that were available at the start the logging cycle. Look for this message if a PowerExchange Logger log file does not receive change data within the time period you expect. Delays can occur for various reasons. For example, if you cold start the PowerExchange Logger from the current restart point but a recent copy of the Oracle online catalog is not available in the archived redo logs, PowerExchange might need to read many archived logs before determining the point from which to begin capturing changes.

This message also indicates the point at which the PowerExchange Logger CATINT, CATBEGIN, and CATEND parameters take effect. These parameters control how often Oracle copies the catalog and the time period within which the copy operation occurs.

## PowerExchange Logger Verbose Messages

If you enter VERBOSE=Y in the pwxcl.cfg configuration file, the PowerExchange Logger for Linux, UNIX, and Windows produces more detailed messages during initialization, condense, fileswitch, record expiration, and shutdown processing. You can use these messages to assess PowerExchange performance and processing status.

For example, the following verbose messages indicate CPU use by the Writer subtask:

- Message PWX-33274 is issued before the Writer subtask starts reading source data after initialization and before the PowerExchange Logger shuts down:

```
PWX-33274 CPU Total number. CAPI Read number. Writing number. File switching number.
Other number
```

- Message PWX-33279 issued after each file switch:

```
PWX-33279 CPU total number. This file total number. CAPI Reads number. Writing file
number. Other number
```

If you do not use verbose messaging, you can use the DL, DG, DISPLAY CPU, and DISPLAY RECORDS commands to gather statistics that are useful for assessing PowerExchange Logger performance and status.

## PWXUCDCT Commands for Maintaining the PowerExchange Logger CDCT and Log Files

You can use the PWXUCDCT utility to maintain the PowerExchange Logger CDCT file and log files.

The following table describes the PWXUCDCT commands that you can use to perform maintenance tasks:

Command	Description
CONVERT_CDCT	<p>If you upgrade to 9.5.1 HotFix 1 or later from an earlier release, you can issue this command to manually perform a one-time conversion of the CDCT file to the new format. Alternatively, the first time the PowerExchange Logger is warm started, it automatically converts the CDCT file to the new format.</p> <p>The conversion creates a CDCT_<i>dbid</i> file instance from the original CDCT file. Ensure that the <i>dbid</i> value in the CDCT file name matches the DBID parameter value in the PowerExchange Logger pwxcl configuration file under which you run the command.</p> <p><b>Note:</b> If the old CDCT file contains information for multiple database instances, you must run this command multiple times, once for each instance. Each time you run the command, ensure that the CS parameter points to the correct pwxcl configuration file for the instance.</p>
CREATE_CDCT_BACKUP	<p>Manually creates a backup of all records in a CDCT file instance for a source database based on the latest configuration incarnation.</p> <p><b>Note:</b> The PowerExchange Logger automatically generates a backup at initialization and at termination.</p>



Command	Description
DELETE_EXPIRED_CDCT	This command is deprecated but is still supported for backward compatibility. Use DELETE_EXPIRED_FILES instead.
DELETE_EXPIRED_FILES	Delete the log files for which the retention period has expired and the CDCT records that reference those expired logs. For this command to work, you must set the <code>LOGGER_DELETES_EXPIRED_CDCT_RECORDS</code> parameter to N in the <code>pxwccf</code> configuration file.
DELETE_ORPHAN_FILES	Delete PowerExchange Logger log files that are not referenced by any record in the CDCT file.
DERIVE_CDCT_BACKUP	<p>If the CDCT file is corrupted or deleted and if a CDCT backup is not available or the latest available backup would result in significant reprocessing of data, use this command to derive a backup text file for recovery purposes.</p> <p>The command uses the <code>EXTERNAL_CAPTURE_MASK</code> parameter value from the PowerExchange Logger configuration file or the <code>external_capture_mask</code> positional parameter from the group definition file to generate a list of PowerExchange Logger log files. The command then uses the content of these log files to generate a text file that can be used as input to the RESTORE CDCT command.</p> <p>Do not use this command if the PowerExchange Logger log files were also corrupted or deleted.</p> <p><b>Tip:</b> Use the <code>PREVBACKUPFILE</code> parameter to supply the name of the last available backup file. By using a previous backup file, you preserve more historic information in the CDCT file. Also, the utility will add any log files that were created since the backup was taken to the derived backup file.</p>
REPORT_CDCT	<p>Print the contents of the CDCT file. This information is primarily for debugging purposes. For the current Logger configuration incarnation, the report shows:</p> <ul style="list-style-type: none"> <li>- Incarnation identifier, status, and reason (Rsn) for creation. The reason can be a cold start or a change in the configuration.</li> <li>- Source instance (or DBID) name and image type.</li> <li>- Number of groups defined in the group definition file. If no groups are defined, the default of 1 is used.</li> <li>- Begin and end timestamps.</li> <li>- Begin and end restart and sequence tokens.</li> </ul> <p>For each Logger group, the report shows:</p> <ul style="list-style-type: none"> <li>- Group number and name.</li> <li>- Incarnation to which the group belongs.</li> <li>- Path to the group log files.</li> <li>- Registration count.</li> <li>- Log file count and the first and current log sequence numbers.</li> <li>- Oldest log file timestamp.</li> </ul> <p>For each registration, the report shows:</p> <ul style="list-style-type: none"> <li>- Registration tag name and status.</li> <li>- Incarnation and group to which the registration belongs.</li> <li>- Default schema name.</li> </ul>
REPORT_CDCT_FILES	<p>Report the following information for each log file that is recorded in the CDCT:</p> <ul style="list-style-type: none"> <li>- Log file name and sequence number.</li> <li>- Configuration incarnation and group to which the log file belongs.</li> <li>- Record count, commit count, and whether any uncommitted data exists.</li> <li>- Begin and end timestamps.</li> <li>- Begin and end restart and sequence tokens.</li> <li>- File open timestamp.</li> <li>- File close timestamp.</li> </ul>

Command	Description
REPORT_CONFIG	List the parameter settings that are defined in the associated PowerExchange Logger pwxcl configuration file. If you created a group definition file and specified it in the GROUPDEFS parameter in the pwxcl file, the command also reports the group statements in the group definition file.
REPORT_EXPIRED_CDCT	This command is deprecated but is still supported for backward compatibility. Use REPORT_EXPIRED_FILES instead.
REPORT_EXPIRED_FILES	List the PowerExchange Logger log files for which the retention period has elapsed.
REPORT_FILES_BY_NAME	List PowerExchange Logger log files by file name. This information is based on directory information for the log files and not on the CDCT file. For each file, the command reports the following information: <ul style="list-style-type: none"> <li>- Date and time when the file was written.</li> <li>- Sequence number</li> <li>- Path and file name.</li> </ul> Also, the command reports the number of log files that match the default mask that is specified in the EXT_CAPT_MASK parameter of the pwxcl configuration file. If you specified a group definition file in the GROUPDEFS parameter of the pwxcl file, the command reports the number of log files that match any masks in the group definition file.
REPORT_FILES_BY_TIME	List PowerExchange Logger log files in the order in which they were created, from earliest to latest. This information is based on directory information for the log files and not on the CDCT file. For each file, the command reports the following information: <ul style="list-style-type: none"> <li>- Date and time when the file was written.</li> <li>- Sequence number.</li> <li>- Path and file name.</li> </ul> Also, the command reports the number of log files that match the default mask that is specified in the EXT_CAPT_MASK parameter of the pwxcl configuration file. If you specified a group definition file in the GROUPDEFS parameter of the pwxcl file, the command reports the number of log files that match any masks in the group definition file.
REPORT_ORPHAN_FILES	List PowerExchange Logger log files that are not referenced by any record in the CDCT file.
RESTORE_CDCT	Restore the CDCT file from a backup, up to a specific point in time. The PowerExchange Logger will reprocess any data that is later than this point in time. After the restore operation completes, run the DELETE_ORPHAN_FILES command.

In the command syntax, include the CS parameter to specify the path to a specific PowerExchange Logger pwxcl configuration file. For more information about the PWXUCDCT utility commands, see the *PowerExchange Utilities Guide*.

## Backing Up PowerExchange Logger Files

The PowerExchange Logger automatically creates a backup of the CDCT file at initialization and normal termination. You must manually back up the log files.

If a recent generated CDCT backup is not available, you can use the PWXUCDCT utility to manually create a CDCT backup. Use the PWXUCDCT utility CREATE\_CDCT\_BACKUP command to back up the CDCT file based on the latest Logger configuration incarnation.

Alternatively, if a failure causes the CDCT file and its recent backups to become damaged or deleted, you can use the `DERIVE_CDCT_BACKUP` command to derive a backup based on the available PowerExchange Logger log files, optionally in conjunction with the last available backup file. This previous backup can be an automatically generated backup or one that you created. By using the previous backup file, you preserve more historic information in the CDCT file. The utility will add any log files that were created since that backup was taken to the derived backup file.

**Tip:** If you manually back up the CDCT file and the log files, try to perform the backup during a period of low database activity, when no or little data is being written to the log files.

## Re-creating the CDCT File After a Failure

If the CDCT file and its recent backups are damaged or deleted, you can re-create the CDCT file based on the available PowerExchange Logger log files. After you derive the CDCT backup, you can use it to restore the CDCT file.

This procedure assumes that PowerExchange Logger log files are available. Do not use this procedure if the log files were also damaged or deleted.

1. Issue the `PWXUCDCT` utility `DERIVE_CDCT_BACKUP` command to derive a backup from available PowerExchange Logger log files.  
**Tip:** Include the `PREVBACKUPFILE` parameter to supply the name of the last available CDCT backup file. By using this previous backup file, you preserve more historic information in the CDCT file. The utility will add any log files that were created after this backup was taken to the derived backup file.
2. To restore the CDCT file from the derived backup, issue the `PWXUCDCT` utility `RESTORE_CDCT` command.
3. Verify that the restore operation was successful as follows:
  - Verify that the return code from the `PWXUCDCT` utility is zero.
  - Verify that messages `PWX-25140` through `PWX-25145` provide reasonable record counts for the records read from the backup file and for the records that were changed in the CDCT file.
4. Run the `DELETE_ORPHAN_FILES` command to delete log files that are no longer referenced by the restored CDCT file.

After you warm start the PowerExchange Logger, it re-creates the CDCT content for those files.

For more information about the `PWXUCDCT` utility commands, see the *PowerExchange Utilities Guide*.

# Part III: PowerExchange CDC Data Sources

This part contains the following chapters:

- [DB2 for Linux, UNIX, and Windows CDC, 77](#)
- [Microsoft SQL Server CDC, 94](#)
- [Oracle CDC with LogMiner, 109](#)
- [Express CDC for Oracle, 139](#)
- [Remote Logging of Data, 188](#)

## CHAPTER 4

# DB2 for Linux, UNIX, and Windows CDC

This chapter includes the following topics:

- [DB2 for Linux, UNIX, and Windows CDC Overview, 77](#)
- [Planning for DB2 CDC, 78](#)
- [Configuring DB2 for CDC, 81](#)
- [Configuring PowerExchange for DB2 CDC, 82](#)
- [Using a DB2 Data Map, 89](#)
- [Managing DB2 CDC, 90](#)
- [DB2 for Linux, UNIX, and Windows CDC Troubleshooting, 92](#)

## DB2 for Linux, UNIX, and Windows CDC Overview

PowerExchange captures change data from the DB2 for Linux, UNIX, and Windows database logs for the database that contains your source tables. PowerExchange uses the PowerExchange Client for PowerCenter (PWXPC) to coordinate with PowerCenter to move the captured change data to one or more targets.

For PowerExchange to capture DB2 change data, you must perform the following configuration tasks in DB2:

- Ensure that archive logging is active for the database.
- Create a PowerExchange capture catalog table in the database. The capture catalog table stores information about all tables in the source database, including column definitions and DB2 log positions.

Also, perform the following configuration tasks in PowerExchange:

- Define a capture registration for each source table. In the capture registration, you can select a subset of columns for which to capture data. PowerExchange generates a corresponding extraction map. Optionally, you can define an additional extraction map.
- If a source table contains columns in which you store data in a format that is inconsistent with the column datatype, you can optionally create a data map to manipulate that data with expressions. For example, if you store packed data in a CHAR column, you can create a data map to manipulate and prepare that data for loading to a target. You must merge the data map with the extraction map for the source table during capture registration creation.

- If you want to use the PowerExchange Logger for Linux, UNIX, and Windows to capture change data and write it to PowerExchange Logger log files, configure the PowerExchange Logger. The change data is then extracted from the PowerExchange Logger log files. Benefits of the PowerExchange Logger include fewer database accesses, faster CDC restart, and no need to prolong retention of DB2 log files for change capture. PowerExchange works in conjunction with PowerCenter to extract change data from DB2 database logs or PowerExchange Logger log files and load that data to one or more targets.

#### RELATED TOPICS:

- [“Planning for DB2 CDC” on page 78](#)
- [“Configuring PowerExchange for DB2 CDC” on page 82](#)
- [“Using a DB2 Data Map” on page 89](#)
- [“Managing DB2 CDC” on page 90](#)
- [“Introduction to Change Data Extraction” on page 208](#)
- [“PowerExchange Logger Overview” on page 35](#)

## Planning for DB2 CDC

Before you configure DB2 for Linux, UNIX, and Windows CDC, verify that the following prerequisites and user authority requirements are met. Also, review the restrictions so that you can properly configure CDC.

### Prerequisites

PowerExchange CDC has the following prerequisites:

- Archive logging must be active for the database that contains the source tables from which change data is to be captured.
- DB2 source tables must be defined with the `DATA CAPTURE CHANGES` clause for capture processing to occur.
- A valid DB2 environment must exist for the PowerExchange user. On Linux, and UNIX, the path to the DB2 client must be specified in the `PATH` and library path environment variables.

### Required User Authority

For PowerExchange to read change data from DB2 logs, the user ID that you specify for database access must have SYSADM or DBADM authority. Usually, you specify this user ID in the `UDB CAPL_CONNECTION` statement in the `dbmover.cfg` file.

## DB2 Datatypes Supported for CDC

PowerExchange supports most DB2 for Linux, UNIX, and Windows datatypes for CDC.

The following table identifies the DB2 for Linux, UNIX, and Windows source datatypes that PowerExchange supports and does not support for CDC:

DB2 Datatype	Supported for CDC?	Comments
BIGINT	Yes	-
BLOB	No	If you register a table with large object (LOB) columns, PowerExchange does not capture changes for the LOB columns but can capture changes for other columns in the table.
CHAR	Yes	-
CLOB	No	If you register a table with LOB columns, PowerExchange does not capture changes for the LOB columns but can capture changes for other columns in the table.
DATE	Yes	-
DBCLOB	No	If you register a table with LOB columns, PowerExchange does not capture changes for the LOB columns but can capture changes for other columns in the table.
DECFLOAT	No	If you register a table with DECFLOAT columns, PowerExchange does not capture changes for the DECFLOAT columns but can capture changes for other columns in the table.
DECIMAL	Yes	-
DOUBLE	Yes	-
GRAPHIC	Yes	-
INTEGER	Yes	-
LONG VARCHAR	Yes	-
LONG VARGRAPHIC	Yes	-
REAL	Yes	-
REF	No	DB2 does not allow change data capture for tables with REF columns.
SMALLINT	Yes	-
TIME	Yes	-
TIMESTAMP	Yes	-
UDTs <sup>1</sup>	No	PowerExchange does not capture change data for tables with UDT columns.
VARCHAR	Yes	-

DB2 Datatype	Supported for CDC?	Comments
VARGRAPHIC	Yes	-
XML	No	If you register a table with XML columns, PowerExchange does not capture changes for the XML columns but does capture changes for other columns in the table.
1. User-defined datatypes, such as DISTINCT and STRUCT.		

## DB2 CDC Considerations

Consider the following CDC capabilities and restrictions when planning DB2 CDC processing.

- To extract change data on a DB2 client machine that is remote from the DB2 server where the change data is captured, both machines must have the same architecture. Otherwise, change data capture processing might fail with the error message PWX-20628.
- For DB2 9.7 and later sources, PowerExchange can capture change data from tables that use DB2 row compression. These tables were created or altered with the COMPRESS YES option.
- If the source tables are compressed, ensure that you have a compression dictionary that is compatible with the compressed DB2 log records from which PowerExchange reads change data for the tables. Otherwise, DB2 cannot decompress the log records for PowerExchange read requests. Usually, the compatible compression dictionary is available because DB2 maintains the current compression dictionary and a backup of the previous compression dictionary on disk.

If you run the DB2 REORG TABLE utility or the DB2 LOAD utility with the REPLACE or RESUME NO option against compressed source tables, Informatica recommends that you specify the KEEPDICTIONARY option for the utility. The KEEPDICTIONARY option forces DB2 to retain the current compression dictionary if it exists. If you use the RESETDICTIONARY option, DB2 rebuilds compression dictionary. In this case, the previous compression dictionary that matches the DB2 log records might not be available any longer.

- PowerExchange cannot capture change data for the following DB2 datatypes:
  - DECFLOAT, LOB, and XML datatypes. You can create a capture registration for a table that includes columns with DECFLOAT, LOB, and XML datatypes. However, the registration does not include these columns, and PowerExchange does not capture change data for them. PowerExchange does capture change data for the other columns in the registered table that have supported datatypes.
  - User-defined datatypes. Tables that include columns with user-defined datatypes cannot be registered for change data capture. PowerExchange cannot capture change data for these tables.
- To add or drop partitions in a partitioned database and then redistribute table data across the updated partition group, or to reconfigure a database partition group, you must use a special procedure. Otherwise, PowerExchange might not be able to resume change data capture properly.
- If you alter a column datatype to or from FOR BIT DATA, PowerExchange does not detect the datatype change. PowerExchange continues to use the datatype that is specified in the existing capture registration.
- If you alter a source table to change the DEFAULT value of a DB2 column of CDC interest, PowerExchange does not detect this DDL change during capture processing. As a result, the correct DEFAULT value is not available when PowerExchange performs the following operations:
  - Delivers the pre-existing short rows for a table to which columns were added.



- Delivers rows for source tables that use the VALUE COMPRESSION option and that include a column with the COMPRESS SYSTEM DEFAULT option and a default value.
- In a partitioned database, if an UPDATE to a table row changes the partition key and that change causes the row to move to another partition, PowerExchange processes the UPDATE as two operations: a DELETE and an INSERT. However, based on the DB2 log information, PowerExchange cannot predictably determine the order in which to perform the DELETE and INSERT operations. If the INSERT is processed first, both the original row and the updated row appear on the target until the DELETE is processed.
- The maximum length of a row from which PowerExchange can capture change data is 128,000 bytes.
- PowerExchange uses multithreaded processing for change data capture. By default, PowerExchange uses up to nine threads. To configure the number of threads, specify the THREADING parameter in the UDB CAPI CONNECTION statement. If you have a DB2 partitioned database, you can use a maximum of one thread for each database partition node plus two additional threads for the CAPI and merge processing.
- PowerExchange does not support DB2 10.5 sources on zLinux.

## Configuring DB2 for CDC

To configure DB2 for Linux, UNIX, or Windows for PowerExchange CDC, perform the following tasks:

1. In the DB2 Control Center Configure Database Logging Wizard, enable archive logging for the DB2 database. For more information, see the IBM DB2 documentation.  
If archive logging is not enabled, PowerExchange issues the error messages PWX-20204 and PWX-20229 during CDC.
2. Set the following user environment variables in any process that runs PowerExchange CDC or the DTLUCUDB program:
  - Set DB2NOEXITLIST to ON.
  - Set DB2CODEPAGE to 1208.
3. Verify that the DB2 source tables are defined with the DATA CAPTURE CHANGES clause.
4. To enable PowerExchange to report the authorization ID and application that is associated with a DB2 transaction in monitoring message PWX-20177, set the DB2\_LOGGING\_DETAIL registry variable to APPLINFO in DB2.  
To set this variable for the current DB2 instance, enter the following command:
 

```
db2set DB2_LOGGING_DETAIL=APPLINFO
```

 To set this variable for all DB2 instances on the system, enter the following command:
 

```
db2set -g DB2_LOGGING_DETAIL=APPLINFO
```
5. If a table that is selected for change data capture includes columns with a LONG datatype, use the INCLUDE LONGVAR COLUMNS clause to alter the table so that PowerExchange can capture data for the LONG columns. Otherwise, PowerExchange might issue the error message PWX-20094 during CDC processing.

# Configuring PowerExchange for DB2 CDC

The tasks that you perform to configure PowerExchange for DB2 for Linux, UNIX, and Windows CDC depend on whether you want to use the PowerExchange Logger for Linux, UNIX, and Windows and the extraction mode you plan to use.

## RELATED TOPICS:

- [“Configuring PowerExchange CDC without the PowerExchange Logger” on page 82](#)
- [“Configuring PowerExchange CDC with the PowerExchange Logger” on page 82](#)
- [“Creating the Capture Catalog Table” on page 83](#)
- [“Initializing the Capture Catalog Table” on page 83](#)
- [“Customizing the dbmover Configuration File for DB2 CDC” on page 84](#)

## Configuring PowerExchange CDC without the PowerExchange Logger

If you plan to run extractions in real-time extraction mode and *not* use the PowerExchange Logger for Linux, UNIX, and Windows, complete the following tasks to configure PowerExchange CDC:

1. Create the PowerExchange capture catalog table.
2. Run the DTLUCUDB SNAPSHOT command to initialize the capture catalog table.
3. When you configure the dbmover.cfg file, include the following statements:
  - CAPT\_PATH
  - CAPT\_XTRA
  - UDB CAPI\_CONNECTION
4. In the PowerExchange Navigator, create a capture registration for each source table. The PowerExchange Navigator generates a corresponding extraction map. Optionally, create a data map if you want to perform field-level processing.

**Tip:** Set the **Condense** option to **Part** even though you do not plan to use the PowerExchange Logger, unless you have a specific reason not to do so. This practice prevents having to edit the capture registrations later if you decide to use the PowerExchange Logger. You might want to set the **Condense** option to **None** if you plan to run both real-time and continuous extractions against tables defined by the same capture registrations and you do not want the PowerExchange Logger to capture change data for some registered tables.

If capture registrations already exist for the source tables, delete the existing registrations and extraction maps and create new ones.

5. Activate the capture registrations. Usually, you do this task after materializing the targets.

**Next Step:** Configure and start extractions. You must use real-time extraction mode.

## Configuring PowerExchange CDC with the PowerExchange Logger

If you plan to use the PowerExchange Logger for Linux, UNIX, and Windows and run extractions in batch or continuous extraction mode, complete the following tasks to configure PowerExchange CDC:

1. Create the PowerExchange capture catalog table.
2. Run the DTLUCUDB SNAPSHOT command to initialize the capture catalog table.

3. When you configure the dbmover.cfg file, include the following statements:
  - CAPT\_PATH
  - CAPT\_XTRA
  - UDB CAPI\_CONNECTION
  - CAPX CAPI\_CONNECTION (for continuous extraction mode only)
4. Configure the pwxcl.cfg file for the PowerExchange Logger.
5. In the PowerExchange Navigator, create a capture registration for each DB2 source table. You must select **Part** in the **Condense** drop-down list. The PowerExchange Navigator generates a corresponding extraction map.  
  
If capture registrations already exist for these tables, delete the existing registrations and extraction maps and create new ones.
6. Activate the capture registrations. Usually, you do this task after materializing the targets.
7. Start the PowerExchange Logger.

**Next Step:** Configure and start extractions. You can use either batch extraction mode or continuous extraction mode.

## Creating the Capture Catalog Table

The PowerExchange capture catalog table stores information about the CDC source tables, column definitions, and valid DB2 log positions. You must create this table in the same database that contains the source tables from which change data is captured.

If the database has multiple partitions, the capture catalog table stores positioning information for each partition. If the database has only a single partition, the capture catalog table still contains positioning information for the partition.

Use the following DDL to create the capture catalog table:

```
CREATE TABLE DTLCCATALOG (
    VTSTIME    TIMESTAMP      NOT NULL,
    VTSACC     INTEGER        NOT NULL,
    NODENUM    SMALLINT       NOT NULL,
    SEQ        INTEGER        NOT NULL,
    TBSHEMA    VARCHAR(128),
    TBNAME     VARCHAR(128),
    OP         VARCHAR(1024)  NOT NULL,
    PRIMARY KEY (VTSTIME, VTSACC, NODENUM, SEQ) )
;
```

In this DDL, the table name is DTLCCATALOG. If necessary, you can specify another table name.

**Tip:** Informatica recommends that you place the PowerExchange capture catalog table in the DB2 catalog partition.

## Initializing the Capture Catalog Table

To initialize the PowerExchange capture catalog table, run the DTLUCUDB utility with the SNAPSHOT command. You should need to do this task only once.

To specify the command, use the following syntax:

```
DTLUCUDB SNAPSHOT [DB=database_name] [CCATALOG=capture_catalog_name] [UID=user_id]
[EPWD=encrypted_password] [REPLACE=Y|N]
```

If the capture catalog table contains existing rows of data, you must set the REPLACE parameter to Y to enable PowerExchange to overwrite the data. For a new capture catalog table, accept the default of N.

After the snapshot successfully completes, back up the capture catalog table to create a point of consistency for recovery.

**Note:** If you run the DTLUCUDB SNAPSHOT command while the DB2 catalog is being updated, the snapshot fails. If this failure occurs, run the SNAPSHOT command again after the DB2 catalog updates are complete.

## Customizing the dbmover Configuration File for DB2 CDC

In the dbmover configuration file, add statements for DB2 for Linux, UNIX, and Windows CDC.

You must include the following statements for DB2 CDC:

### **CAPT\_PATH statement**

Path to the local directory on a Linux, UNIX, or Windows system that contains the control files for CDC.

These files are: the CCT file for capture registrations, the CDEP file for application names that are used for ODBC extractions, and the CDCT file for the PowerExchange Logger for Linux, UNIX, and Windows.

### **CAPT\_XTRA statement**

Path to the local directory on a Linux, UNIX, or Windows system that stores extraction maps for CDC.

### **UDB CAPI\_CONNECTION statement**

A named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and to control CDC processing for DB2 for Linux, UNIX, and Windows sources.

Add this statement to the dbmover.cfg file on the system where DB2 capture registrations reside. This location corresponds to the **Location** node that you specify when defining a registration group. Usually, this location is where the source database resides.

If you plan to use continuous extraction mode, you must also include the following statement:

### **CAPX CAPI\_CONNECTION statement**

A named set of parameters that the CAPI uses for continuous extraction of change data from PowerExchange Logger for Linux, UNIX, and Windows log files.

Also, Informatica recommends including the LOGPATH and TRACING statements to make finding messages easier. The LOGPATH statement defines a directory specifically for PowerExchange message log files, and the TRACING statement enables PowerExchange to create an alternative set of message log files for each PowerExchange process.

For more information about all DBMOVER statements, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“CAPI\\_CONNECTION - CAPX Statement” on page 57](#)
- [“CAPI\\_CONNECTION - UDB Statement” on page 85](#)
- [“CAPT\\_PATH Statement” on page 30](#)
- [“CAPT\\_XTRA Statement” on page 30](#)

## Example dbmover Statements for DB2

This example shows dbmover statements that are typically defined for DB2 for Linux, UNIX, and Windows CDC.

```
CAPT_PATH=c:/pwxcapt/Vnnn
CAPT_XTRA=c:/pwxcapt/Vnnn/extrmaps
CAPI_CONN_NAME=UDBCC
CAPI_CONNECTION=(NAME=UDBCC
```

```
,DLLTRACE=bbbb
,TYPE=(UDB
      ,CCATALOG=mylib.captcat_tbl
      ,USERID=db2admin
      ,PASSWORD=db2admin))
```

## CAPI\_CONNECTION - UDB Statement

The UDB CAPI\_CONNECTION statement specifies a named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and to control CDC processing for DB2 for Linux, UNIX, and Windows sources.

Add this statement to the dbmover.cfg file on the system where DB2 capture registrations reside. This location corresponds to the **Location** node that you specify when defining a registration group. Usually, this location is where the source database resides.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** DB2 for Linux, UNIX, and Windows

**Related Statements:** N/A

**Required:** Yes for CDC

**Syntax:**

```
CAPI_CONNECTION=( [DLLTRACE=trace_id]
                  ,NAME=capi_connection_name
                  [,TRACE=trace_name]
                  ,TYPE=(UDB
                        [,CCATALOG={capture_catalog|creator.DTLCCATALOG}]
                        [,DBCONN=database_name]
                        [,EPWD=encrypted_password]
                        [,LARGEOPS=number_of_operations]
                        [,MEMCACHE={cache_size|1024}]
                        [,MONITORINT={minutes|5}]
                        [,PASSWORD=password]
                        [,RSTRADV=seconds]
                        [,SPACEPRI={AUTO|MAX|NONE|nn}]
                        [,THREADING={AUTO|MAX|NONE|nn}]
                        [,UDBSchema=schema]
                        [,UPDINT={seconds|600}]
                        [,UPDREC={records|1000}]
                        [,USERID=user_id]
                        )
                  )
```

**Parameters:**

**DLLTRACE=trace\_id**

Optional. A user-defined name for the TRACE statement that activates internal DLL tracing for this CAPI.

Specify this parameter only at the direction of Informatica Global Customer Support.

**NAME=capi\_connection\_name**

Required. A unique user-defined name for this CAPI\_CONNECTION statement.

Maximum length is eight alphanumeric characters.

**TRACE=trace\_name**

Optional. A user-defined name for the TRACE statement that activates the common CAPI tracing.

Specify this parameter only at the direction of Informatica Global Customer Support.

**TYPE=(UDB, ... )**

Required. The type of CAPI\_CONNECTION statement. For DB2 for Linux, UNIX, and Windows sources, this value must be UDB.

**CCATALOG={capture\_catalog|creator.DTLCCATALOG}**

Optional. The name of the PowerExchange capture catalog table in the format *creator.table\_name*.

Default is *creator.DTLCCATALOG*, where *creator* is the user ID that is used to connect to the database.

**DBCONN=database\_name**

Optional. The name of the override database that you want to connect to for data extraction instead of the database that is specified for the registration group. The override database must contain tables and columns that are identical to those in the original database. The registration tag names and extraction map names include the original database name.

**EPWD=encrypted\_password**

Optional. An encrypted password that is used with the user ID in the USERID parameter for database access.

If you specify this parameter, you must also specify either the USERID parameter. However, do not also specify the PASSWORD parameter.

**Tip:** You can create encrypted passwords in the PowerExchange Navigator.

**LARGEOPS=number of operations**

Optional. Overrides the default value that PowerExchange uses to identify transactions as large transactions for reporting purposes. Enter the number of DML operations (inserts, updates, and deletes), in thousands, that a transaction must contain to be considered a large transaction.

PowerExchange issues status messages for large transactions that meet this criteria. If PowerExchange issues too many messages, you can increase this value to limit the number of messages.

Valid values are 1 through 2147483 (1000 through 2,147,483,000 operations). The default value is one half of the MEMCACHE parameter value rounded up to the nearest thousand. Based on the default MEMCACHE value of 1024 KB, the default LARGEOPS value is 1000 (1,000,000 operations).

**Note:** If a committed transaction spans multiple partitions in a DB2 database, PowerExchange reports the number of SQL operations and transaction size across all of the partitions.

**MEMCACHE={cache\_size|1024}**

Optional. The maximum memory cache size, in kilobytes, that PowerExchange can allocate to reconstruct complete UOWs.

Enter a number from 0 through 2147483647. Default is 1024 KB. If you enter 0, the memory cache size is limited only by the available memory on the system. Informatica recommends that you enter 0.

For each extraction session, PowerExchange keeps all changes for each pending UOW in the memory cache until it processes the end-UOW record. PowerExchange incrementally allocates memory cache up to the limit that this parameter specifies. If the memory cache is too small to hold all of the changes in the pending UOWs, PowerExchange spills the changes in a UOW to sequential files, called UOW spill files, on disk.

Each UOW spill file contains change data from one UOW. A UOW might require multiple UOW spill files to hold all of the changes for that UOW. If the change stream contains multiple large UOWs and the memory cache is insufficient, PowerExchange might create numerous UOW spill files.

PowerExchange processes the change stream more efficiently if it does not need to use UOW spill files. A large number of UOW spill files can degrade extraction performance and cause disk space shortages.

**Important:** If the change stream contains small UOWs, the default value might be sufficient. However, the default value is often too small to eliminate UOW spill files.

The location in which PowerExchange allocates the UOW spill files varies by operating system, as follows:

- For Linux and UNIX, PowerExchange uses the current directory by default. To use a different directory, you must specify the TMPDIR environment variable.

PowerExchange names the UOW spill files using the prefix "dtlq" and the operating system function tempnam.

**Note:** The UOW spill files are temporary files that are deleted when PowerExchange closes them. These files are not visible in the directory while they are open.

- For Windows, PowerExchange uses the current directory by default for UOW spill files. To use a different directory, specify the TMP environment variable.

PowerExchange names the UOW spill file names using the prefix "dtlq" and the Windows \_tempnam function.

**Warning:** PowerExchange allocates the cache size for each extraction operation. If you use a large MEMCACHE value and run many concurrent extraction sessions, memory constraints can occur.

**MONITORINT=minutes**

Optional. The time interval, in minutes, at which PowerExchange checks transaction activity for long outstanding transactions and large transactions. A long outstanding transaction is one that remains active for two monitoring intervals, and a large transaction is one that meets the LARGEOPS criteria. When this interval elapses, PowerExchange issues messages that identify the large transactions and long outstanding transactions and report their processing activity. PowerExchange also issues messages that identify the current position in the change stream. Valid values are 0 through 720. A value of 0 disables monitoring. Default is 5.

**PASSWORD=password**

Optional. A clear text password that is used with the user ID in the USERID parameter for database access.

If you specify this parameter, you must also specify either the USERID parameter. However, do not also specify the EPWD parameter.

**RSTRADV=seconds**

The time interval, in seconds, that PowerExchange waits before advancing restart and sequence tokens for a registered data source during periods when UOWs do not include any changes of interest for the data source. When the wait interval expires, PowerExchange returns the next committed "empty UOW," which includes only updated restart information.

Enter a number from 0 through 86400. No default is provided.

PowerExchange resets the wait interval to 0 when one of the following events occur:

- PowerExchange completes processing a UOW that includes changes of interest.
- PowerExchange returns an empty UOW because the wait interval expired without PowerExchange receiving any changes of interest.

For example, if you specify 5, PowerExchange waits 5 seconds after it completes processing the last UOW or after the previous wait interval expires. Then PowerExchange returns the next committed empty UOW that includes the updated restart information and resets the wait interval to 0.

If you do not specify RSTRADV, PowerExchange does not advance restart and sequence tokens for a registered source during periods when PowerExchange receives no changes of interest. When PowerExchange warm starts, it reads all changes, including those not of interest for CDC, from the restart point.

**Attention:** A value of 0 can degrade performance. In addition to the UOWs that contain changes for registered sources of interest, PowerExchange returns an empty UOW for every UOW that does not contain changes for the registered sources of interest.

**SPACEPRI={primary\_space|2147483647}**

Optional. The amount of disk space, in bytes, that PowerExchange uses to allocate UOW spill files as temporary files.

Enter a number from 1 through 2147483647. Default is 2147483647 bytes.

**THREADING={AUTO|MAX|NONE|nn}**

Optional. Controls the number of threads that the UDB CAPI uses to capture change data from a DB2 database. Use this parameter to improve the performance of capture processing. If you have a partitioned database, you can use a maximum of one thread for each database partition node plus two additional threads for CAPI and merge processing.

Valid values are:

- **AUTO.** Use up to nine threads.
- **MAX.** Use one thread for each database partition plus two additional threads for CAPI and merge processing. The maximum number of threads is 99.
- **NONE.** Do not use multiple threads for capture processing.
- **nn.** A user-specified number of threads. Valid values are 1 to 99. For a partitioned database, if you enter a value that exceeds the sum of (*number\_of\_database\_partitions* + 2), the CAPI does not use the excess threads.

Default is AUTO.

**UDBSchema=schema**

Optional. A schema name that overrides the schema name in capture registrations.

**UPDINT={seconds|600}**

Optional. The minimum number of seconds that PowerExchange must wait after encountering a virtual time stamp (VTS) in the DB2 log records for a partition before writing a positioning entry to the PowerExchange capture catalog table. The positioning entry, which contains a log sequence number (LSN) and VTS, indicates the location in the DB2 logs.

Enter a number from 1 through 2147483647. Default is 600 seconds.



**Note:** The minimum number of records that is specified in the UPDREC parameter must also be met before PowerExchange can write positioning entries to the capture catalog table.

**UPDREC={records|1000}**

Optional. The minimum number of DB2 log records that PowerExchange must read for a partition before writing a positioning entry to the PowerExchange capture catalog table. The positioning entry contains an LSN and VTS and indicates a location in the DB2 logs.

Enter a number from 1 through 2147483647. Default is 1000 records.

**Note:** The minimum wait period that is specified in the UPDINT parameter must also be met before PowerExchange can write positioning entries to the capture catalog table.

**USERID=user\_id**

Optional. A database user ID that has SYSADM or DBADM authority.

If you specify this parameter, you must also specify the PASSWORD or EPWD parameter.

## Using a DB2 Data Map

If you want PowerExchange to perform field-level processing on some records in a DB2 for Linux, UNIX, and Windows source table, you must use a data map.

For example, in some DB2 environments, a table can contain a single column that stores an array of fields in a format that is not consistent with the column datatype, such as a CHAR or VARCHAR column that stores multiple packed data fields. You can use an expression to modify this data before PowerCenter replicates it to a target. Also, if you add a user-defined field to a table in record view, you can build an expression to populate it. In the PowerExchange Navigator, you can define expressions only for data maps.

You might have data maps available for your source tables if you used PowerExchange bulk data movement to materialize your data targets. Bulk data movement requires data maps. You can use the bulk data maps for CDC if you merge them with the extraction maps for your data sources. The PowerExchange Navigator automatically generates an extraction map when you create a capture registration. Alternatively, you can manually add an extraction map.

**Note:** The field names in the data map must match the actual column names, as indicated in the DB2 capture registration.

## Task Flow for DB2 Data Map Use

Perform the following tasks to use a DB2 data map for change data capture:

1. In the PowerExchange Navigator, create a capture registration for the DB2 source table.
2. Create a DB2 data map for the same DB2 source table if one is not available from a previous bulk data movement operation.
3. Merge the DB2 data map with the extraction map for the table.
4. Perform a row test on the merged extraction map.

For more information, see the *PowerExchange Navigator User Guide*.

# Managing DB2 CDC

You might need to stop DB2 for Linux, UNIX, and Windows CDC for source tables occasionally, for example, to change the table definitions.

## RELATED TOPICS:

- [“Stopping DB2 CDC” on page 90](#)
- [“Reconfiguring a Partitioned Database or Database Partition Group” on page 91](#)
- [“DB2 for Linux, UNIX, and Windows CDC Troubleshooting” on page 92](#)
- [“Changing a DB2 Source Table Definition” on page 90](#)

## Stopping DB2 CDC

You might need to stop change data capture for a DB2 source table to perform troubleshooting or routine maintenance tasks, such as maintenance on the capture catalog table or redistribution of table data across reconfigured database partitions.

To stop change data capture, use one of the following methods:

- Open the capture registration for a source table, and change the **Status** option from **Active** to **History**.  
**Warning:** After you set the status to **History**, you cannot activate the registration again. This status change permanently stops change data capture based on the capture registration.
- To temporarily stop change data capture, alter the DB2 table to specify the DATA CAPTURE NONE clause:

```
ALTER owner.table_name DATA CAPTURE NONE
```

When DATA CAPTURE NONE is specified, DB2 no longer writes changes to the DB2 log files in expanded format. Because CDC requires expanded format, PowerExchange can no longer capture change data for the table from the log files. If you set it back to DATA CAPTURE CHANGES, you might need to rematerialize the targets.

## RELATED TOPICS:

- [“Stopping PowerCenter CDC Sessions” on page 248](#)

## Changing a DB2 Source Table Definition

Occasionally, you might need to change the structural definition of a DB2 for Linux, UNIX, and Windows source table that is registered for change data capture.

If you make table definition changes that affect the columns from which change data is captured, perform this procedure to enable PowerExchange to switch to the updated table definition, while preserving access to previously captured data. These table definition changes include adding, altering, or dropping columns. Do not perform this procedure if you are selectively capturing change data for a subset of columns and none of the selected columns are affected by the table definition changes.

**Tip:** If you no longer need to capture change data from a column in a table, you can remove the column from the extraction map without changing the capture registration. Change data for that column is still captured but is not extracted.

1. Stop all transactions, applications, and other activity that update the source table.
2. Verify that any change data that was captured under the previous table definition has completed extraction processing. Then stop all workflows that extract change data for the table.
3. If you use the PowerExchange logger for Linux, UNIX, and Windows, shut down the Logger.

4. In the PowerExchange Navigator, delete the original capture registration and extraction map.
5. Use DDL to make the table changes.
6. In the PowerExchange Navigator, create a new capture registration that reflects the metadata changes and set its status to **Active**. A corresponding extraction map is generated.

**Tip:** When you create the capture registration, use the original registration name so that you will not need to edit the map name that is used by the PowerCenter CDC workflows that contain the changed source table.

7. If you shut down the PowerExchange Logger, warm start it.  
The PowerExchange Logger begins capturing changes based on the new capture registration.
8. If necessary, change the target table definition to reflect the source table metadata changes.
9. In PowerCenter Designer, import the extraction map for the altered source table to create a new source definition. Edit the mapping if necessary.

If you also changed the target table, edit or re-create the target definition. Then edit the mapping, if necessary.

10. If necessary, rematerialize the target tables. After materialization completes, create new restart tokens.
11. Enable update activity on the source table again.
12. Restart the PowerCenter workflows.

If the table definition changes affected columns of CDC interest or you needed to edit the mapping, cold start the session. Otherwise, warm start the session.

## Reconfiguring a Partitioned Database or Database Partition Group

In a DB2 for Linux, UNIX, and Windows partitioned database environment, you might need to perform some reconfiguration tasks.

Common tasks are:

- Add a new partition to a partitioned database, or drop an existing partition. Then reconfigure the database partition group or groups to reflect the change.
- Reconfigure a database partition group by adding or removing existing partitions.

Typically, after making these types of changes, you run the DB2 REDISTRIBUTE DATABASE PARTITION GROUP command to redistribute table data among the partitions in the updated database partition group.

If PowerExchange change data capture is active in the partitioned database environment, you must use the following procedure to properly resume change data capture after making the reconfiguration changes.

### Adding or Dropping Database Partitions

Use this procedure to create a new partition in a partitioned database or to drop an existing partition, and then update the appropriate database partition group for the change.

1. Stop updates to the source tables.
2. Verify that any change data that was captured has completed extraction processing.
3. In PowerCenter, stop all CDC sessions that extract change data for the tables in the partitioned database instance.
4. For each table for which the DATA CAPTURE CHANGES clause is specified, specify DATA CAPTURE NONE.

**Note:** This step temporarily disables DB2 capture of changes to its log files. If you do not perform this step, DB2 records the data redistribution changes that result from the REDISTRIBUTE command as regular change data activity.

5. Execute the SQL for adding the new database partition or for dropping an existing partition.
6. Execute the ALTER DATABASE PARTITION GROUP SQL to add the new partition to or remove the dropped partition from the appropriate database partition group.
7. Run the DB2 REDISTRIBUTE DATABASE PARTITION GROUP command to redistribute table data among the partitions in the altered database partition group.
8. Back up the PowerExchange capture catalog table.
9. Run the PowerExchange DTLUCUDB SNAPUPDT command. Set the REPLACE option set to Y. This step updates the PowerExchange capture catalog table to reflect the reconfigured partitioned database.  
**Tip:** Informatica recommends that you first perform a test run with the REPLACE option set to N.
10. For each table for which you specified DATA CAPTURE NONE in step 2, reinstate the DATA CAPTURE CHANGES clause.
11. Restart the PowerCenter CDC sessions to resume extraction processing.

## Reconfiguring a Database Partition Group

Use this procedure to add a partition to or remove a partition from a database partition group without changing the partitioning of the partitioned database instance.

1. Stop updates to the source tables.
2. Verify that any change data that was captured has completed extraction processing.
3. In PowerCenter, stop all CDC sessions that extract change data for the tables in the partitioned database instance.
4. For each table for which the DATA CAPTURE CHANGES clause is specified, specify DATA CAPTURE NONE.

**Note:** This step temporarily disables DB2 capture of changes to its log files. If you do not perform this step, DB2 records the data redistribution changes that result from the REDISTRIBUTE command as regular change data activity.

5. Execute the ALTER DATABASE PARTITION GROUP SQL to add the new partition to or remove the dropped partition from the appropriate database partition group.
6. Run the DB2 REDISTRIBUTE DATABASE PARTITION GROUP command to redistribute table data among the partitions in the altered database partition group.
7. For each table for which you specified DATA CAPTURE NONE in step 2, reinstate the DATA CAPTURE CHANGES clause.
8. Restart the PowerCenter CDC sessions to resume extraction processing.

# DB2 for Linux, UNIX, and Windows CDC Troubleshooting

If you encounter the following issue when running DB2 for Linux, UNIX, and Windows CDC, attempt the solution that is described. If you cannot resolve the problem, contact Informatica Global Customer Support.

## Workaround for SQL1224 Error on AIX

On AIX systems only, you might receive the following PowerExchange message for a DB2 SQL1224 error when you connect locally to a DB2 database that has multiple other local connections:

```
PWX-20604 State=08001, Code=-1224, Msg=[IBM][CLI Driver] SQL1224N A database agent
could not be started to service a request, or was terminated as a result of a database
system shutdown or a force command.  SQLSTATE=55032.
```

To circumvent this problem, implement a loopback TCP/IP connection for the local DB2 database. The database can then function as a remote client that uses TCP/IP instead of interprocess communications (IPC) over shared memory.

To implement a loopback connection without changing the database alias that users enter for database connection, issue the following DB2 commands:

```
db2 catalog tcpip node node_name1 remote server_name1 server port_number1
db2 uncatalog database database_name1
db2 catalog database database_name1 at node node_name1
db2 catalog database database_name1 as database_alias1
db2 catalog database database_alias1 as database_name1 at node node_name1
```

For more information about these commands, see your IBM DB2 documentation.

## CHAPTER 5

# Microsoft SQL Server CDC

This chapter includes the following topics:

- [Microsoft SQL Server CDC Overview, 94](#)
- [Planning for SQL Server CDC, 95](#)
- [Configuring SQL Server for CDC, 99](#)
- [Configuring PowerExchange for SQL Server CDC, 99](#)
- [Managing SQL Server CDC, 106](#)

## Microsoft SQL Server CDC Overview

PowerExchange uses SQL Server transactional replication to capture change data from SQL Server distribution databases. PowerExchange uses the PowerExchange Client for PowerCenter (PWXPC) to coordinate with PowerCenter to move the captured change data to one or more targets.

For CDC to work, you must enable SQL Server Replication on the system from which change data is to be captured. If your database has a high volume of change activity, you should use a distributed server as the host of the distribution database.

To configure CDC in PowerExchange, you must define a capture registration for each source table. In the capture registration, you can select a subset of columns for which to capture data. PowerExchange generates a corresponding extraction map.

If you want to use the PowerExchange Logger for Linux, UNIX, and Windows to capture change data and write it to PowerExchange Logger log files, configure the PowerExchange Logger. The change data is then extracted from the PowerExchange Logger log files. Benefits of the PowerExchange Logger include fewer database accesses and faster CDC restart.

PowerExchange works with PowerCenter to extract change data from the SQL Server distribution database or PowerExchange Logger log files and load that data to one or more targets.

### RELATED TOPICS:

- [“Planning for SQL Server CDC” on page 95](#)
- [“Configuring SQL Server for CDC” on page 99](#)
- [“Configuring PowerExchange for SQL Server CDC” on page 99](#)
- [“Managing SQL Server CDC” on page 106](#)
- [“PowerExchange Logger for Linux, UNIX, and Windows” on page 35](#)
- [“Introduction to Change Data Extraction” on page 208](#)

# Planning for SQL Server CDC

Before you configure SQL Server change data capture (CDC), verify that the following prerequisites and user authority requirements are met. Also, review the restrictions so that you can properly configure CDC.

## SQL Server CDC Prerequisites

PowerExchange CDC has the following SQL Server prerequisites:

- PowerExchange CDC requires an edition of Microsoft SQL Server 2008 or later that supports transactional replication. You must configure and enable transactional replication on the source system to do CDC.
- PowerExchange uses SQL Server 2008 Server Management Objects (SMO) to manage SQL Server publications. SMO is a .NET Framework object model that applications can use to manage and administer SQL Server objects and services. You must install SQL Server 2008 SMO and some related packages before you create or edit capture registrations. For your convenience, PowerExchange provides the installation packages for these SQL Server objects and packages under the *PowerExchange\_base\_installation\_directory\packages\mssqlInstalls/2008* directory.
- The Microsoft SQL Server Agent and Log Reader Agent must be running on the Windows machine from which change data is extracted. Usually, the SQL Server Agent remains running after it is initially started. For more information, see your SQL Server documentation.
- Each source table in the distribution database must have a primary key.
- If the PowerExchange Navigator is installed on a Windows machine than is different from the one where the Microsoft SQL Server source runs, you must install the SQL Server client components on the PowerExchange Navigator machine.
- Ensure that the PowerExchange Listener runs on the Windows system that contains the Microsoft SQL Server source or on another Windows system. The PowerExchange Listener cannot run on a Linux or UNIX system.
- If you use the PowerExchange Logger for Linux, UNIX, and Windows, you can run it on the system where the PowerExchange Listener runs or on a Linux, UNIX, or Windows system for which you configure remote logging.

## Required User Authority for SQL Server CDC

Verify that you have the proper authority level to complete registration and SQL Server configuration tasks.

The following user authority levels are required:

- To enable transactional replication on the publication database, you must have System Administrator authority, or the SA role.
- To create registration groups and capture registrations from the PowerExchange Navigator, you must have DB\_OWNER authority.

The first time you add a registration group in the PowerExchange Navigator, PowerExchange creates a publication named "PowerExchange Change Capture." If you did not previously enable transactional replication on the publication database, PowerExchange also attempts to enable transactional replication. In this situation, you must have System Administrator authority, or the SA role, to enable transactional replication and create the first registration group. Thereafter, users with DB\_OWNER authority can add registration groups under their own user IDs and passwords and add capture registrations.

- To run change data extractions against a SQL Server distribution database, you must have read access to that database.

If you do not specify a user ID and password when creating a registration group, the PowerExchange Navigator and the extraction processes attempt to use your Windows user ID and password to connect to the SQL Server distribution database.

## SQL Server Datatypes Supported for CDC

PowerExchange supports most SQL Server datatypes for CDC, with some exceptions.

The following table indicates the SQL Server datatypes that PowerExchange supports and does not support for CDC:

Datatype	Supported for CDC?	Comments
bigint	Yes	-
binary	Yes	-
bit	Yes	-
char	Yes	-
date	Yes	In PowerCenter, when you import source metadata from PowerExchange to create a source definition, PowerExchange converts date columns to timestamp columns. This conversion is for consistency with PowerCenter datatype handling.
datetime	Yes	-
datetime2	Yes	-
datetimeoffset	Yes	PowerCenter treats this datatype as varchar.
decimal	Yes	-
float	Yes	-
geography	No	-
geometry	No	-
hierarchyid	No	-
image <sup>1</sup>	No	Use varbinary(MAX) instead.
int	Yes	-
money	Yes	-
nchar	Yes	-
ntext <sup>1</sup>	No	Use nvarchar(MAX) instead.
numeric	Yes	-
nvarchar	Yes	-



Datatype	Supported for CDC?	Comments
real	Yes	-
smalldatetime	Yes	-
smallint	Yes	-
smallmoney	Yes	-
sql_variant	No	PowerExchange does not capture change data for sql_variant columns but does capture change data for other columns in the same table.
text <sup>1</sup>	No	Use varchar(MAX) instead.
time	Yes	-
timestamp	Yes	-
tinyint	Yes	-
uniqueidentifier	Yes	PowerCenter imports the uniqueidentifier datatype as a varchar datatype of 38 characters.
user-defined datatypes (UDTs)	Yes	PowerExchange treats a UDT in the same way as the datatype on which the UDT is based.
varbinary	Yes	-
varchar	Yes	-
xml	Yes	PowerExchange treats this datatype as varchar(MAX).
<p>1. PowerExchange might not be able to capture change data for columns that have the image, ntext, or text datatype because of SQL Server transactional replication restrictions on these column types. Instead, use the alternative datatypes that Microsoft recommends, as indicated in the Comments column.</p>		

## SQL Server CDC Operational Considerations

PowerExchange for SQL Server CDC has the following operational considerations:

- PowerExchange can capture change data from SQL Server distribution databases for which Transparent Data Encryption (TDE) is enabled. No special configuration tasks are required.
- PowerExchange does not capture change data for SQL Server system tables.
- The maximum length of a row for which PowerExchange can capture and process change data is 128,000 bytes.
- PowerExchange does not capture the user ID that is associated with the original transaction that updated the database.
- The timestamp that PowerExchange records for each captured change indicates when the change was captured, not when the original transaction occurred.

- PowerExchange does not capture change data for derived columns that are not persisted. SQL Server computes values for these columns at run-time based on an expression but does not store the values in a table.
- SQL Server publishes *deferred updates* to SQL Server tables as DELETEs followed by INSERTs rather than as UPDATEs. Consequently, PowerExchange propagates deferred updates as DELETEs followed by INSERTs, even if you select **AI** for the **Image Type** attribute in the CDC connection. PowerExchange does not include before image (BI) and change indicator (CI) information in DELETE and INSERT records. If you must capture a deferred update as an UPDATE for business reasons, set the SQL Server 8207 trace flag. This flag causes the SQL Server Replication Log Reader to combine the DELETE and INSERT pair into a single UPDATE. For more information about SQL Server processing of deferred updates and the SQL Server 8207 trace flag, see the SQL Server documentation
- PowerExchange does not support the use of local aliases when connecting to SQL Server and creating publications at registration creation.
- If you need to switch the status of multiple SQL Server capture registrations from active to inactive or from inactive to active, use the DTLUCBRG utility with the MSSOPTS UPDATESTATUS parameter. This optional parameter enables you to switch the status of many registrations in one operation and regenerate the associated SQL Server publications.
- You can run the PowerExchange Listener on a Windows system other than the Windows system where the SQL Server source distribution database runs. On the PowerExchange Listener system, you must install the Microsoft SQL Server Native Client and define a MSQL CAPI CONNECTION statement that provides connection information for the SQL Server source distribution database in the dbmover configuration file. Also, in the dbmover configuration files on the PowerCenter Integration Service machine and on the PowerExchange Navigation system, define a NODE statement that points to the PowerExchange Listener system.

**Tip:** To install the SQL Server 2008 Native Client, use the sqlncli.msi file in the `PowerExchange_root_installation\packages\mssqlInstalls\2008` directory.

## Extracting Data for Multiple Publication Databases

If you plan to extract data from a Microsoft SQL Server distribution database that contains information for articles in multiple publication databases in a single pass, you must set some parameters for extraction processing to succeed.

This requirement applies to extraction of data directly from the change stream in real time extraction mode or from PowerExchange Logger for Linux, UNIX, and Windows log files in continuous extraction mode.

Specify the following parameters:

- In the MSQL CAPI\_CONNECTION statement in the dbmover configuration, verify that the MULTIPUB parameter is set to Y, the default setting. If you do not use this setting, the extraction will fail with message PWX-15757.  
**Tip:** If you plan to extract data for single publication database, set this parameter to N for more efficient extraction processing.
- If you use the PowerExchange Logger for Linux, UNIX, and Windows, set the following parameters:
  - In the PowerExchange Logger configuration file, pwxcccl, ensure that you define the DBID, DISTDB, and DISTSRV parameters. These parameters are usually required for PowerExchange Logger processing of SQL Server sources.
  - In the PowerCenter PWX MSSQL CDC Change or Real Time connection definition, specify the PowerExchange Logger DBID parameter value for the **Logger DBID** attribute.
  - To perform a database row test, specify the PowerExchange Logger DBID parameter value in the **MSS LUW DBId** field in the CAPXRT Advanced Parameters dialog box.

# Configuring SQL Server for CDC

You must perform a few configuration tasks to prepare SQL Server for PowerExchange change data capture (CDC).

If your SQL Server tables have a high level of update activity, use a distributed server as the host of the distribution database from which to capture change data. This practice prevents competition between PowerExchange CDC and your production database for CPU use and disk storage.

PowerExchange uses SQL Server Management Objects (SMO) to manage SQL Server publications. SMO is a .NET Framework object model that applications can use to manage and administer SQL Server objects and services. You must install SMO and some related packages before you create or edit capture registrations. PowerExchange provides the 64-bit installation packages for the required SMO objects and packages for your convenience.

1. On the system where you run the PowerExchange Navigator and PowerExchange utilities, install the following SQL Server 2008 SMO objects and packages by using the installer .msi files in the *PowerExchange\_root\_installation\packages\mssqlInstalls\2008* directory:
  - Microsoft SQL Server 2008 Management Objects
  - Microsoft SQL Server 2008 Replication Management Objects
  - Microsoft SQL Server 2008 Native Client
  - Microsoft SQL Server System CLR Types

**Note:** You must install the Microsoft SQL Server 2008 Replication Management Objects even if you use SQL Server 2012 or SQL Server 2014.
2. Start the SQL Server Agent and Log Reader Agent if they are not running. For more information, see your Microsoft SQL Server documentation.
3. Configure and enable SQL Server transactional replication on the publication database. For more information, see your Microsoft SQL Server documentation.

**Tip:** The default transactional retention period at the Distributor is 72 hours. If you are using the PowerExchange Logger, accept this default retention period. If you are not using the PowerExchange Logger, Informatica recommends that you increase the retention period to 14 days. However, you might need to decrease the value if you have a high volume of transactions or space constraints.
4. Verify that each source table in the distribution database has a primary key.

## Configuring PowerExchange for SQL Server CDC

The tasks that you perform to configure PowerExchange for change data capture (CDC) depend on whether you want to use the PowerExchange Logger for Linux, UNIX, and Windows and the extraction mode you plan to use.

### RELATED TOPICS:

- [“Customizing the dbmover Configuration File for SQL Server CDC” on page 101](#)
- [“Configuring PowerExchange CDC with the PowerExchange Logger” on page 100](#)
- [“Configuring PowerExchange CDC without the PowerExchange Logger” on page 100](#)

## Configuring PowerExchange CDC without the PowerExchange Logger

If you plan to run extractions in real-time extraction mode and *not* use the PowerExchange Logger for Linux, UNIX, and Windows, complete the following tasks to configure PowerExchange CDC:

1. When you configure the dbmover.cfg file, define the following statements:
  - CAPT\_PATH
  - CAPT\_XTRA
  - MSQL CAPI\_CONNECTION
2. In the PowerExchange Navigator, create a capture registration for each SQL Server source table. The PowerExchange Navigator generates a corresponding extraction map.

**Tip:** Set the **Condense** option to **Part** even though you do not plan to use the PowerExchange Logger, unless you have a particular reason not to do so. This practice prevents having to change the capture registrations later if you decide to use the PowerExchange Logger. You might want to set the **Condense** option to **None** if you run both real-time and continuous extractions against tables defined by the same capture registrations and do not want the PowerExchange Logger to capture change data for certain registered tables.

If capture registrations already exist for these tables, delete the existing registrations and extraction maps and create new ones.

The PowerExchange Navigator generates a corresponding extraction map for each capture registration.

3. Activate the capture registrations. Usually, you do this task after materializing the targets.

**Next Step:** Configure and start extractions. You must use real-time extraction mode.

## Configuring PowerExchange CDC with the PowerExchange Logger

If you plan to run extractions in batch or continuous extraction mode and use the PowerExchange Logger for Linux, UNIX, and Windows, complete this procedure to configure PowerExchange CDC.

1. When you configure the dbmover configuration file, define the following statements:
  - CAPT\_PATH
  - CAPT\_XTRA
  - MSQL CAPI\_CONNECTION
  - CAPX CAPI\_CONNECTION (for continuous extraction mode only)
2. Configure the pwxcl configuration file for the PowerExchange Logger.
3. In the PowerExchange Navigator, create a capture registration for each SQL Server source table.

You must set the **Condense** option to **Part**. If capture registrations already exist for these tables, you can edit the **Condense** option without affecting the registration version.

The PowerExchange Navigator generates a corresponding extraction map.
4. Start the PowerExchange Logger.
5. Activate the capture registrations.

Usually, you do this task after materializing the targets.

**Next Step:** Configure and start extractions. You can use either batch extraction mode or continuous extraction mode.

# Customizing the dbmover Configuration File for SQL Server CDC

In the dbmover configuration file, add statements for Microsoft SQL Server CDC.

You must include the following statements for SQL Server CDC:

## **CAPT\_PATH statement**

Path to the local directory on a Linux, UNIX, or Windows system that contains the control files for CDC.

These files are: the CCT file for capture registrations, the CDEP file for application names that are used for ODBC extractions, and the CDCT file for the PowerExchange Logger for Linux, UNIX, and Windows.

## **CAPT\_XTRA statement**

Path to the local directory on a Linux, UNIX, or Windows system that stores extraction maps.

## **MSQL CAPI\_CONNECTION statement**

A named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and control CDC processing for SQL Server sources.

Add this statement to the dbmover.cfg file on the system where SQL Server capture registrations are stored. This location corresponds to the **Location** node that you specify when defining a registration group. Usually, this location is where the source database resides.

If you plan to use continuous extraction mode, you must also include the following statement:

## **CAPX CAPI\_CONNECTION statement**

A named set of parameters that the CAPI uses for continuous extraction of change data from PowerExchange Logger for Linux, UNIX, and Windows log files.

Also, Informatica recommends including the LOGPATH and TRACING statements to make finding messages easier. The LOGPATH statement defines a directory specifically for PowerExchange message log files, and the TRACING statement enables PowerExchange to create an alternative set of message log files for each PowerExchange process.

For more information about all DBMOVER statements, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“CAPI\\_CONNECTION - CAPX Statement” on page 57](#)
- [“CAPI\\_CONNECTION - MSQL Statement” on page 102](#)
- [“CAPT\\_PATH Statement” on page 30](#)
- [“CAPT\\_XTRA Statement” on page 30](#)

## Example dbmover Statements for SQL Server

The following statements are typical of those included in a dmover.cfg for SQL Server CDC:

```
LOGPATH="C:\Informatica\PowerExchangeVnnn\Logs"
CAPT_XTRA="C:\Informatica\PowerExchangeVnnn\Capture\camaps"
CAPT_PATH="C:\Informatica\PowerExchangeVnnn\Capture"
CAPI_CONN_NAME=CAPIMSSC
CAPI_CONNECTION=(NAME=CAPIMSSC
                  ,TYPE=(MSQL,DISTSRV=AUX159908\PWXPC
                        ,DISTDB=distribution
                        ,RSTRADV=30))
```

**Note:** You must use non-curly double quotation marks around values that include a space.

## CAPI\_CONNECTION - MSQL Statement

The MSQL CAPI\_CONNECTION statement specifies a named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and control CDC processing for Microsoft SQL Server sources.

**Operating Systems:** Windows

**Data Sources:** Microsoft SQL Server

**Required:** Yes for Microsoft SQL Server CDC

**Syntax:**

```
CAPI_CONNECTION=( [DLLTRACE=trace_id]
                  ,NAME=capi_connection_name
                  [,TRACE=trace_name]
                  ,TYPE=(MSQL
                        ,DISTDB=distribution_database
                        ,DISTRV=distribution_server
                        [,BATCHSIZE=number]
                        [,DWFLAGS={flag1flag2flag3flag4|NNNN}]
                        [,ENABLELWM={N|Y}]
                        [,EOF={N|Y}]
                        [,MEMCACHE={cache_size|256}]
                        [,MULTIPUB={N|Y}]
                        [,POLWAIT={seconds|1}]
                        [,RSTRADV=seconds]
                        [,UIDFMT={DBNAME|NONE}]
                        )
                  )
```

**Parameters:**

**DLLTRACE=trace\_id**

Optional. User-defined name of the TRACE statement that activates internal DLL tracing for this CAPI. Specify this parameter only at the direction of Informatica Global Customer Support.

**NAME=capi\_connection\_name**

Required. Unique user-defined name for this CAPI\_CONNECTION statement.

Maximum length is eight alphanumeric characters.

**TRACE=trace\_name**

Optional. User-defined name of the TRACE statement that activates the common CAPI tracing. Specify this parameter only at the direction of Informatica Global Customer Support.

**TYPE=(MSQL, ... )**

Required. Type of CAPI\_CONNECTION statement. For Microsoft SQL Server sources, this value must be MSQL.

**DISTDB=distribution\_database\_name**

Required. The name of the distribution database.

**DISTRV=distribution\_database\_server**

Required. The network name of the server that hosts the distribution database.

**Important:** This name is different from the network name of the SQL Server instance if the distribution database resides on a different server.

**BATCHSIZE=number**

Optional. The number of rows from which PowerExchange captures change data before closing the cursor and then reopening it. This parameter allows resources to be released periodically to reduce the capture processing load on system memory and to reduce temporary tables in the tempdb database. Valid values are 0 through 2147483647. No default is provided.

Specify this parameter only at the direction of Informatica Global Customer Support. It can degrade CDC performance because PowerExchange issues the data read query more often.

**DWFLAGS={flag1flag2flag3flag4|NNNN}**

Optional. Series of four positional parameters that control whether processing stops or continues when data loss, truncation, schema changes, or unrecognized transaction log records occur.

Specify this statement only at the direction of Informatica Global Customer Support.

Enter the following positional parameters:

- *flag1*. Controls whether PowerExchange stops a change data extraction when PowerExchange retrieves data of an unexpected length from the distribution database. Enter Y to continue processing or enter N to stop processing.
- *flag2*. Controls whether PowerExchange stops a change data extraction when it detects a schema change. Enter Y to continue processing or enter N to stop processing.
- *flag3*. Controls whether PowerExchange stops a change data extraction when PowerExchange does not find the requested start sequence in the transaction log. Enter Y to continue processing or enter N to stop processing.
- *flag4*. Controls whether PowerExchange stops a change data extraction when PowerExchange finds an unrecognized record in the transaction log. Enter Y to continue processing after error message PWX-15742 or enter N to stop processing.

Default is NNNN, which indicates none of the parameters are set.

**ENABLELWM={N|Y}**

Optional. When you use the PowerExchange Logger for Linux, UNIX, and Windows, controls whether the PowerExchange consumer API (CAPI) connection process deletes data read from the SQL Server distribution database after the data has been hardened to PowerExchange Logger log files or after the PowerExchange publication expiry time has elapsed. You can use this parameter to improve distribution database performance and to prevent the distribution database from growing too large in size when the PowerExchange Logger is in use.

Enter one of the following options:

- **N**. The distribution database cleanup job deletes data from the distribution database after the expiry time for the PowerExchange publications elapses. This option might degrade the performance of the distribution-database cleanup job and cause excessive growth of the distribution database.

- **Y.** The CAPI connection process deletes processed data from the distribution database after the data has been hardened to the PowerExchange Logger log files. After a log file switch, the PowerExchange Logger sends a low water marker to the CAPI connection process to identify the last end UOW prior to the file switch. At the end of the next capture cycle, after the CAPI connection process has read to the end of the available data in the distribution database, the CAPI deletes all of the processed data for the PowerExchange publications up to and including the low water mark data from the distribution.dbo.MSrepl\_commands table in the distribution database.

**Note:** The user ID under which the PowerExchange Logger runs must have delete authority on the MSrepl\_commands table.

This option can help improve distribution-database performance and control distribution-database size. However, if the SQL Server Log Reader Agent is writing very large UOWs to the distribution database when the CAPI connection processes the low water mark data, the performance of the distribution database might be temporarily degraded because the CAPI connection process must wait for a lock on the MSrepl\_commands table.

**Note:** If you run multiple extractions against a single distribution database for different publication databases and use ENABLELWM=Y for one CAPI connection and ENABLELWM=N with a RSTRADV value for another CAPI connection, PowerExchange might issue error message PWX-15756 for the connection with ENABLELWM=N. The message incorrectly reports that change data has been lost. To suppress this error, add the DWFLAGS=NNYN parameter to MSQL CAPI\_CONNECTION statement.

Default is N.

#### **EOF={N|Y}**

Optional. Controls whether PowerExchange stops change data extractions when the end-of-log (EOL) is reached.

Enter one of the following options:

- **N.** PowerExchange does not stop change data extractions when the EOL is reached.
- **Y.** PowerExchange stops change data extractions when the EOL is reached.

Default is N.

Because this parameter affects all users of the MSQL CAPI\_CONNECTION statement, Informatica recommends that you use one of the following alternative methods to stop change data extractions at the EOL:

- For CDC sessions that use real-time extraction mode, enter 0 for the **Idle Time** attribute of the PWX MSSQL CDC Real Time application connection.
- For the PowerExchange Logger for Linux, UNIX, and Windows, enter 1 for the COLL\_END\_LOG statement in the pwxcl.cfg configuration file.
- For CDC sessions that use ODBC connections, enter 0 for the WAITTIME parameter in the ODBC data source.

#### **MEMCACHE={cache\_size|256}**

The maximum size, in kilobytes, of the memory cache that stores change data for a single SQL operation that is captured from the SQL Server distribution database. The memory cache stores the full row image, which can include both the before image and after image and any LOB data.

Valid values are 0 through 2147483647. Default is 256. If you enter 0, the default value is used.



**MULTIPUB={N|Y}**

Optional. Indicates whether you capture change data from the distribution database for articles in a single publication database or in multiple publication databases. This option can affect the performance of CDC processing in real time extraction mode and in continuous extraction mode with the PowerExchange Logger for Linux, UNIX, and Windows.

Enter one of the following options:

- **N.** Specify this option if you capture change data for articles in a single publication database. Informatica recommends this option in this scenario because it causes PowerExchange to extract changes much more efficiently. It can also help reduce resource usage.
- **Y.** Use this option to extract change data for articles in multiple publication databases in a single CDC session or in a single PowerExchange Logger for Linux, UNIX, and Windows pass. If you do not use this option in this scenario, extraction processing fails with message PWX-15757.

This option might cause change records to be written to the distribution database more slowly. To improve performance, add the following index to the distribution database:

```
USE [distribution]
GO
/***** Object: Index [IX_MSrepl_transactions] Script Date: 03/31/2012
11:56:07 *****/
CREATE NONCLUSTERED INDEX [IX_MSrepl_transactions] ON [dbo].[MSrepl_transactions]
(
    [entry_time] ASC,
    [publisher_database_id] ASC,
    [xact_seqno] ASC,
    [xact_id] ASC
)WITH (PAD_INDEX = OFF, STATISTICS_NORECOMPUTE = ON, SORT_IN_TEMPDB = OFF,
IGNORE_DUP_KEY = OFF, DROP_EXISTING = OFF, ONLINE = OFF, ALLOW_ROW_LOCKS =
ON, ALLOW_PAGE_LOCKS = ON) ON [PRIMARY]
GO
```

Default is Y.

You can change the MULTIPUB setting after capturing changes. To maintain data integrity, follow the procedure for changing the MULTIPUB parameter setting in the *PowerExchange CDC Guide for Linux, UNIX, and Windows*. If you use the PowerExchange Logger for Linux, UNIX, and Windows and change the setting from Y to N, you must cold start the PowerExchange Logger.

**POLWAIT={seconds|1}**

Optional. The maximum number of seconds that PowerExchange waits after reaching the end of log before polling the source database for more change data.

For Microsoft SQL Server sources, the polling frequency also depends on the PowerExchange Logger NO\_DATA\_WAIT2 parameter, or if you do not use the PowerExchange Logger, the polling frequency depends on the PWX Latency attribute on the PWX CDC application connection. If the NO\_DATA\_WAIT2 or PWX Latency value is less than the POLWAIT value, the lesser value takes precedence. In this case, PowerExchange polls the source more frequently than expected based on the POLWAIT parameter only.

For the *seconds* variable, enter a number from 1 through 2147483647. Default is 1.

**RSTRADV=seconds**

Time interval, in seconds, that PowerExchange waits before advancing restart and sequence tokens for a registered data source during periods when UOWs do not include any changes of interest for the data source. When the wait interval expires, PowerExchange returns the next committed "empty UOW," which includes only updated restart information.

PowerExchange resets the wait interval to 0 when one of the following events occur:

- PowerExchange completes processing a UOW that includes changes of interest.
- PowerExchange returns an empty UOW because the wait interval expired without PowerExchange receiving any changes of interest.

For the *seconds* variable, enter a number from 0 through 86400. No default is provided.

For example, if you specify 5, PowerExchange waits 5 seconds after it completes processing the last UOW or after the previous wait interval expires. Then PowerExchange returns the next committed empty UOW that includes the updated restart information and resets the wait interval to 0.

If you do not specify RSTRADV, PowerExchange does not advance restart and sequence tokens for a registered source during periods when PowerExchange receives no changes of interest. In this case, when PowerExchange warm starts, it reads all changes, including those not of interest for CDC, from the restart point.

**Warning:** A value of 0 can degrade performance. In addition to the UOWs that contain changes for registered sources of interest, PowerExchange returns an empty UOW for every UOW that does not contain changes for the registered sources of interest.

#### **UIDFMT={DBNAME |NONE}**

Optional. Controls the type of value that PowerExchange uses to populate the generated DTL\_\_CAPXUSER column in each change record. Options are:

- **DBNAME.** Returns the Microsoft SQL Server publication database name.
- **NONE.** Returns a null because a user ID is not available.

Default is NONE.

## Managing SQL Server CDC

You might need to stop CDC for source tables occasionally, for example, to change the table definitions.

### Disabling Publication of Change Data for a SQL Server Source

You can disable publication of change data for a SQL Server source. For example, you might disable publication to perform some database maintenance, change the table definition, or avoid capturing unwanted changes.

- Open the capture registration for the table, and change the **Status** setting from **Active** to **History**.

This action disables publication of the SQL Server article for the table to the distribution database, which causes change capture to stop.

**Warning:** After the registration status is set to **History**, you cannot activate the registration for CDC use again.

### Changing a SQL Server Source Table Definition

If you change the definition of a SQL Server source table that is registered for change data capture, use this procedure to enable PowerExchange to use the updated table definition and preserve access to previously captured data. Table definition changes include adding, altering, or dropping columns.

**Tip:** If you no longer need to capture change data from a column in a table, you can remove the column from the extraction map without changing the capture registration. Change data for that column is still captured but is not extracted.

To change a SQL Server source table definition:

1. Stop DELETE, INSERT, and UPDATE activity against the table.
2. Verify that any change data that was captured under the previous table definition has completed extraction processing. Then stop all workflows that extract change data for the table.
3. Delete the capture registration and extraction map.
4. Use DDL to change the table definition in SQL Server.
5. In the PowerExchange Navigator, create a new capture registration that reflects the metadata changes and set its status to **Active**. PowerExchange creates a corresponding extraction map.  
The newly activated capture registration becomes eligible for change data capture.
6. If necessary, change the target table definition to reflect the source table metadata changes.
7. In the PowerCenter Designer, import the altered source and target definitions. Edit the mapping if necessary.
8. If necessary, rematerialize the target tables. After materialization completes, create new restart tokens.
9. Create new restart tokens for the altered table.
10. Re-enable DELETE, INSERT, and UPDATE activity against the table.
11. Cold start the extraction workflows.

## Changing the MULTIPUB Parameter Setting After Running Extractions

After you run change data extraction processing, you can change the MULTIPUB parameter setting in the MSQL CAPI\_CONNECTION statement. You might need to do this task if you add or remove publication databases that include sources of CDC interest. To preserve data integrity, you must use the proper procedure.

The MULTIPUB parameter indicates whether you extract data for articles in a single publication database or in multiple publications. For a single publication database, Informatica recommends that you set MULTIPUB to N so that PowerExchange can use more efficient extraction processing. For multiple publications, you must set MULTIPUB to Y, the default setting. This parameter applies to real time extractions directly from the change stream, and PowerExchange Logger for Linux, UNIX, and Windows extractions in continuous extraction mode.

### To switch the MULTIPUB setting from Y to N:

Use this procedure to switch the MULTIPUB from the default of Y to N. If you use the PowerExchange Logger for Linux, UNIX, and Windows, you must cold start it after making this change.

1. Stop extraction workflows that process the SQL Server distribution database and that are running in real-time extraction mode or continuous extraction mode.
2. If you use the PowerExchange Logger for Linux, UNIX, and Windows, stop the PowerExchange Logger.
3. In the dbmover configuration file, edit the MSQL CAPI\_CONNECTION statement to switch the MULTIPUB parameter setting from Y to N.
4. Cold start the PowerExchange Logger.
5. Restart the extraction workflows.

**Note:** The sequence tokens no longer include a timestamp.

### To switch the MULTIPUB setting from N to Y:

Use this procedure to switch the MULTIPUB from N back to Y. If you use the PowerExchange Logger for Linux, UNIX, and Windows, you do not need to cold start it after making this change.

1. Stop DELETE, INSERT, and UPDATE activity on the SQL Server source tables.
2. Wait for the extraction workflows to reach the end of log and then stop them.
3. In the dbmover configuration file, edit the MSQL CAPI\_CONNECTION statement to switch the MULTIPUB parameter setting from Y to N.
4. To help avoid performance degradation, define the following index on the distribution database:

```
USE [distribution]
GO
/***** Object: Index [IX_MSrepl_transactions] Script Date: 03/31/2012 11:56:07
*****/
CREATE NONCLUSTERED INDEX [IX_MSrepl_transactions] ON [dbo].[MSrepl_transactions]
(
    [entry_time] ASC,
    [publisher_database_id] ASC,
    [xact_seqno] ASC,
    [xact_id] ASC
)WITH (PAD_INDEX = OFF, STATISTICS_NORECOMPUTE = ON, SORT_IN_TEMPDB = OFF,
IGNORE_DUP_KEY = OFF, DROP_EXISTING = OFF, ONLINE = OFF, ALLOW_ROW_LOCKS = ON,
ALLOW_PAGE_LOCKS = ON) ON [PRIMARY]
GO
```

5. To get the current restart tokens for the end of log, use one of the following methods:
  - Run the DTLUAPPL utility with the GENERATE RSTKKN option.
  - In the PowerExchange Navigator, perform a database row test with a SELECT CURRENT\_RESTART SQL statement.
  - Specify the CURRENT\_RESTART option on the RESTART1 and RESTART2 special override statements in the PWXPC restart token file. When the CDC session runs, PWXPC requests that PowerExchange provide restart tokens for the current EOL. PWXPC uses this restart information to locate the extraction start point.
6. Add the current restart tokens for the extractions to the restart token file.
7. Allow DELETE, INSERT, and UPDATE activity to resume on the SQL Server tables.
8. Cold start the extraction workflows.

**Note:** PowerExchange adds a timestamp in the sequence token to combine the data from multiple publication databases during extraction processing.

## CHAPTER 6

# Oracle CDC with LogMiner

This chapter includes the following topics:

- [Oracle CDC with LogMiner Overview, 109](#)
- [Planning for Oracle CDC with LogMiner, 110](#)
- [Oracle Configuration for CDC with LogMiner, 114](#)
- [PowerExchange Configuration for CDC with LogMiner, 120](#)
- [Managing Oracle CDC with LogMiner, 137](#)

## Oracle CDC with LogMiner Overview

PowerExchange Oracle CDC with LogMiner uses Oracle LogMiner to read change data from Oracle archived redo logs. PowerExchange then makes the data available to PowerCenter CDC sessions for propagation to targets.

Oracle CDC with LogMiner is an alternative to PowerExchange Express CDC for Oracle.

**Important:** You cannot use both PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle in the same PowerExchange instance with the same dbmover and pwxcccl configuration files. In many environments, PowerExchange Express CDC for Oracle can provide faster CDC processing and avoid LogMiner reinitialization issues.

Oracle CDC with LogMiner supports Oracle Real Application Cluster (RAC) and Automatic Storage Management (ASM) environments.

Oracle CDC with LogMiner uses the UOW Cleanser component to reconstruct UOWs from archived redo logs into complete and consecutive UOWs that are in chronological order by end time.

To implement Oracle CDC with LogMiner, you need to perform configuration tasks in Oracle, PowerExchange, and PowerCenter.

In Oracle, ensure that ARCHIVELOG mode and global minimal supplemental logging are enabled so that change data can be retrieved from archived redo logs. Also, ensure that a copy of the Oracle online catalog exists in the archived log destination. PowerExchange requires a copy of the catalog to determine restart points for change data extraction processing.

In PowerExchange, define a capture registration for each source table. In the capture registration, you can select a subset of columns for which to capture data. PowerExchange generates a corresponding extraction map.

Use of the PowerExchange Logger for Linux, UNIX, and Windows is optional but strongly recommended. The PowerExchange Logger captures change data from Oracle logs and write only the successful units of work

(UOWs), in chronological order based on commit time, to PowerExchange Logger log files. The change data is then extracted from the PowerExchange Logger log files in either continuous extraction mode or batch extraction mode. Benefits of using the PowerExchange Logger include fewer database accesses, faster CDC restart, and no need to prolong retention of the Oracle redo files for change capture.

**Caution:** If you use real-time extraction mode without the PowerExchange Logger, PowerExchange starts a separate Oracle LogMiner session for each extraction session. Running multiple, concurrent sessions can significantly degrade performance of the system where LogMiner runs.

PowerExchange works with the PowerExchange Client for PowerCenter (PWXPC) and PowerCenter to extract change data from Oracle redo logs or PowerExchange Logger log files and load that data to one or more targets.

## RELATED TOPICS:

- [“Planning for Oracle CDC with LogMiner” on page 110](#)
- [“Oracle Configuration for CDC with LogMiner” on page 114](#)
- [“PowerExchange Configuration for CDC with LogMiner” on page 120](#)
- [“Managing Oracle CDC with LogMiner” on page 137](#)
- [“Express CDC for Oracle” on page 139](#)
- [“PowerExchange Logger for Linux, UNIX, and Windows” on page 35](#)
- [“Introduction to Change Data Extraction” on page 208](#)

# Planning for Oracle CDC with LogMiner

Before you configure Oracle CDC with LogMiner change data capture, review the following restrictions, requirements, and performance considerations.

## Oracle CDC with LogMiner Implementation Considerations

The following implementation considerations apply to PowerExchange Oracle CDC with LogMiner:

- Verify that a valid Oracle environment exists for the PowerExchange user. On Linux, and UNIX, the path to the Oracle client must be specified in the PATH and library path environment variables.
- The Oracle source instance must be running in ARCHIVELOG mode, and Oracle global minimal supplemental logging must be enabled.
- A copy of the Oracle catalog must exist in the Oracle archived logs.
- The COMPATIBLE initialization parameter in the init.ora file must be set to 9.2 or later.
- Oracle LogMiner continuous mining reads archived redo logs only from the directory to which they were originally written.
- If PowerExchange is not installed on the same machine as the Oracle instance, configure a TNS entry on the client machine with SERVER=DEDICATED in the CONNECT\_DATA section of the connect descriptor. This specification is also required if the network is configured for Multi-Threaded Server (MTS) mode.
- PowerExchange requires the Oracle Client binaries. When you install Oracle, the Client binaries are installed by default. To use SQL\*Net connectivity on a machine that does not have an installed Oracle instance, you must install the Oracle Client.
- The maximum length of a row for which PowerExchange can capture and process change data is 128,000 bytes.

- In a Transparent Data Encryption (TDE) environment, PowerExchange can capture change data from source tablespaces and columns that are encrypted. Before starting CDC, verify that the encryption wallet has been opened.

**Attention:** If you need to change the encryption key, wait until Oracle CDC with LogMiner finishes reading all of the change data from the Oracle archived redo logs for a CDC session and then change the encryption key. Otherwise, an Oracle error might occur and cause the CDC session to end abnormally.

- If you use Oracle materialized views, PowerExchange can capture change data from the master tables that underlie those views. PowerExchange supports change capture for any type of materialized view. The view and its underlying table have a one-to-one correspondence and share the same name.

If you issue `DTLDESCRIBE tables` from the **Database Row Test** dialog box in the PowerExchange Navigator, the results include a row for the materialized view and a row for the underlying table. The **Type** column indicates which row is for the materialized view and which row is for the table.

After the master tables receive data changes, Oracle refreshes the associated materialized views based on the REFRESH option that you specified in the CREATE MATERIALIZED VIEW statement. REFRESH FAST performs an incremental refresh that contains only the changes since the last refresh operation, and REFRESH COMPLETE refreshes all data. The REFRESH option that you use affects the level of PowerExchange change capture activity for the target table, as reported in the session log.

For example, the following table compares the effects of the REFRESH FAST and REFRESH COMPLETE options when the master table contains 400 rows:

REFRESH Option	Changes Made to the Master Table (rows)	Changes Written to the Target Table (rows)
REFRESH FAST	100 Inserts	100 Inserts
	100 Updates	100 Deletes followed by 100 Inserts
	100 Deletes	100 Deletes
REFRESH COMPLETE	100 Inserts	400 Deletes followed by 500 Inserts
	100 More Inserts	500 Deletes followed by 600 Inserts
	Then 100 Updates	600 Deletes followed by 600 Inserts
	Then 100 Deletes	600 Deletes followed by 500 Inserts

- PowerExchange Oracle CDC with LogMiner can capture change data from Oracle Data Guard logical standby database, which use SQL Apply to synchronize with the primary database. However, you must create a logon trigger for the Oracle CDC user (ORACAPT). The logon trigger disables Data Guard for the user logon. The PowerExchange DataGuard\_Logon\_Trigger.sql file contains sample SQL statements for creating this trigger. If you do not create a logon trigger, the following problems can occur because Data Guard locks the logical standby database to keep it consistent with the primary production database:

- PowerExchange Navigator row tests and the PowerExchange Logger might fail with the following Oracle error message:

```
ORA-01031: insufficient privileges on the SELECT from V$LOGMNR_CONTENTS
```

- You cannot create supplemental log groups for Oracle data sources.

- For tables that do not have row movement enabled, you can populate the PowerExchange-generated DTL\_\_CAPXROWID column in change records with Oracle physical rowid values. You must set the ROWID parameter to Y or ALLOW in the ORCL CAPI\_CONNECTION statement in the dbmover configuration file. This feature is useful for processing rows in unkeyed tables during CDC extraction sessions.

- PowerExchange uses Oracle LogMiner to read change data from archived logs. If you use an archived log destination other than the LOG\_ARCHIVE\_DEST\_1 path and LogMiner processing lags behind, problems might occur. In this situation, LogMiner starts reading change data from the archived logs in the LOG\_ARCHIVE\_DEST\_1 directory. If these archived logs are inaccessible from the machine with the Oracle instance to which you are connected, the LogMiner session might fail.

## Oracle CDC with LogMiner Restrictions

The following restrictions apply to PowerExchange Oracle CDC with LogMiner:

- PowerExchange Oracle CDC with LogMiner does not support change data capture from Oracle 12c multitenant environments.
- PowerExchange Oracle CDC with LogMiner does not capture change data from Oracle Data Guard physical databases, including active and passive standby databases that use Redo Apply to synchronize with the primary database.
- PowerExchange Oracle CDC with LogMiner does not support change data capture from the following Oracle objects:
  - Tables that use the system partitioning or reference partitioning scheme
  - Tables in a sorted hash clusters
  - Virtual columns, which contain derived data that Oracle does not log in the redo logs
  - Columns that have unsupported datatypes

Because you cannot include these columns in capture registrations, PowerExchange Express CDC for Oracle does not capture change data for them. However, PowerExchange Express CDC for Oracle can capture change data for other columns in the same registered table.
- If you truncate Oracle source tables from which change data is captured or if you drop and re-create source tables, PowerExchange cannot continue to extract change data for these tables. In these situations, you must rematerialize the corresponding targets.
- PowerExchange can capture data that the SQL\*Loader utility loaded into Oracle tables. However, the following restrictions apply:
  - The load type must be *conventional path*. PowerExchange cannot capture data that was loaded by a direct path load because Oracle LogMiner does not support direct path loads.
  - The load method must be Insert, Append, or Replace. Do not use Truncate. Truncate causes SQL\*Loader to issue TRUNCATE TABLE DDL. Because PowerExchange does not capture DDL, it cannot capture any row deletions that result from TRUNCATE TABLE DDL.



## Oracle Datatypes Supported for CDC

PowerExchange uses Oracle LogMiner to retrieve changes from the Oracle redo logs. Because Oracle does not log, or does not completely log, data with some datatypes in the redo logs, PowerExchange Oracle CDC with LogMiner cannot retrieve change data for all Oracle datatypes.

The following table identifies the Oracle datatypes that PowerExchange Oracle CDC with LogMiner supports and does not support:

Datatype	Supported for CDC?	Comments
BFILE	No	Data for columns that have this datatype are not completely logged in the Oracle redo logs and cannot be captured.
BINARY_DOUBLE	Yes	-
BINARY_FLOAT	Yes	-
CHAR	Yes	-
DATE	Yes	-
FLOAT	Yes	-
LOBs	No	-
LONG	No	Columns of this datatype cannot be included in capture registrations.
LONG RAW	No	Columns of this datatype cannot be included in capture registrations.
NCHAR	Yes	For CDC support of this datatype, you must have PowerExchange 8.5 or later.
NUMBER	Yes	PowerExchange handles NUMBER columns as follows: <ul style="list-style-type: none"><li>- Numbers with a scale of 0 and a precision value less than 10 are treated as INTEGER.</li><li>- Numbers with a defined precision and scale are treated as NUMCHAR.</li><li>- Numbers with an undefined precision and scale are treated as DOUBLE.</li></ul>
NVARCHAR2	Yes	For CDC support of this datatype, you must have PowerExchange 8.5 or later.
RAW	Yes	-
ROWID	Yes	-
TIMESTAMP	Yes	-
TIMESTAMP WITH TIME ZONE	No	-
TIMESTAMP WITH LOCAL TIME ZONE	No	-

Datatype	Supported for CDC?	Comments
UROWID	No	-
VARCHAR2	Yes	-
XML types	No	-

## Performance Considerations

Review the following considerations that pertain to PowerExchange Oracle CDC with LogMiner performance:

- Use real-time extraction mode only if you run very few concurrent change data extractions. PowerExchange CDC creates an Oracle LogMiner session for each real-time extraction. Because LogMiner sessions are resource intensive, they can impact overall system performance. Instead, use continuous extraction mode. For continuous extraction mode, PowerExchange extracts change data from PowerExchange Logger log files.
- If you use continuous extraction mode, minimize the size of the CDCT file. The CDCT file contains information about the PowerExchange Logger log files. PowerExchange reads the CDCT file each time the interval that is specified in the FILEWAIT parameter of the CAPX CAPL\_CONNECTION statement elapses. If a CDCT file is large, PowerExchange read operations can result in a high level of I/O activity, increased use of system resources, and increased extraction latency. To manage the CDCT file size, use the COND\_CDCT\_RET\_P statement in the pwxcl.cfg configuration file for the PowerExchange Logger for Linux, UNIX, and Windows.

## Oracle Configuration for CDC with LogMiner

PowerExchange provides sample script files to help you configure Oracle for PowerExchange Oracle CDC with LogMiner.

### RELATED TOPICS:

- [“Configuration Readme File and Script Files” on page 114](#)
- [“Configuring Oracle for CDC” on page 115](#)
- [“Configuration in an Oracle RAC Environment” on page 119](#)

## Configuration Readme File and Script Files

To configure PowerExchange Oracle CDC with LogMiner, use the readme\_oracapt.txt file and the sample SQL script files in the PowerExchange installation directory.

The readme\_oracapt.txt file describes the following sample SQL script files that PowerExchange provides for preparing Oracle for CDC:

- **Setup\_Archive\_Logging.sql** (non-RAC environments) or **Setup\_Archive\_Logging\_rac.sql** (RAC environments). Defines the archive log destination and enables ARCHIVELOG mode.
- **Setup\_Minimal\_Supplemental\_Logging.sql**. Enables global minimal supplemental logging.

- **Setup\_Logminer\_Tablespace.sql.** Creates a tablespace for exclusive LogMiner use.
- **Create\_ORACAPT\_user.sql.** Defines an Oracle CDC user and grants user privileges.
- **Optional\_GRANT\_for\_registrations.sql.** Grants the privilege that the Oracle CDC user requires to run generated DDL for creating a supplemental log group at capture registration completion in the PowerExchange Navigator.
- **Copy\_Oracle\_Dictionary.sql.** Copies the Oracle catalog to the archived redo logs.
- **DataGuard\_Logon\_Trigger.sql.** Creates a Data Guard logon trigger for the Oracle CDC user to prevent database row test or PowerExchange Logger failures when capturing changes from an Oracle Data Guard logical standby database.

## Configuring Oracle for CDC

This section describes steps for configuring Oracle for PowerExchange Oracle CDC with LogMiner. For sample SQL and DDL, refer to the sample script files that are specified in `readme_oracapt.txt`.

### Step 1. Set Up Archive Logging

For PowerExchange to capture changes from archive logs, you must define the archive log destination and enable ARCHIVELOG mode in Oracle. By default, ARCHIVELOG mode is not enabled.

For a non-RAC environment, use the sample SQL statements in the `Setup_Archive_Logging.sql` file. For a RAC environment, use the statements in the `Setup_Archive_Logging_rac.sql` file. Both files include detailed comments.

First, specify the archive log destination in one of the following ways:

- If you use the Oracle `init.ora` initialization parameter file, edit the appropriate parameters in this file to identify the archive log destination and file name format. For more information, see the Oracle database administrator's guide for your Oracle version.
- If you use a server parameter file (spfile), execute the `ALTER SYSTEM SET SQL` statements that are included in the sample `.sql` files. The specific SQL and configuration steps vary for RAC and non-RAC environments.

To enable ARCHIVELOG mode, issue the following statements:

```
SHUTDOWN IMMEDIATE;
STARTUP MOUNT;
ALTER DATABASE ARCHIVELOG;
ALTER DATABASE OPEN;
SHUTDOWN IMMEDIATE;
STARTUP;
```

**Tip:** Back up your database after both `SHUTDOWN` commands.

### Step 2. Set Up Oracle Minimal Global Supplemental Logging

PowerExchange requires that Oracle use minimal global supplemental logging so that Oracle LogMiner can properly handle chained rows.

To enable minimal global supplemental logging, log in to the Oracle database and execute the following SQL statement:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
COMMIT;
```

You can copy this statement from the sample `Setup_Minimal_Supplemental_Logging.sql` file.

**Note:** If you do not know whether minimal global supplemental logging is enabled for your database, you can still execute this ALTER statement. The statement has no effect if minimal supplemental logging is already enabled.

If the Oracle database is open when you execute this statement, Oracle waits for in-flight transactions to complete, which can affect database performance. This issue is likely to occur for databases that have a high level of user activity. To avoid this problem, you can close and re-open the database and then issue the statement manually.

You must also define a supplemental log group for each Oracle source table to cause Oracle to log full before- and after-images of the data that changed. When you register an Oracle table in the PowerExchange Navigator, PowerExchange generates DDL that adds a supplemental log group for the table.

### Step 3. Create a Tablespace for Oracle LogMiner Use

Create a tablespace exclusively for Oracle LogMiner use. Enable logging for this tablespace if you use logical standby databases, Oracle Streams, or Oracle native change capture processes.

Use the sample DDL in the Setup\_Logminer\_Tablespace.sql file.

**Note:** The creation of a separate LogMiner tablespace prevents the SYSAUX tablespace from becoming full, which could cause service problems during PowerExchange CDC.

1. To create the tablespace, issue the appropriate DDL statements for your environment:

- To create a tablespace in a file system:

```
CREATE TABLESPACE "LOGMNRTS"
  NOLOGGING
  DATAFILE
    '/oracle_path/datafilename.ora' SIZE
    50M REUSE AUTOEXTEND
  ON NEXT 10M MAXSIZE 100M EXTENT MANAGEMENT LOCAL;
```

- To create a tablespace in an Oracle ASM instance:

```
CREATE TABLESPACE "LOGMNRTS"
  NOLOGGING
  DATAFILE
    '+ORAA1RC1' SIZE 10M
  REUSE AUTOEXTEND
  ON NEXT 10M MAXSIZE 100M EXTENT MANAGEMENT LOCAL
  SEGMENT SPACE MANAGEMENT AUTO;
```

In the sample statements, specify NOLOGGING if you use Oracle LogMiner only for PowerExchange CDC and an occasional query. Specify LOGGING if you use any of the following Oracle features: logical standby databases, Oracle Streams, or native Oracle change capture processes.

Change the DATAFILE name to one that Oracle expects, based on your local Oracle database file naming standards.

2. Enter the following command:

```
EXECUTE SYS.DBMS_LOGMNR_D.SET_TABLESPACE('LOGMNRTS');
```

If this statement fails with the ORA\_01353 message, see the comments related to this message in the Setup\_Logminer\_Tablespace.sql file.

3. Recompile the SYS.DBMS\_LOGMNR\_D package. Use the following command:

```
ALTER PACKAGE SYS.DBMS_LOGMNR_D COMPILE BODY;
```

**Tip:** LogMiner opens a number of cursors for internal processing. When you configure LogMiner the first time, you might receive messages that state “number of open cursors exceeded.” In this case, increase the maximum number of open cursors to handle the extra LogMiner processing.

## Step 4. Define a CDC User and Grant User Privileges

Define an Oracle user that can run PowerExchange Oracle CDC with LogMiner. Then grant to this CDC user the specific Oracle system and object privileges that PowerExchange requires to extract change data from Oracle redo logs. You must also grant certain object privileges to the user who is responsible for creating and managing registrations and extraction maps for the Oracle source tables from the PowerExchange Navigator.

### Oracle CDC User

As the Oracle CDC user, use an existing user who has the required privileges, or create a user and grant the required privileges to that user.

To create an Oracle CDC user and grant the required privileges, use the sample SQL in the Create\_ORACAPT\_user.sql file. You can edit the sample SQL, as needed, for your environment.

The following table identifies the system privileges to grant to the Oracle CDC user:

System Privilege Granted	Description
CONNECT	Required. Grant this privilege to perform Oracle CDC data extraction in real time and to run PowerExchange Logger tasks. <b>Note:</b> The Oracle CDC user needs only the CREATE SESSION privilege on the CONNECT role.
LOCK ANY TABLE	If you specify GENRLOCK=Y in the ORCL CAPI_CONNECTION statement of the dbmover.cfg file, either grant the LOCK ANY TABLE system privilege or grant the SELECT object privilege on each table that is registered for change data capture.
SELECT ANY TRANSACTION	Required. Grant this privilege to perform Oracle CDC data extraction in real time and to run PowerExchange Logger tasks.

The following table identifies the minimum object privileges to grant to the Oracle CDC user:

Object Name	Object Privilege
Source tables	If you specify GENRLOCK=Y in the ORCL CAPI_CONNECTION statement of the dbmover.cfg file, you must either grant the LOCK ANY TABLE system privilege or grant the SELECT object privilege on each table that is registered for change data capture.
PUBLIC.V\$ARCHIVED_LOG	SELECT
PUBLIC.V\$DATABASE	SELECT
PUBLIC.V\$INSTANCE	SELECT
PUBLIC.V\$LOGMNR_CONTENTS	SELECT
PUBLIC.V\$NLS_PARAMETERS	SELECT
PUBLIC.V\$PARAMETER	SELECT
PUBLIC.V\$TRANSACTION	SELECT
SYS.DBA_LOG_GROUPS	SELECT
SYS.DBA_LOG_GROUP_COLUMNS	SELECT

Object Name	Object Privilege
SYS.DBMS_FLASHBACK	EXECUTE
SYS.DBMS_LOGMNR	EXECUTE
SYS.DBMS_LOGMNR_D	EXECUTE

### PowerExchange Navigator User

The user who uses the PowerExchange Navigator to create and manage capture registrations must have certain SELECT privileges to create and manage registrations. If you want the user to be able to run the generated DDL that creates supplemental log groups at the end of registration, also grant the privilege that is specified in the `Optional_GRANT_for_registrations.sql` file.

Grant the following privileges to the PowerExchange Navigator user who creates and manages registrations:

- To create capture registrations and perform other tasks in the PowerExchange Navigator, grant the following privileges:

```
GRANT SELECT ON "PUBLIC"."V$PARAMETER" TO "registration_user";
GRANT SELECT ON table TO "registration_user";    <<Repeat for each table of CDC
interest.
```

Instead of granting SELECT on each table of interest, you can specify GRANT SELECT ON ANY TABLE for the *registration user* if your site security rules allow it.

- To run the SQL for creating supplemental log groups at the end of registration, grant the following system privilege:

```
GRANT ALTER ANY TABLE TO "registration_user";
```

If your site security rules do not allow this level of authority to be granted to the registration user, you can give the SQL file to your DBA. The DBA can then use the SQL to create the supplemental log groups.

## Step 5. Copy the Oracle Catalog to the Archived Logs

PowerExchange Oracle CDC with LogMiner requires a copy of the Oracle online catalog in the Oracle archived redo logs to determine the point from which to restart change data extractions.

PowerExchange reads the last catalog copy in the archived logs, even if you specify `ONLINECAT=Y` in the `ORCL CAPI_CONNECTION` statement. You should copy the catalog on a routine basis to minimize CDC restart times.

To copy the catalog, issue the following command in an SQL\*Plus session:

```
begin
SYS.DBMS_LOGMNR_D.BUILD(
  options => sys.dbms_logmnr_d.store_in_redo_logs);
end;
/
```

**Tip:** Periodically, PowerExchange requests Oracle to recopy the catalog to the Oracle archived redo logs. To control how often Oracle copies the catalog and the time period within which the copy operation can occur, set the `CATBEGIN`, `CATEND`, and `CATINT` parameters in the `ORCL CAPI_CONNECTION` statement of the `dbmover` configuration file.

## Step 6. Create a Logon Trigger for an Oracle Data Guard Logical Standby Database (Optional)

If you use an Oracle Data Guard logical standby database as the source database for PowerExchange Oracle CDC with LogMiner, you must create a logon trigger for the Oracle CDC user. The logon trigger disables Data Guard for the user logon.

If you do not create a logon trigger, PowerExchange Navigator database row tests and the PowerExchange Logger fail with Oracle error message ORA-01031. Also, you cannot create supplemental log groups, as required for CDC, for your Oracle data sources.

To create a trigger on DATABASE, you must have the ADMINISTER DATABASE TRIGGER system privilege.

Use the following sample SQL from the PowerExchange DataGuard\_Logon\_Trigger.sql file:

```
CREATE OR REPLACE TRIGGER SYS.DISABLE_GUARD_4ORACAPT
AFTER LOGON
ON DATABASE
BEGIN
    IF USER IN('ORACAPT')
    THEN
        EXECUTE IMMEDIATE 'alter session disable guard ';
    END IF;
EXCEPTION
    WHEN OTHERS THEN
        NULL;
END;
/
```

In this SQL, the trigger name is SYS.DISABLE\_GUARD\_4ORACAPT, and the Oracle CDC user name is ORACAPT. You can change these names.

## Configuration in an Oracle RAC Environment

PowerExchange Oracle CDC with LogMiner can process change data for database instances in a real application cluster (RAC) environment.

For CDC to work properly, you might need to apply certain Oracle patches. In particular, Oracle bug 6596564 can cause data loss or data integrity problems in an Oracle RAC environment where CDC runs. If you use Oracle 11.1.0.7 on Windows, install Patch 24 before starting change capture. An Oracle 11.1.0.7 patch is not available for Linux and UNIX. If you use Oracle 11.2.0.1 or later, you do not need to apply a patch because the 6596564 bug fix is part of the base version.

Also complete the following configuration tasks:

- Ensure that the Oracle instance from which you run PowerExchange CDC can access the Oracle archived redo logs for all Oracle instances in the RAC from which change data is captured.
- In the init.ora file for each of these Oracle instances, define the LOG\_ARCHIVE\_DEST\_1 parameter to point to the directory in which Oracle creates the archived logs.

**Note:** PowerExchange uses Oracle LogMiner to read change data from the archived logs. If you use an archived log destination other than the LOG\_ARCHIVE\_DEST\_1 path and LogMiner processing lags behind, problems might occur. In this situation, LogMiner starts reading change data from the archived logs in the LOG\_ARCHIVE\_DEST\_1 directory. If these archived logs are inaccessible from the machine with the Oracle instance to which you are connected, the LogMiner session might fail.

- Ensure that all of the Oracle instances in the RAC that participate in CDC have access to the Oracle online redo logs. Usually, these redo logs reside on shared storage.

The tasks for configuring access to archived redo logs vary by operating system, as follows:

- On Windows, set up an Oracle flash recovery area on the shared file system that contains all of the table data for the RAC. For each Oracle instance in the RAC, set the LOG\_ARCHIVE\_DEST\_1 parameter to point to that recovery area.
- On Linux and UNIX, use any of the following methods:
  - Set up an Oracle flash recovery area in the same manner as for Windows.
  - Store all archived redo logs on shared storage.
  - Set up Network File System (NFS) access to the archive logs.

If you use shared storage or NFS access, the Oracle instance from which you run CDC must access the archived logs of the other RAC member instances. This access uses the mount points that match the archive log directories defined for those member instances. For example, ORA1 is the Oracle instance that runs CDC, and ORA2 is another Oracle instance in a RAC. ORA2 has a LOG\_ARCHIVE\_DEST\_1 parameter that points to the following archive log directory:

```
/ora/arch2/
```

The mount point that the ORA1 uses to access the ORA2 archive logs must also be `/ora/arch2/`.

## PowerExchange Configuration for CDC with LogMiner

The tasks that you perform to configure PowerExchange for Oracle CDC with LogMiner depend on the extraction mode and whether you use the PowerExchange Logger for Linux, UNIX, and Windows.

### RELATED TOPICS:

- [“Customizing the dbmover Configuration File for Oracle CDC” on page 122](#)
- [“Configuring Oracle CDC with LogMiner - With the PowerExchange Logger” on page 121](#)
- [“Configuring Oracle CDC with LogMiner - Without the PowerExchange Logger” on page 120](#)

## Configuring Oracle CDC with LogMiner - Without the PowerExchange Logger

If you plan to run extractions in real-time extraction mode and *not* use the PowerExchange Logger for Linux, UNIX, and Windows, complete the following tasks to configure PowerExchange Oracle CDC with LogMiner:

1. When you configure the dbmover configuration file on the Oracle source machine, include the following statements:
  - CAPT\_PATH
  - CAPT\_XTRA
  - ORACLEID
  - ORCL CAPI\_CONNECTION
  - UOWC CAPI\_CONNECTION

For more information, see the *PowerExchange Reference Manual*.

2. In the PowerExchange Navigator, create a capture registration for each Oracle source table.



If capture registrations already exist for these tables, delete the existing registrations and extraction maps and create new ones.

You must enter a name in the **Supplemental Log Group Name** field.

**Tip:** Set the **Condense** option to **Part** even though you do not plan to use the PowerExchange Logger, unless you have a specific reason not to do so. This practice prevents having to edit the capture registrations later if you decide to use the PowerExchange Logger. You might want to set the **Condense** option to **None** if you plan to run both real-time and continuous extractions against tables defined by the same capture registrations and do not want the PowerExchange Logger to capture change data for some registered tables.

The PowerExchange Navigator generates a corresponding extraction map and the DDL for creating a supplemental log group. If you selected the **Execute DDL now** option, PowerExchange executes the DDL for creating a supplemental log group when you click **Finish**. If you did not select this option, you must execute the DDL prior to starting extraction processing.

3. Activate the capture registrations. Usually, you do this task after materializing the targets.

**Next Step:** Configure and start extractions. You must use real-time extraction mode.

## Configuring Oracle CDC with LogMiner - With the PowerExchange Logger

If you plan to use the PowerExchange Logger for Linux, UNIX, and Windows and run extractions in batch or continuous extraction mode, complete the following tasks to configure PowerExchange Oracle CDC with LogMiner:

1. When you configure the dbmover configuration file used to access the source tables, include the following statements:

- CAPT\_PATH
- CAPT\_XTRA
- ORACLEID
- ORCL CAPI\_CONNECTION
- UOWC CAPI\_CONNECTION
- CAPX CAPI\_CONNECTION (for continuous extraction mode only)

For more information, see the *PowerExchange Reference Manual*.

2. Configure the pwxcl.cfg file for the PowerExchange Logger.
3. Start the PowerExchange Listener on the source machine.
4. Customize the dbmover configuration files on the Windows machine where the PowerExchange Navigator runs and on the PowerCenter Integration Service machine, if these machines are separate from the Oracle source machine.

In each of these dbmover configuration files, specify a **NODE** statement that points to the machine that contains the Oracle source tables. On the Windows machine, you must also specify an **ORACLEID** statement.

5. In the PowerExchange Navigator, create a capture registration for each Oracle source table.

If capture registrations already exist for these tables, delete the existing registrations and extraction maps and create new ones.

You must select **Part** in the **Condense** list, and enter a name in the **Supplemental Log Group Name** field. You can also set the **Status** option to **Active**, or wait until after you materialize the target tables.

The PowerExchange Navigator generates a corresponding extraction map and the DDL for creating a supplemental log group. If you selected the **Execute DDL now** option, PowerExchange executes the DDL for creating a supplemental log group when you click **Finish**. If you did not select this option, you must execute the DDL prior to starting extraction processing.

6. In the PowerExchange Navigator, perform a database row test on the extraction maps to verify that PowerExchange can access the source data.
7. After stopping updates to the source tables, materialize the target tables.
8. Start the PowerExchange Logger.
9. Allow changes to be written to the source tables.

**Next Step:** Configure and start extractions. You can use either batch extraction mode or continuous extraction mode.

## Customizing the dbmover Configuration File for Oracle CDC

Customize the dbmover configuration file to add some statements for PowerExchange Oracle CDC with LogMiner.

Enter the following statements for Oracle CDC with LogMiner:

### **CAPT\_PATH statement**

Path to the local directory on a Linux, UNIX, or Windows system that contains the control files for CDC.

These files are: the CCT file for capture registrations, the CDEP file for application names that are used for ODBC extractions, and the CDCT file for the PowerExchange Logger for Linux, UNIX, and Windows.

### **CAPT\_XTRA statement**

Path to the local directory on a Linux, UNIX, or Windows system that stores extraction maps for CDC.

### **ORACLEID statement**

The Oracle source instance, database, and connection information for CDC.

### **ORCL CAPI\_CONNECTION statement**

A named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and control CDC processing for Oracle sources.

### **UOWC CAPI\_CONNECTION statement**

A named set of parameters that the CAPI uses for the UOW Cleanser.

In this statement, the CAPINAME parameter points to an ORCL CAPI\_CONNECTION statement.

If you plan to use continuous extraction mode, you must also include the following statement:

### **CAPX CAPI\_CONNECTION statement**

A named set of parameters that the CAPI uses for continuous extraction of change data from PowerExchange Logger for Linux, UNIX, and Windows log files.

Define the CAPI\_CONNECTION statements in the dbmover configuration file on the system where the Oracle capture registrations are stored. This location corresponds to the **Location** node that you specify when defining a registration group. Usually, this location is where the source database resides.

Additionally, Informatica recommends including the LOGPATH and TRACING statements to make finding messages easier. The LOGPATH statement defines a directory specifically for PowerExchange message log files, and the TRACING statement enables PowerExchange to create an alternative set of message log files for each PowerExchange process.

For more information about all dbmover statements, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“CAPI\\_CONNECTION - CAPX Statement” on page 57](#)
- [“CAPI\\_CONNECTION - ORCL Statement” on page 125](#)
- [“CAPI\\_CONNECTION - UOWC Statement” on page 132](#)
- [“ORACLEID Statement” on page 123](#)
- [“CAPT\\_PATH Statement” on page 30](#)
- [“CAPT\\_XTRA Statement” on page 30](#)

## Example dbmover Statements for Oracle CDC with LogMiner

The following statements are typical of those included in a dmover configuration file for Oracle CDC with LogMiner:

```
LOGPATH=/pwx/logs
TRACING=/PFX=PWXLOG,RECLLEN=255,FILENUM=3,APPEND=Y,FLUSH=99)
CAPT_XTRA=/pwx/capture/vnnn/camaps
CAPT_PATH=/aus/pwx/capture/vnnn
ORACLEID=(FOX123,FO920DTL)
CAPI_SRC_DFLT=(ORA,CAPIUOWC)
CAPI_CONN_NAME=CAPIUOWC
/*
/* CAPI connection statements
/*
/* Both UOWC and ORCL CAPI_CONNNECTON statements are required for Oracle CDC.
CAPI_CONNECTION=(NAME=CAPIORA
,DLTRACE=ORA2
,TYPE=(ORCL
,ARRAYSIZE=1000
,BYPASSUF=Y
,CATBEGIN=00:01
,CATEND=23:59
,CATINT=1440
,ORACOLL=FOX123
,SELRETRY=0))
CAPI_CONNECTION=(NAME=CAPIUOWC
,TYPE=(UOWC
,CAPINAME=CAPIORA
,MEMCACHE=50000
,RSTRADV=1800))
/* Additional CAPX CAPI_CONNECTION statement is required for continuous extraction mode.
CAPI_CONNECTION=(NAME=CAPXORA
,TYPE=(CAPX
,DFLTINST=FOX920))
```

## ORACLEID Statement

The ORACLEID statement specifies the Oracle source instance, database, and connection information for CDC.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Oracle CDC sources

**Related Statements:** CAPI\_CONNECTION - ORCL and CAPI\_CONNECTION - ORAD

**Required:** Yes, for PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle

**Syntax:**

```
ORACLEID=(collection_id
,oracle_db
[,source_connect_string]
[,capture_connect_string])
```

```

        [,fifth_positional_parameter]
        [,USEDATABASE]
    )

```

#### **Parameters:**

##### ***collection\_id***

Required. User-defined identifier for this ORACLEID statement. This value must match the ORACOLL parameter value in the ORCL CAPI\_CONNECTION or ORAD CAPI\_CONNECTION statement, the collection ID in the registration group defined for the source tables, and the DBID value in the PowerExchange Logger pwxocl configuration file.

Maximum length is eight characters.

##### ***oracle\_db***

Required. Name of the Oracle database that contains the source tables that are registered for change data capture. If you use PowerExchange Express CDC for Oracle to capture change data from a pluggable database (PDB) in an Oracle multitenant environment, this value is the name of the database that contains the PDB.

##### ***source\_connect\_string***

Optional. Oracle connection string, defined in TNS, that is used to connect to the Oracle database that contains the source tables. This connection string must be defined in the Oracle Client tnsnames.ora file on the system with the source database.

For PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle, the source connection string is used only for PowerExchange Navigator access to the Oracle source database. Enter this parameter in the dbmover configuration file on the machine from which the PowerExchange Listener retrieves data for PowerExchange Navigator requests. If you plan to run a database row test on extraction maps for the source tables, also specify the *capture\_connect\_string* parameter.

**Note:** The source connection string is not used to transfer change data.

If this value is null and the Oracle source is *not* a PDB in a multitenant environment, the value of the ORACLE\_SID environment variable is used by default. If you use PowerExchange Express CDC for Oracle to capture change data from a PDBs, you must enter a value in this parameter.

##### ***capture\_connect\_string***

Optional. Oracle connection string, defined in TNS, that the PowerExchange Logger uses to connect to the Oracle database with the source tables for PowerExchange Oracle CDC with LogMiner or PowerExchange Express CDC with LogMiner. This connection string must be specified in the Oracle Client tnsnames.ora file that is used for connection to the Oracle source database. If you use PowerExchange Express CDC to capture change data from a PDB in an Oracle multitenant environment, specify the name of the PDB service entry in the tnsnames.ora file.

If this value is null and the Oracle source is *not* a PDB in a multitenant environment, the value of the ORACLE\_SID environment variable is used by default.

If this value is null and the Oracle source is a PDB, PowerExchange cannot capture change data for the source. If you use PowerExchange Express CDC for Oracle to capture change data from PDBs, you must enter a value in this parameter.

Also, for PowerExchange Oracle CDC with LogMiner or Express CDC for LogMiner, if you have multiple Oracle databases and capture changes from a database other than the default database, you must specify both the *source\_connect\_string* and *capture\_connect\_string* parameters.

**Tip:** If possible, bypass the use of SQL\*Net to improve PowerExchange Logger performance, even if the PowerExchange Logger is running on the same machine as the Oracle source database. Set the

following environment variables, whenever possible, to enable connection to the appropriate Oracle database without using the *capture\_connect\_string* parameter and SQL\*Net:

- ORACLE\_HOME
- ORACLE\_SID
- PATH
- On Linux or UNIX, one of the following:
  - LD\_LIBRARY\_PATH
  - LIBPATH
  - SHLIB\_PATH

#### ***fifth\_positional\_parameter***

Not used. Add a comma as a placeholder if you specify the USEDDBNAME positional parameter, for example:

```
ORACLEID=(collection_id,oracle_db,src_connect_string,capture_connect_string,,USEDDBNAME)
```

#### **USEDDBNAME**

Optional. Specify this parameter only under all of the following conditions:

- You upgrade to PowerExchange 9.1.0 or later from an earlier release.
- You use Oracle 11g or later.
- You run the following SQL query on the V\$DATABASE view and the query returns different values for the NAME and DB\_UNIQUE\_DATABASE fields, including values that vary in case only such as ORAABC1 and oraabc1:

```
select name, db_unique_name from v$database;
```

In this situation, the USEDDBNAME parameter can prevent potential restart errors that are caused by the difference in the NAME and DB\_UNIQUE\_DATABASE values.

**Tip:** Alternatively, you can specify the DB\_UNIQUE\_NAME value in the second positional parameter, *oracle\_db*.

#### **Usage Notes:**

- PowerExchange requires an ORACLEID statement for each Oracle database for which you want to capture and extract change data. You can define a maximum of 20 ORACLEID statements in a single dbmover configuration file.
- Define the ORACLEID statement in the dbmover configuration file on the system where the PowerExchange Logger runs, or if you plan to perform Oracle CDC without the PowerExchange Logger, on the system where your PowerExchange extractions run.

## **CAPI\_CONNECTION - ORCL Statement**

The ORCL CAPI\_CONNECTION statement specifies a named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and control PowerExchange Oracle CDC with LogMiner processing for Oracle sources.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Oracle

**Related Statements:** CAPI\_CONNECTION - UOWC and ORACLEID

**Required:** Yes for PowerExchange Oracle CDC with LogMiner

**Syntax:**

```

CAPI_CONNECTION=( [DLLTRACE=trace_id]
                  ,NAME=capi_connection_name
                  [,TRACE=trace_name]
                  ,TYPE=(ORCL
                        [,ARRAYSIZE=array_size|100]
                        [,BYPASSUF={N|Y}]
                        [,CATBEGIN=hh:mm|00:00]
                        [,CATEND=hh:mm|24:00]
                        [,CATINT=minutes|1400]
                        [,COMMITINT=minutes|5]
                        [,GENRLOCK={N|Y}]
                        [,IGNUFMSG={N|Y}]
                        [,LOGDEST=logdest_id]
                        [,LGTHREAD=instance_number]
                        [,ONLINECAT={N|Y}]
                        ,ORACOLL=collection_id
                        [,ROWID={N|Y|ALLOW}]
                        [,SELRETRY=retry_number|1000]
                        [,SGLINST={N|Y}]
                  )
                )

```

**Parameters:****DLLTRACE=trace\_id**

Optional. User-defined name of the TRACE statement that activates internal DLL tracing for this CAPI.

Specify this parameter only at the direction of Informatica Global Customer Support.

**NAME=capi\_connection\_name**

Required. Unique user-defined name for this CAPI\_CONNECTION statement.

Maximum length is eight alphanumeric characters.

**TRACE=trace\_name**

Optional. User-defined name of the TRACE statement that activates the common CAPI tracing.

Specify this parameter only at the direction of Informatica Global Customer Support.

**TYPE=(ORCL, ... )**

Required. Type of CAPI\_CONNECTION statement. For PowerExchange Oracle CDC with LogMiner sources, this value must be ORCL.

**ARRAYSIZE={array\_size|100}**

Optional. The number of rows of the prefetch array that PowerExchange uses to read the Oracle redo logs.

Enter a number from 0 through 2147483647. Default is 100.

**Note:** A value less than 100 can degrade Oracle CDC with LogMiner performance. A value of 0 disables prefetch. Enter 0 only at the direction of Informatica Global Customer Support.

**BYPASSUF={N|Y}**

Optional. Controls whether PowerExchange ends abnormally or issues a warning message when Oracle LogMiner returns an unformatted log record.

LogMiner returns unformatted log records when Global Temporary Tables are updated, or when ONLINECAT=Y and the log data that is being read is inconsistent with the catalog.

Enter one of the following options:

- **N.** PowerExchange ends with an error whenever it receives an unformatted log record from Oracle LogMiner.
- **Y.** PowerExchange writes a warning message to the PowerExchange message log that indicates unformatted log data was found and then continues processing. Depending on the amount of unformatted log data, PowerExchange might write many warning messages. To suppress these warning messages, specify Y for the IGNUFMSG parameter.

Default is N.

**Tip:** Enter Y if the Oracle instance contains Global Temporary tables. Otherwise, do not include the BYPASSUF parameter.

**CATBEGIN={hh:mm|00:00}**

Optional. Earliest time of day, in a 24-hour clock format, at which PowerExchange requests Oracle to write a copy of the Oracle catalog to the redo logs.

If you specify a value for the CATBEGIN parameter, you must also specify a value for the CATEND parameter.

Default is 00:00.

**CATEND={hh:mm|24:00}**

Optional. Latest time of day, in a 24-hour clock format, at which PowerExchange requests Oracle to write a copy of the Oracle catalog to the redo logs.

If you specify a value for the CATEND parameter, you must also specify a value for the CATBEGIN parameter.

Default is 24:00.

**CATINT={minutes|1440}**

Optional. Time interval, in minutes, between requests to copy the Oracle catalog to the redo logs.

Enter a number from 1 through 1440. Default is 1440.

If this interval elapses but the time is outside the time period that is specified by the CATBEGIN and CATEND parameters, PowerExchange does not request Oracle to make a copy of the Oracle catalog. Instead, PowerExchange waits the amount of time that you specify in the CATBEGIN parameter to request a catalog copy.

**COMMITINT={minutes|5}**

Optional. Time interval, in minutes, between the SQL COMMIT operations that PowerExchange issues to commit the transactions that the Oracle LogMiner session generates.

Enter a number from 1 through 60. Default is 5.

Although PowerExchange does not update data in user tables while reading change data from the redo logs, the Oracle LogMiner interface automatically generates transactions for the LogMiner sessions that PowerExchange initiates. Oracle leaves these transactions open, or in-flight, until the LogMiner session ends.

To be able to restart change data extraction operations efficiently, PowerExchange must occasionally issue SQL COMMIT operations to end these in-flight transactions. Otherwise, the restart of all future real-time extraction operations might be impacted because PowerExchange always begins reading change data at the beginning of the oldest in-flight UOW.

**GENRLOCK={N|Y}**

Optional. Controls whether PowerExchange generates a safe restart point for requests for restart points that match the current end-of-log (EOL).

Enter one of the following options:

- **N.** PowerExchange generates restart points that match the current EOL, ignoring any in-flight transactions for the source tables.
- **Y.** PowerExchange generates safe restart points for source tables.

Default is N.

A safe restart point for a source table is a point in the change stream that does not skip any in-flight UOWs for that table. To generate a safe restart point for a source table, PowerExchange gets an exclusive lock on the table to stop further changes. PowerExchange then searches the Oracle catalog for the point in the change stream that matches the earliest active transaction for the table and uses this point as the restart point. If no in-flight UOWs exist for a table, PowerExchange uses the current EOL. PowerExchange releases the lock on the source table after restart point generation completes. Then changes can be written to the table again.

PowerExchange generates restart tokens that match the current EOL in the following situations:

- You could start the PowerExchange Logger for Linux, UNIX, and Windows and the `pwxccl` configuration file does not specify the `SEQUENCE_TOKEN` and `RESTART_TOKEN` parameters.  
PowerExchange gets locks for all of the tables that are associated with active capture registrations to be used for PowerExchange Logger processing.
- The restart token file for a CDC session specifies the `CURRENT_RESTART` option on the `RESTART1` and `RESTART2` special override statements.  
PowerExchange gets locks only for the tables in the CDC session to which the special override statements apply.
- A database row test in the PowerExchange Navigator that uses the `SELECT CURRENT_RESTART` SQL statement.  
PowerExchange gets a lock for the table represented by capture registration associated with the extraction map used in the database row test.
- A `DTLUPPL` utility operation that uses the `RSTTKN GENERATE` option.  
PowerExchange gets a lock for the table represented by the capture registration specified in the utility control statements.

**IGNUFMSG={N|Y}**

Optional. Controls whether PowerExchange writes warning messages to the PowerExchange message log file for unformatted data records.

Enter one of the following options:

- **N.** PowerExchange does not write any warning messages.
- **Y.** PowerExchange writes warning messages.

Default is N.

**LOGDEST=logdest\_id**

Optional. The numeric identifier for the archive log destination that you want to force PowerExchange to use. This archive log destination must be local to the Oracle instance that PowerExchange is using.



Enter a number from 1 through 10.

For example, to use archived logs from the destination set by the LOG\_ARCHIVE\_DEST\_3 parameter in the init.ora file, specify LOGDEST=3.

The SNGINST parameter affects how PowerExchange uses the archive log destination and the Oracle instance that you specify in LOGDEST and LGTHREAD.

If you specify Y for the ONLINECAT parameter, PowerExchange validates and then ignores the LOGDEST and LGTHREAD parameters.

**LGTHREAD=instance\_number**

Optional. The instance number for the Oracle instance. PowerExchange uses this instance number to identify the archived redo logs to process.

Enter a number from 1 through 2147483647.

For non-RAC environments, if you specify this parameter, set it to 1.

The SNGINST parameter affects how PowerExchange uses the archive log destination and the Oracle instance that you specify in LOGDEST and LGTHREAD.

If you specify Y for the ONLINECAT parameter, PowerExchange validates and then ignores the LOGDEST and LGTHREAD parameters.

**ONLINECAT={N|Y}**

Optional. Controls whether PowerExchange directs Oracle LogMiner to use the Oracle online catalog or the copy of the catalog in the redo logs to format log data for CDC.

Enter one of the following options:

- **N.** Oracle LogMiner uses the copy of the catalog from the archived redo logs and PowerExchange tracks schema changes to ensure that data loss does not occur.
- **Y.** Oracle LogMiner uses the online catalog and PowerExchange cannot track schema changes.

Default is N.

When you configure PowerExchange to use the online catalog for formatting log data, PowerExchange still uses catalog copies to determine the restart point for change data extraction operations. Therefore, you must copy the online catalog to the Oracle redo logs on a regular basis.

Change data extraction operations generally initialize faster when you configure PowerExchange to create LogMiner sessions with the online catalog instead of a catalog copy. However, when LogMiner uses the online catalog, it does not track DDL changes, and cannot format log records for tables that have schema changes.

If LogMiner uses the online catalog and you make schema changes while LogMiner is reading log data, LogMiner passes unformatted log records for subsequent changes to PowerExchange. If you specify N for the BYPASSUF parameter or accept the default value of N, PowerExchange fails the extraction request after Oracle passes the first unformatted record. Otherwise, PowerExchange skips the unformatted record and continues processing, which results in change data loss. Therefore, specify N for the ONLINECAT parameter, or allow it to default, if you have the following requirements:

- You specify Y for the BYPASSUF parameter and need to change the schema of tables registered for capture while change data extraction operations are running.
- You need to start an extraction from a point in the Oracle redo logs that contains table data that PowerExchange captured under a previous schema.

**ORACOLL=*collection\_id***

Required. Oracle collection identifier, which must match the value specified in the ORACLEID statement.

**ROWID={*N*|*Y*|*ALLOW*}**

Controls whether Oracle physical rowid values are included in captured change records for tables that do not have Oracle row movement enabled. PowerExchange writes the rowid values to the PowerExchange-generated DTL\_\_CAPXROWID column. For example, use this parameter if you have unkeyed source tables on which you need to perform some processing that requires a unique row ID when extraction sessions run.

Enter one of the following options:

- **N.** Do not capture rowid values. The DTL\_\_CAPXROWID column contains null values.
- **Y.** Capture rowid values for tables that do not have row movement enabled and write the values to the DTL\_\_CAPXROWID column in change records. If a table has row movement enabled, capture processing ends abnormally.
- **ALLOW.** Capture rowid values for tables that do not have row movement enabled and write the values to the DTL\_\_CAPXROWID column in change records. If a table has row movement enabled, return null values to the DTL\_\_CAPXROWID column and continue capture processing. You might want capture processing to continue if you do not need rowid values for the tables that have row movement enabled.

**Note:** This parameter pertains to PowerExchange Oracle CDC with LogMiner. If you use PowerExchange Express CDC for Oracle, set the similar ROWID parameter in the OPTIONS statement of the Express CDC configuration file instead.

Default is N.

**SELRETRY={*retry\_number*|1000}**

Optional. The number of times that PowerExchange immediately loops back to the Oracle LogMiner call before implementing a graduated-scale wait loop.

After PowerExchange retries the call to LogMiner the specified number of times, PowerExchange implements a wait interval between each subsequent retry. The wait interval begins at one millisecond and gradually increases to one second. When LogMiner returns data, PowerExchange resets the wait interval to 0, and the process begins again for the next call to LogMiner.

For the *retry\_number* variable, enter a number from 0 through 2147483647. Default is 1000.

If you specify a nonzero value, PowerExchange uses nonblocking SQL to ensure that it can process a user request to shut down an extraction session in a timely manner.

If you specify 0, PowerExchange does not use nonblocking SQL. This setting improves CPU consumption but can prolong shutdown of an extraction session. On quiescent Oracle instances, PowerExchange does not honor a shutdown request until Oracle returns log data. On Oracle instances where update activity is occurring, shutdown behavior does not noticeably change.

**Important:** If you capture change data from an Oracle 12.1 or later source, set the SELRETRY parameter to 0. Otherwise, the Oracle LogMiner sessions for PowerExchange CDC fail when trying to fetch change data.

**SINGLINST={*N*|*Y*}**

Optional. Controls whether PowerExchange uses only the archived redo logs from a specific Oracle instance and archive log destination.

Enter one of the following options:

- **N.** PowerExchange uses the specified Oracle instance to search for the archived redo logs that contain copies of the Oracle catalog. After PowerExchange passes these logs to an Oracle LogMiner session, LogMiner determines the other archived redo logs to read.
- **Y.** PowerExchange uses only the archive log destination and Oracle instance that you specify in LOGDEST and LGTHREAD parameters to read archived redo logs. LogMiner does not read any other archived redo logs. After PowerExchange processes the logs from the specified location, the change data extraction operation ends.

If you specify Y, you must also specify the LOGDEST and LGTHREAD parameters to identify the archive log destination and Oracle instance to use. In a RAC environment, you must run separate change data extraction processes for all remaining Oracle instances in the RAC and determine how to properly merge the change data so that it can be applied to the targets.

Default is N.

## Oracle Catalog Parameters in the ORCL CAPI\_CONNECTION Statement

The CATINT, CATBEGIN, and CATEND parameters in the ORCL CAPI\_CONNECTION statement can affect PowerExchange performance.

These parameters control the frequency at which the Oracle catalog is copied to the Oracle redo logs and the time period within which these copy operations can occur. Because catalog copy operations are resource intensive, they can degrade the performance of both PowerExchange and the Oracle instance.

When you restart PowerExchange extraction processing, PowerExchange directs Oracle LogMiner to begin reading change data from the redo logs. The read starts from the SCN of the last Oracle catalog copy written to the logs before the previous extraction session ended. Restart processing might become inefficient if you copy the catalog too infrequently.

To configure the CATINT, CATBEGIN, and CATEND parameters, try various settings until you find a combination that provides for efficient restart processing without degrading Oracle and PowerExchange performance to an unacceptable level. The default frequency of once a day might not be sufficient if you have a high volume of transaction activity.

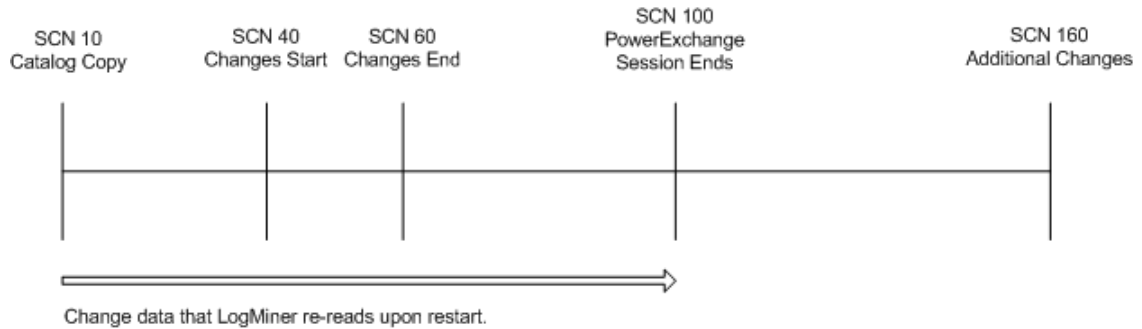
The following examples demonstrate how copying the Oracle catalog multiple times can affect the amount of change data that is reread from the archived redo logs when PowerExchange extraction processing is restarted.

### Example 1

Assume that the Oracle catalog was initially copied to the Oracle redo logs at SCN 10 and another copy has not yet been written to the logs. Change data was logged starting at SCN 40 and ending at SCN 60. A PowerExchange extraction session extracted these changes before ending at SCN 100. Since the extraction session ended, additional changes have been logged starting at SCN 160.

When you restart PowerExchange extraction processing, LogMiner must begin reading change data from the initial catalog copy at SCN 10 because it is the latest catalog copy prior to the session end at SCN 100. As a result, PowerExchange reprocesses the data between SCN 10 and SCN 100, before continuing to the new change data that begins at SCN 160. This reprocessing can affect PowerExchange performance.

The following figure shows, in a linear manner, that the reprocessing starts at SCN 10 and ends at SCN 100:

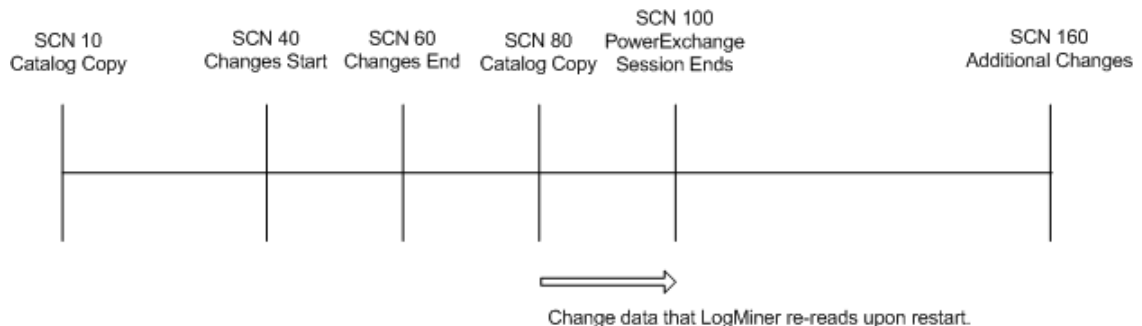


## Example 2

Assume that the Oracle catalog was copied to the Oracle redo logs twice: at SCN 10 and at SCN 80. Change data was logged starting at SCN 40 and ending at SCN 60. A PowerExchange extraction session extracted these changes before ending at SCN 100. Since the extraction session ended, additional changes have been logged starting at SCN 160.

When you restart PowerExchange extraction processing, LogMiner begins reading change data from the second catalog copy at SCN 80 because it is the latest catalog copy prior to the session end at SCN 100. As a result, PowerExchange reprocesses only the data between SCN 80 and SCN 100, before continuing to the new change data that begins at SCN 160.

The following figure shows, in a linear manner, that reprocessing starts from the last catalog copy at SCN 80 and stops at SCN 100:



In this case, less data is reprocessed.

## CAPI\_CONNECTION - UOWC Statement

The UOWC CAPI\_CONNECTION statement specifies a named set of parameters that the Consumer API (CAPI) uses for the UOW Cleanser.

In the change stream for some data sources, changes from multiple UOWs are intermingled. The UOW Cleanser reconstructs the intermingled changes read from the change stream into complete UOWs in chronological order based on end time.

**Operating Systems:** i5/OS, Linux, UNIX, Windows, and z/OS

**Data Sources:** DB2 for i5/OS CDC sources, Oracle CDC with LogMiner sources, and z/OS CDC sources

**Related Statements:** CAPI\_CONNECTION - AS4J, CAPI\_CONNECTION - LRAP, and CAPI\_CONNECTION - ORCL

**Required:** Yes, for CDC for the specified sources

**Syntax:**

```

CAPI_CONNECTION=( [DLLTRACE=trace_id]
                  ,NAME=capi_connection_name
                  [,TRACE=trace_name]
                  ,TYPE=(UOWC
                        ,CAPINAME=source_capi_name
                        [,BLKSIZE=block_size]
                        [,DATACLASS=data_class]
                        [,LARGEOPS=number_of_operations]
                        [,MEMCACHE={cache_size|1024}]
                        [,MONITORINT={minutes|5}]
                        [,RSTRADV=seconds]
                        [,SPACEPRI={primary_space|50}]
                        [,SPACETYP={BLK|TRK|CYL}]
                        [,SPILLKEEP=number_of_spill_files]
                        [,STORCLASS=storage_class]
                        [,TIMESTAMP={LOG|COMMIT}]
                        [,UNIT=unit]
                  )
                )

```

**Parameters:****DLLTRACE=trace\_id**

Optional. A user-defined name for the TRACE statement that activates internal DLL tracing for this CAPI.

Specify this parameter only at the direction of Informatica Global Customer Support.

**NAME=capi\_connection\_name**

Required. A unique user-defined name for this CAPI\_CONNECTION statement.

Maximum length is eight alphanumeric characters.

**TRACE=trace\_name**

Optional. A user-defined name for the TRACE statement that activates the common CAPI tracing.

Specify this parameter only at the direction of Informatica Global Customer Support.

**TYPE=(UOWC, ... )**

Required. The type of CAPI\_CONNECTION statement. For the UOW Cleanser, this value must be UOWC.

**CAPINAME=capi\_name**

Required. The value of the NAME parameter in the related source-specific CAPI\_CONNECTION statement, which can be one of the following statement types:

- AS4J CAPI\_CONNECTION statement for DB2 for i5/OS sources
- ORCL CAPI\_CONNECTION statement for Oracle CDC with LogMiner sources
- LRAP CAPI\_CONNECTION statement for data sources on z/OS

**BLKSIZE=block\_size**

Optional. The block size, in bytes, for the sequential UOW spill files that the UOW Cleanser creates when the memory cache cannot hold all changes for a UOW.

The following table shows valid values by CDC source type:

Data Source Type	Valid Values	Default Value
DB2 for i5/OS	A number from 8 through 32760	32760
Oracle	A number from 8 through 65535	32768
z/OS data sources	A number from 8 through 32760	18452

**DATACLASS=*data\_class***

Optional. On z/OS, the SMS data class that the UOW Cleanser uses when allocating the sequential UOW spill files. If you do not specify this parameter, the SMS ACS routines can assign the data class.

**LARGEOPS=*number of operations***

Optional. Overrides the default value that PowerExchange uses to identify transactions as large transactions for reporting purposes. Enter the number of DML operations (inserts, updates, and deletes), in thousands, that a transaction must contain to be considered a large transaction.

PowerExchange issues status messages for large transactions that meet this criteria. If PowerExchange issues too many messages, you can increase this value to limit the number of messages.

Valid values are 1 through 2147483 (1000 through 2,147,483,000 operations). The default value is one half of the MEMCACHE parameter value rounded up to the nearest thousand. Based on the default MEMCACHE value of 1024 KB, the default LARGEOPS value is 1000 (1,000,000 operations).

**MEMCACHE={*cache\_size*|1024}**

Optional. The maximum memory cache size, in kilobytes, that PowerExchange allocates to reconstruct complete UOWs.

Enter a number from 0 through 2147483647. Default is 1024 KB. If you enter 0, the memory cache size is limited only by the available memory on the system.

For each extraction session, PowerExchange keeps all changes for each UOW in the memory cache until it processes the end-UOW record. PowerExchange incrementally allocates memory cache up to the limit that this parameter specifies. If the memory cache is too small to hold all of the changes in a UOW, PowerExchange spills the changes to a sequential files on disk, called UOW spill files.

Each UOW spill file contains one UOW. A UOW might require multiple UOW spill files to hold all of the changes for that UOW. If the change stream contains multiple large UOWs and the memory cache is insufficient, PowerExchange might create numerous UOW spill files.

PowerExchange processes the change stream more efficiently if it does not need to use UOW spill files. A large number of UOW spill files can degrade extraction performance and cause disk space shortages.

**Important:** If the change stream contains small UOWs, the default value might be sufficient. However, Informatica recommends that you specify a larger value because the default value is often too small.

The location in which PowerExchange allocates the UOW spill files varies by operating system, as follows:

- For i5/OS, PowerExchange uses CRTPF command to create a physical file for UOW spill files.  
PowerExchange names the UOW spill files using the C/C++ `tmpnam()` function.
- For Linux and UNIX, PowerExchange uses the current directory by default for UOW spill files. To use a different directory, specify the `TMPDIR` environment variable.  
PowerExchange names the UOW spill file names using the prefix "dtlq" and the operating system function `tempnam`.  
**Note:** The UOW spill files are temporary files that are deleted when PowerExchange closes them. These files are not visible in the directory while they are open.
- For Windows, PowerExchange uses the current directory by default for UOW spill files. To use a different directory, specify the `TMP` environment variable.  
PowerExchange names the UOW spill file using the prefix "dtlq" and the Windows `_tempnam` function.
- For z/OS, PowerExchange uses dynamic allocation to allocate temporary data sets for the UOW spill files. Generally, SMS controls the location of temporary data sets. If you do not use SMS to control temporary data sets, the UNIT parameter controls the location for the UOW spill files.  
Because PowerExchange allocates temporary data sets for the UOW spill files, z/OS assigns these files system-generated data set names, which begin with `SYSydd.Thhmmss.RA000.jobname`.

**Warning:** PowerExchange allocates the cache size for each extraction operation. If you use a large MEMCACHE value and run many concurrent extraction sessions, memory constraints can occur.

#### **MONITORINT=minutes**

Optional. The time interval, in minutes, at which PowerExchange checks transaction activity for long outstanding transactions and large transactions. A long outstanding transaction is one that remains active for two monitoring intervals, and a large transaction is one that meets the LARGEOPS criteria. When this interval elapses, PowerExchange issues messages that identify the large transactions and long outstanding transactions and report their processing activity. PowerExchange also issues messages that identify the current position in the change stream. Valid values are 0 through 720. A value of 0 disables monitoring. Default is 5.

#### **RSTRADV=seconds**

The time interval, in seconds, that PowerExchange waits before advancing restart and sequence tokens for a registered data source during periods when UOWs do not include any changes of interest for the data source. When the wait interval expires, PowerExchange returns the next committed "empty UOW," which includes only updated restart information.

Enter a number from 0 through 86400. No default is provided.

PowerExchange resets the wait interval to 0 when one of the following events occur:

- PowerExchange completes processing a UOW that includes changes of interest.
- PowerExchange returns an empty UOW because the wait interval expired without PowerExchange receiving any changes of interest.

For example, if you specify 5, PowerExchange waits five seconds after it completes processing the last UOW or after the previous wait interval expires. Then PowerExchange returns the next committed empty UOW that includes the updated restart information and resets the wait interval to 0.

If you do not specify RSTRADV, PowerExchange does not advance restart and sequence tokens for a registered source during periods when PowerExchange receives no changes of interest. When PowerExchange warm starts, it reads all changes, including those not of interest for CDC, from the restart point.

For DB2 for i5/OS sources, Informatica recommends that you use this parameter if the change records that PowerExchange reads from i5/OS journal receivers are created under commitment control. If the change records are created without commitment control, do not specify this parameter.

**Attention:** A value of 0 can degrade performance. In addition to the UOWs that contain changes for registered sources of interest, PowerExchange returns an empty UOW for every UOW that does not contain changes for the registered sources of interest.

**SPACEPRI={primary\_space|50}**

Optional. On z/OS, the amount of primary space that the UOW Cleanser uses for allocating UOW spill files. The SPACETYP parameter indicates the type of space units.

Enter a number from 1 through 16777215. Default is 50 blocks.

The UOW Cleanser does not use secondary space. Instead, when a spill file becomes full, the UOW Cleanser allocates another spill file of the same size.

SMS ACS routines can override the UOW spill file size.

**Note:** On i5/OS, the UOW Cleanser allocates UOW spill files as physical files with SIZE(\*NOMAX), which means that the maximum spill file size is controlled by the system maximum file size. On Linux, UNIX, and Windows, PowerExchange allocates UOW spill files as temporary files that are 2 GB in size.

**SPACETYP={BLK|TRK|CYL}**

Optional. On z/OS, the type of units in which the primary space for UOW Cleanser allocation of UOW spill files is expressed.

Options are:

- **BLK.** Blocks.
- **CYL.** Cylinders.
- **TRK.** Tracks.

Default is BLK.

**SPILLKEEP=number\_of\_spill\_files**

Optional. The number of spill files that the UOW Cleanser retains for re-assignment. The UOW Cleanser retains spill files instead of deallocating them so that the files are available to be reassigned to new transactions. This feature is intended to prevent excessive file deallocation and allocation activity.

Valid values are 0 through 999. On z/OS and i5/OS, the default is 3. On Linux, UNIX, and Windows, the default is 0.

**STORCLASS=storage\_class**

Optional. On z/OS, the SMS storage class name that the UOW Cleanser uses to allocate UOW spill files.



**TIMESTAMP={LOG|COMMIT}**

Optional. The type of timestamp that PowerExchange records in the generated DTL\_\_CAPXTIMESTAMP column of each change record for a transaction. Specify this parameter only if you want to display the commit timestamp instead of the timestamp from the source logs or data sets.

Options are:

- **LOG.** The timestamp of a change on the source database, as recorded by the DBMS in the source database logs or data sets near the time when the change is made. For more information, see [Appendix A, “DTL\\_\\_CAPXTIMESTAMP Time Stamps” on page 279](#).
- **COMMIT.** The timestamp of the transaction commit on the source database. Specify this option if you use the timestamp to calculate latency.

Default is LOG.

**UNIT=unit**

Optional. On z/OS, the generic or esoteric unit name that the UOW Cleanser uses to allocate UOW spill files.

## Managing Oracle CDC with LogMiner

You might need to perform some CDC management tasks.

### Stopping Oracle CDC with LogMiner Processing for Selected Tables

You can stop change data capture for one or more tables temporarily or permanently.

1. In the PowerExchange Navigator, open the capture registration for each table that is no longer of CDC interest, and change the **Status** option from **Active** to one of the following settings:

- 
- **History.** This option permanently deactivates the registration and stops all change data capture based on the registration. You cannot use the registraton for change data capture again.

2. If you use the PowerExchange Logger for Linux, UNIX, and Windows for Oracle CDC, shut down the PowerExchange Logger and then warm start it.

This step refreshes the registration information that the PowerExchange Logger uses.

3. Drop the supplemental log groups for the tables with the deactivated registrations. Use the following DDL:

```
ALTER TABLE schema.table_name DROP SUPPLEMENTAL LOG GROUP
```

Oracle stops recording before- and after-images of the change data for the tables. If you reinstate a supplemental log group later, you must rematerialize the target database.

In PowerCenter, you must also delete or update CDC workflows, as needed, to make sure that no workflow processes the removed tables.

## RELATED TOPICS:

- [“Stopping PowerCenter CDC Sessions” on page 248](#)

## Stopping Oracle CDC with LogMiner Processing Temporarily

If you use the PowerExchange Logger for Linux, UNIX, and Windows and want to temporarily stop PowerExchange Oracle CDC with LogMiner processing for all source tables, shut down the PowerExchange Logger.

For example, you might want to stop CDC processing for troubleshooting or target database maintenance. You can warm start the PowerExchange Logger later without any change data loss.

## Changing a Source Table Definition Used in Oracle LogMiner CDC

Occasionally, you might need to change the definition of an Oracle source table that is registered for change data capture.

If your table definition changes affect the columns from which change data is captured, complete this procedure to enable PowerExchange to switch to the updated table definition, while preserving access to previously captured data. These table definition changes include adding, altering, or dropping columns. You do not need to perform this procedure if you are selectively capturing change data for a subset of columns and none of the selected columns are affected by the table definition changes.

**Tip:** If you no longer need to capture change data from a column in a table, you can remove that column from the extraction map without changing the capture registration. Change data for the column is still captured but is not extracted.

1. Stop DELETE, INSERT, and UPDATE activity against the table.
2. Verify that any change data that was captured under the previous table definition has completed extraction processing. Then stop all workflows that extract change data for the table.
3. In the PowerExchange Navigator, open the original capture registration and set its status to **History**. PowerExchange does not capture change data based on capture registrations that have a status of **History** or **Inactive**.
4. Drop the supplemental log group for the table.
5. Use DDL to make the table changes.
6. In the PowerExchange Navigator, create a new capture registration that reflects the metadata changes. When creating the registration, specify these settings:
  - Set the capture registration **Status** to **Active**.
  - Select the **Execute DDL now** option so that when you finish the capture registration, the PowerExchange Navigator runs DDL for creating a new supplemental log group.PowerExchange uses the newly activated capture registration for change data capture.
7. Change the target table definition to reflect the source table metadata changes, if necessary.
8. If you use the PowerExchange Logger for Linux, UNIX, and Windows, restart the PowerExchange Logger process so that it will begin using the new capture registration.
9. In PowerCenter Designer, import the altered source and target tables. Edit the mapping if necessary.
10. If necessary, rematerialize the target tables. After materialization completes, create new restart tokens.
11. Re-enable DELETE, INSERT, and UPDATE activity against the table.
12. Restart extraction processing.

## CHAPTER 7

# Express CDC for Oracle

This chapter includes the following topics:

- [Express CDC for Oracle Overview, 139](#)
- [PowerExchange Express CDC for Oracle Benefits, 140](#)
- [PowerExchange Express CDC for Oracle Architecture, 141](#)
- [PowerExchange Express CDC Configuration Overview, 145](#)
- [Gathering Information About the CDC Environment, 145](#)
- [Planning Considerations, 147](#)
- [Summary Task Flow, 158](#)
- [Configuring Oracle for Express CDC, 159](#)
- [Configuring PowerExchange for Express CDC, 162](#)
- [Managing PowerExchange Express CDC for Oracle, 184](#)

## Express CDC for Oracle Overview

PowerExchange Express CDC for Oracle captures change data directly from Oracle active and archived redo logs and makes that data available to PowerCenter CDC sessions for propagation to targets. It is an alternative to Oracle CDC with LogMiner.

PowerExchange Express CDC for Oracle is more efficient and faster than PowerExchange Oracle CDC with LogMiner in many Oracle environments.

**Important:** You cannot use both PowerExchange Express CDC for Oracle and PowerExchange Oracle CDC with LogMiner in the same PowerExchange instance with the same dbmover.cfg and pwxcl.cfg configuration files.

PowerExchange Express CDC requires the Oracle database to run in ARCHIVELOG mode with minimal global supplemental logging enabled. PowerExchange Express CDC can capture change data from most types of source environments, including those that use Oracle RAC and Automatic Storage Management (ASM), Oracle Data Guard logical or physical standby databases, Oracle Exadata database machines, or an Oracle multitenant pluggable database.

PowerExchange Express CDC runs on 64-bit operating systems only.

PowerExchange Express CDC does not use Oracle LogMiner to read the redo logs. It has its own log reader to get change data directly from the redo logs for real-time extraction processing. This architecture avoids the performance degradation that can occur on Oracle systems because of initialization of LogMiner sessions.

Also, PowerExchange Express CDC for Oracle does not use the PowerExchange UOW Cleanser. PowerExchange Express CDC capture components handle UOW Cleanser functions transparently for CDC sessions that run in real-time extraction mode. The capture components produce a change stream in which inserts, updates, and deletes are organized by transaction, and the transactions are arranged in their original commit order.

With PowerExchange Express CDC for Oracle, use of the PowerExchange Logger for Linux, UNIX, and Windows is optional but strongly recommended. The PowerExchange Logger writes successful units of work (UOWs) in chronological order by end time to PowerExchange Logger log files. CDC sessions can then extract the change data from the PowerExchange Logger log files in either continuous extraction mode or batch extraction mode. Benefits of using the PowerExchange Logger include fewer database accesses, faster CDC restart, and no need for prolonged retention of the Oracle redo files for CDC.

To configure PowerExchange Express CDC for Oracle, you must add some Oracle-specific statements to the dbmover configuration file. Also, define the separate PowerExchange Express CDC configuration file, which has the default file name of `pwxorad`. Other configuration tasks are similar to those for Oracle CDC with LogMiner. You create capture registrations, configure PowerCenter CDC sessions, and configure restart processing in the same manner as for Oracle CDC with LogMiner. The restart token format for PowerExchange Express CDC for Oracle sources is different from that for any other data source type, including Oracle CDC with LogMiner sources.

For tables that do not have row movement enabled, you can use the PowerExchange-generated `DTL_CAPXROWID` column in extraction maps and the `OPTIONS ROWID=Y` statement in the PowerExchange Express CDC for Oracle configuration file to include Oracle physical rowid values in change records. This feature is useful for processing rows in unkeyed tables during CDC extraction sessions.

If you currently use Oracle CDC with LogMiner and need better CDC performance, consider migrating to PowerExchange Express CDC for Oracle.

## PowerExchange Express CDC for Oracle Benefits

Consider the benefits of PowerExchange Express CDC for Oracle when comparing it to other Oracle change capture solutions.

- PowerExchange Express CDC for Oracle can perform CDC processing much faster than PowerExchange Oracle CDC with LogMiner in certain types of Oracle environments.
- PowerExchange Express CDC for Oracle processing is multithreaded to improve throughput.
- PowerExchange Express CDC for Oracle can efficiently process changes in environments that have a high volume of changes and large UOWs.
- PowerExchange Express CDC for Oracle supports RAC and ASM environments.
- PowerExchange Express CDC for Oracle avoids the LogMiner reinitialization issues and related degradation to Oracle system performance that can occur with PowerExchange Oracle CDC with LogMiner. Overall, PowerExchange Express CDC affects Oracle system overhead to a lesser degree.
- You do not need to configure periodic dumps of the Oracle data dictionary, as with PowerExchange Oracle CDC with LogMiner. PowerExchange Express CDC for Oracle copies the Oracle data dictionary into memory at initialization.

# PowerExchange Express CDC for Oracle Architecture

Example configurations demonstrate the general PowerExchange Express CDC for Oracle architecture.

The configurations include the following components:

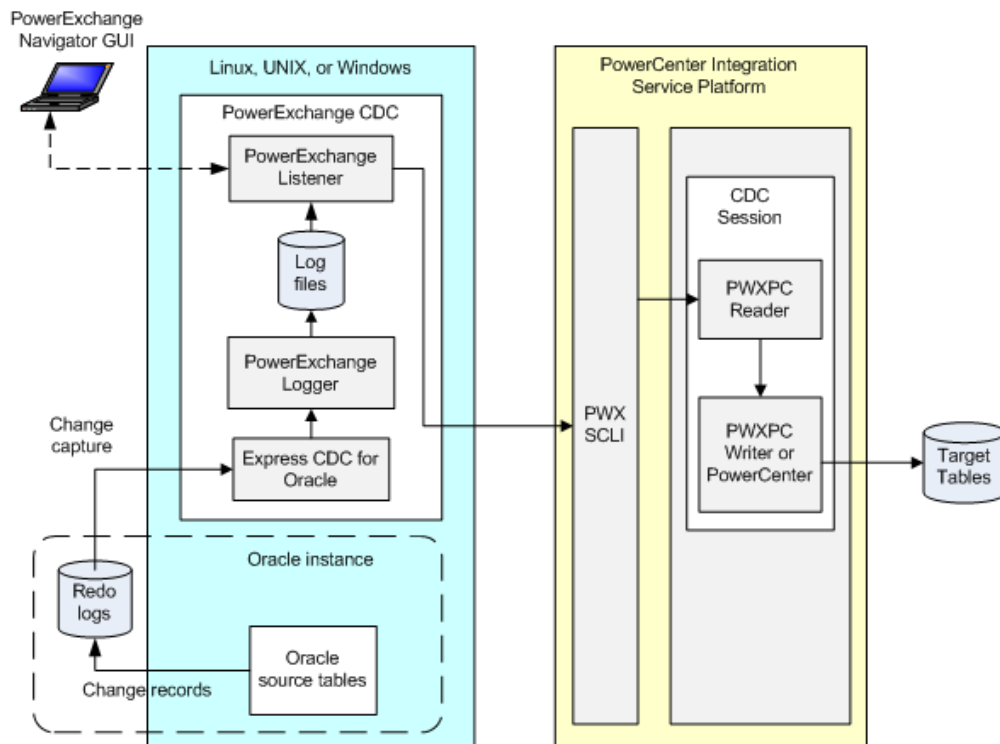
- Oracle source tables and redo logs
- PowerExchange Express CDC capture components, including the log reader
- PowerExchange Listener
- PowerExchange Logger for Linux, UNIX, and Windows, which is optional but strongly recommended
- PowerExchange Navigator
- PowerCenter and the PowerExchange Client for PowerCenter (PWXPC)

**Note:** PowerExchange Express CDC for Oracle does not use the PowerExchange UOW Cleanser.

## Configuration 1: All Capture Components on the Oracle System

If you have a Linux, UNIX, or Windows system that has sufficient CPUs and disk space, Informatica recommends that you run the Oracle instance, Oracle Express change capture components, PowerExchange Listener, and PowerExchange Logger all on that system. The PowerExchange Logger log files reside on the same system. This configuration avoids network-related degradation of performance and is the easiest one to configure and maintain.

The following figure shows Configuration 1:



The Express CDC log reader reads change records directly from the Oracle redo logs and forwards the committed changes to the PowerExchange Logger. The PowerExchange Logger logs the changes to its local log files. When a PowerCenter CDC session runs, change data is pulled from the PowerExchange Logger log files and sent to target tables, which are usually on a different system. Besides handling requests for change

data, the PowerExchange Listener handles PowerExchange Navigator and PWXPC requests for Oracle metadata or data, registrations, and extraction maps for other functions such as database row tests.

With this configuration, the PowerExchange Express CDC system contains the PowerExchange Express CDC for Oracle configuration file, the PowerExchange Logger configuration file, and the dbmover.cfg configuration file. The dbmover.cfg file includes the ORAD CAPI\_CONNECTION, CAPX CAPI\_CONNECTION, ORACLEID, and ORACLE\_CAPTURE\_TYPE statements. The dbmover.cfg file also includes the CAPX CAPI\_CONNECTION if you use continuous extraction mode.

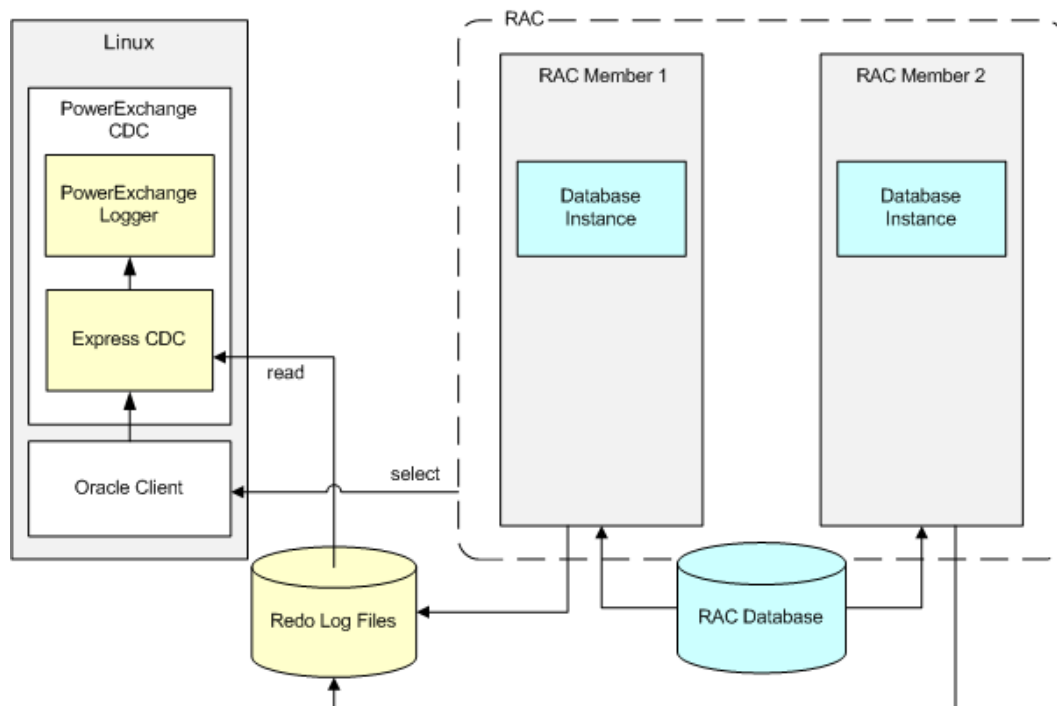
### Configuration 2: PowerExchange Express CDC in a RAC Environment Without ASM

If you want to capture change data in an Oracle RAC environment, Informatica recommends that you run the PowerExchange Express CDC capture process and PowerExchange Logger on a server outside of the RAC. With this configuration, if a RAC member node fails, Express CDC continues running, and Oracle establishes a connection to another active RAC member node.

In this configuration, the Express CDC log reader still reads change records directly from the Oracle redo logs and forwards the committed changes to the PowerExchange Logger. The archived and online redo logs must exist on a shared disk that can be accessed from the server where the PowerExchange Logger runs. The PowerExchange Logger logs the changes to its local log files. When PowerCenter CDC sessions run, PWXPC extracts changes from these log files.

To read change data, the PowerExchange Express CDC log reader must run under a user ID and password that has read access to the online and archived redo logs. Also, the Oracle Client must run under a user ID and password that has been granted SELECT authority on the appropriate database objects, as described in the PowerExchange ora\_orad.sql file.

The following figure shows a RAC with two member nodes and a separate Linux system with the PowerExchange Express CDC capture process and PowerExchange Logger:



**Note:** The PowerExchange Listener also runs on the Linux system.

In this scenario, a tnsnames.ora file resides on the Linux system. It specifies the FAILOVER option and the following connection descriptor that allows connection to either RAC member node:

```
ORATEST2=
  (DESCRIPTION=
    (FAILOVER=ON)
    (ADDRESS_LIST=
      (ADDRESS= (PROTOCOL=TCP) (HOST=rclnxxrac21.informatica.com) (PORT=1521))
      (ADDRESS= (PROTOCOL=TCP) (HOST=rclnxxrac22.informatica.com) (PORT=1521))
    )
    (CONNECT_DATA=
      (SERVICE_NAME=ORATEST2.informatica.com)
    )
  )
```

The following PowerExchange files also reside on the Linux system:

- CCT, CDEP, and CDCT files
- PowerExchange Logger log files
- PowerExchange Express CDC for Oracle configuration file
- A dbmover.cfg configuration file that includes the ORAD CAPI\_CONNECTION, CAPX CAPI\_CONNECTION, ORACLEID, and ORACLE\_CAPTURE\_TYPE statements

In the PowerExchange Express CDC for Oracle configuration file, you must specify the following statement for CDC in a RAC:

```
RAC MEMBERS=2;
```

**Note:** The MEMBERS parameter specifies the total number of member instances, including those that are not active.


### Configuration 3: PowerExchange Express CDC in a RAC Environment with ASM

If you want to capture change data in an Oracle RAC environment, Informatica recommends that you run the PowerExchange Express CDC capture process and PowerExchange Logger on a RAC member node with an ASM instance. This configuration provides the best performance.

The PowerExchange Express CDC log reader must be able to connect to the ASM instance as well as to the database to read the redo logs. A log reader that connects to an ASM instance must use an ASM login user ID that has SYSDBA or SYSASM authority. After the log reader reads the data, Express CDC forwards the committed changes to the PowerExchange Logger. The PowerExchange Logger logs the changes to its local log files. When PowerCenter CDC sessions run, PWXPC extracts changes from these log files.

In this sample configuration, the Express CDC log reader connects to an ASM instance and runs under a user ID and password that have SYSDBA authority. Because the Express log reader also connects to the Oracle database, Express CDC still requires a user ID and password that have the authorities described in the PowerExchange ora\_orad.sql file.

The following image shows a RAC with two member nodes, where each node has a database instance and an ASM instance and one node has the PowerExchange Express CDC capture process and PowerExchange Logger:



In this scenario, a `tnsnames.ora` file on RAC member 1 can specify the following ASM connection descriptor that allows connection to the ASM instance on either RAC member node:

```
ASMTst=
(DESCRIPTION=
  (FAILOVER=ON)
  (ADDRESS_LIST=
    (ADDRESS=(PROTOCOL=TCP) (HOST=rcln rac21.informatica.com) (PORT = 1521))
    (ADDRESS=(PROTOCOL=TCP) (HOST=rcln rac22.informatica.com) (PORT = 1521))
  )
  (CONNECT_DATA=
    (SERVICE_NAME=+ASM)
  )
)
```

The following PowerExchange files reside on RAC member 1, where Express CDC and the PowerExchange Logger run:

- CCT, CDEP, and CDCT files
- PowerExchange Logger log files
- PowerExchange Express CDC for Oracle configuration file
- A `dbmover.cfg` configuration file that includes the `ORAD CAPI_CONNECTION`, `CAPX CAPI_CONNECTION`, `ORACLEID`, and `ORACLE_CAPTURE_TYPE` statements



In the PowerExchange Express CDC for Oracle Configuration file, you must specify the **READER** statement with the **ASM** parameters and the **RAC** statement:

```
RAC MEMBERS=2;
READER
  MODE=ACTIVE
  ASM_CONNECT_STRING=tns_connect_string
  ASM_EPWD=encrypted_password|ASM_PASSWORD=password
  ASM_USERID=user_id
  other optional parameters;
```

**Note:** Specify either **ASM\_EPWD** or **ASM\_PASSWORD** but not both.

## PowerExchange Express CDC Configuration Overview

To implement PowerExchange Express CDC for Oracle, first gather information about your Oracle environment and review the planning considerations. Then complete the required configuration tasks in Oracle and PowerExchange.

### RELATED TOPICS:

- [“Gathering Information About the CDC Environment” on page 145](#)
- [“Planning Considerations” on page 147](#)
- [“Configuring Oracle for Express CDC” on page 159](#)
- [“Configuring PowerExchange for Express CDC” on page 162](#)
- [“Managing PowerExchange Express CDC for Oracle” on page 184](#)
- [“Oracle CDC with LogMiner” on page 109](#)
- [“PowerExchange Logger for Linux, UNIX, and Windows” on page 35](#)
- [“Introduction to Change Data Extraction” on page 208](#)

## Gathering Information About the CDC Environment

To prepare for implementation, gather information about your Oracle CDC environment.

Ask your Oracle DBAs the following questions:

**What is the Oracle database name?**

Answer:

**For which Oracle tables do you need to capture change data?**

Answer:

**Do any of the Oracle objects that will be involved in CDC conflict with PowerExchange Express CDC for Oracle restrictions?**

Answer:

**Is ARCHIVELOG mode and minimal global supplemental logging enabled for the Oracle source database? If not, can they be enabled?**

Answer:

**Do you have read access to the redo logs in your environment?**

Answer:

**If you do not have the authority to read the redo logs directly, can the archived redo log files be copied to a file system from which you can access them?**

Answer:

**Can you create a new Oracle user with the privileges that PowerExchange Express CDC for Oracle requires? What user name do you want to use?**

Answer:

**Are the redo logs in ASM-managed storage? If you plan to connect to an ASM instance to read redo logs, are you allowed to create a login user ID for ASM that has SYSDBA or SYSASM authority?**

You must be able to connect to an ASM instance as well as to the database to read redo logs. The ASM login user ID must have SYSDBA authority, or if you have Oracle 11g, the user ID can have SYSASM authority.

Answer:

**Can you make your archived redo logs available for diagnostic use by Informatica Global Customer Support if an error or anomaly occurs during CDC processing?**

Answer:

**Do you need to capture changes from a RAC? How many member nodes does the RAC contain, including inactive nodes?**

Answer:

**What is the average amount of archive log that is created per hour during peak and non-peak periods for the Oracle database?**

Answer:

**What is the typical size of units of work (UOWs) for the source tables?**

Answer:

**Does the Oracle system have the capacity to run PowerExchange Express CDC for Oracle locally?**

Answer:

**Do you have unkeyed source tables for which you want to capture unique rowids?**

Answer:

**Do you need to capture changes for NUMBER columns that have a precision greater than 28 or an undefined length?**

Answer:

**Do you need to capture changes for FLOAT columns that have a precision greater than 15?**

Answer:

**Do you need to capture change data from an Oracle Data Guard logical or physical standby database?**

Answer:

**Do you need to capture change data from tablespaces that use Oracle Transparent Data Encryption (TDE)? If yes, what is the TDE password?**

Answer:

**Do you need to capture change data from Oracle index-organized tables?**

Answer:

**Do you need to capture change data from a pluggable database (PDB) in an Oracle multitenant environment?**

Answer:

**Do you need to capture Oracle direct-path operations?**

Answer:

## Planning Considerations

Before you configure PowerExchange Express CDC for Oracle, review the restrictions, operational and performance considerations, and supported datatypes table. This information can help you successfully configure and use this CDC solution.

Also, verify that a valid Oracle environment exists for the PowerExchange user. On Linux, and UNIX, the path to the Oracle client must be specified in the PATH and library path environment variables.

### RELATED TOPICS:

- [“PowerExchange Express CDC for Oracle Restrictions” on page 147](#)
- [“Operational Considerations” on page 150](#)
- [“Performance Considerations” on page 153](#)
- [“Oracle Datatypes Supported for Express CDC” on page 148](#)
- [“ASM Considerations” on page 154](#)
- [“RAC Considerations” on page 153](#)

## PowerExchange Express CDC for Oracle Restrictions

The following restrictions apply to PowerExchange Express CDC for Oracle:

- PowerExchange Express CDC for Oracle does not support the following Oracle environments, objects, and features:
  - Any type of encryption of tablespaces other than Oracle Advanced Security Transparent Data Encryption (TDE) of tablespaces.
  - Tables in a sorted hash clusters
  - Columns that use TDE
  - Virtual columns, which contain derived data that Oracle does not log in the redo logs
  - Columns that have unsupported datatypesBecause you cannot include these columns in capture registrations, PowerExchange Express CDC for Oracle does not capture change data for them. However, PowerExchange Express CDC for Oracle can capture change data for other columns in the same registered table.
- PowerExchange Express CDC for Oracle does not capture the following types of operations:
  - Direct-path INSERTs
  - INSERT AS SELECT, MERGE, and PL/SQL FORALL operations that use the APPEND hint or parallel mode
  - SQL\*Loader direct-path loads
  - CREATE TABLE...AS SELECT operations

- PowerExchange Express CDC for Oracle can capture data that the SQL\*Loader utility loaded into Oracle tables with the following restrictions:
  - The load type must be *conventional path*.
  - The load method must be Insert, Append, or Replace. Do not use Truncate. If you use Truncate, the SQL\*Loader issues TRUNCATE TABLE DDL. PowerExchange cannot capture row deletions that result from TRUNCATE TABLE DDL.
- If a PowerExchange Express CDC source instance has redo logs on a RAW device and the logs are *not* managed by ASM, the PowerExchange Logger for Linux, UNIX, and Windows and database row tests fail when they try to process the logs.
- If a source table contains LOB columns that are defined with the STORE AS SECUREFILE storage parameter, PowerExchange Express CDC for Oracle cannot capture change data from these LOB columns but can capture change data from the other columns in the table.

## Oracle Datatypes Supported for Express CDC

Verify that columns in the Oracle source tables from which you plan to capture change data have datatypes that PowerExchange Express CDC for Oracle supports.

Oracle does not log, or does not completely log, data from columns with certain datatypes in the Oracle redo logs. Consequently, PowerExchange cannot retrieve change data for these columns.

The following table identifies the Oracle datatypes that PowerExchange Express CDC for Oracle supports and does not support:

Datatype	Supported for CDC?	Comments
BFILE	No	Column data that has this datatype is not completely logged in the Oracle redo logs and cannot be captured.
BINARY_DOUBLE	Yes	-
BINARY_FLOAT	Yes	-
CHAR	Yes	-
DATE	Yes	-
FLOAT	Yes	If you set the ORACLE_UNHANDLED_NUMASCHAR statement to Y in the dbmover.cfg file, PowerExchange registration processing for Oracle CDC source tables converts FLOAT columns that have a precision greater than 15 to VARCHAR columns. This setting can prevent data loss.
LOBs	No	-
LONG	No	Columns that have this datatype cannot be included in capture registrations.
LONG RAW	No	Columns that have this datatype cannot be included in capture registrations.
NCHAR	Yes	-

Datatype	Supported for CDC?	Comments
NUMBER	Yes	<p>PowerExchange handles NUMBER columns as follows:</p> <ul style="list-style-type: none"> <li>- Numbers that have a precision value less than 10 and a scale of 0 are treated as INTEGER.</li> <li>- Numbers that have a defined precision and scale are treated as NUMCHAR.</li> <li>- Numbers that have an undefined precision and scale are treated as double-precision floating point numbers by default.</li> </ul> <p>If you set the ORACLE_UNHANDLED_NUMASCHAR statement to Y in the dbmover.cfg file, PowerExchange registration processing for Oracle CDC source tables treats NUMBER columns that have a precision greater than 28 or an undefined length as variable length strings. This setting can prevent data loss.</p>
NVARCHAR2	Yes	-
RAW	Yes	-
ROWID	Yes	-
TIMESTAMP	Yes	-
TIMESTAMP WITH TIME ZONE	Yes	<p>PowerExchange captures data with this datatype as a UTC timestamp.</p> <p><b>Note:</b> PowerCenter does not support this datatype. If you use PowerCenter to materialize a target table from a source table that includes this datatype, manually override the datatype in Source Analyzer with the timestamp datatype. Also, edit the generated SQL select statement that PowerCenter sends to PowerExchange to use the sys_extract_utc() function. Syntax:</p> <pre>select sys_extract_utc(tmstmpwith_tz) from schema.source_table</pre>
TIMESTAMP WITH LOCAL TIME ZONE	Yes	<p>PowerExchange captures data with this datatype as a UTC timestamp.</p> <p><b>Note:</b> PowerCenter does not support this datatype. If you use PowerCenter to materialize a target table from a source table that includes this datatype, manually override the datatype in Source Analyzer with the timestamp datatype. Also, edit the generated SQL select statement that PowerCenter sends to PowerExchange to use the sys_extract_utc() function. Syntax:</p> <pre>select sys_extract_utc(tmstmpwith_tz) from schema.source_table</pre>
UROWID	No	-
User-defined types (UDTs)	No	-
VARCHAR2	Yes	-
XML types	No	-

**Note:** PowerExchange Express CDC for Oracle also does not support virtual columns and columns that have user-defined types (UDTs). Because you cannot include these columns in capture registrations, PowerExchange does not capture change data for them. However, PowerExchange can capture change data for other columns in the same registered table.

## Operational Considerations

Review the following operational considerations for PowerExchange Express CDC for Oracle:

- Use of the PowerExchange Logger for Linux, UNIX, and Windows is optional. However, Informatica strongly recommends that you use it to reduce or avoid CDC impacts on your system, such as increased database I/O and the need to retain archive logs longer than normal. You can run the PowerExchange Logger in batch or continuous mode.
- The maximum length of a row for which PowerExchange can capture and process change data is 128,000 bytes.
- The PowerExchange Express CDC for Oracle capture process and the Oracle source instance that contains the redo logs can run on machines that have different operating systems and architectures. You can use any combination of supported Linux, UNIX, and Windows operating systems.

To capture changes across heterogeneous platforms, the Oracle source tables must be registered on the Express CDC system.

- PowerExchange Express for Oracle can read Oracle redo logs that are in ASM-managed storage, on a standard file system, or on a network file system (NFS). However, if the logs are on an NFS and the NFS buffers become stale, Express CDC might fail with error message PWX-36171 when reading the active logs. To avoid this problem, use one of the following options:
  - Configure the Oracle database to write active redo logs to ASM-managed storage.
  - Disable PowerExchange Express CDC use of active redo logs by setting the `READER MODE` parameter to `ARCHIVEONLY` or `ARCHIVECOPY` in the `pwxorad.cfg` file.
  - When you configure the NFS mount point for active redo logs, disable read buffering and attribute caching.

**Note:** On Linux, PowerExchange Express CDC for Oracle cannot guarantee support for the combination of Oracle direct I/O writes and the reading of active redo logs over NFS. The Linux NFS server does not honor the request to open files by using direct I/O. The combination of direct I/O writes and buffered reads might lead to a corrupted page cache. If PowerExchange Express CDC detects a corrupted page cache, Express CDC fails with error message PWX-36171. Under some circumstances, duplexing the active redo logs might allow PowerExchange Express CDC to recover from the failure. To avoid the failure completely, you might need to disable Oracle direct I/O by setting the Oracle `FILESYSTEMIO_OPTIONS` parameter to `none`.

- Oracle `ARCHIVELOG` mode and minimal global supplemental logging must be enabled. Also, the registered Oracle source table columns for which you want to capture change data must be associated with an unconditional supplemental log group, also called an `ALWAYS` supplemental log group. With unconditional supplemental logging, Oracle logs before images of column data to redo logs whenever any column in a row is updated. PowerExchange Express CDC for Oracle requires before images to properly process updates.
- You cannot run both PowerExchange Express CDC for Oracle and PowerExchange Oracle CDC with LogMiner with the same PowerExchange instance and the same PowerExchange Listener and PowerExchange Logger. If you need to use both of these Oracle CDC solutions, run each one in a separate PowerExchange instance with different `dbmover.cfg` and `pwxccl.cfg` configuration files. However, you can use the same capture registrations for both CDC solutions if you want to capture change data from the same Oracle database and tables. For these CDC solutions to share registration information, store the CCT file in a common location.
- The PowerExchange Express CDC for Oracle log reader must have access to the Oracle redo log files. If the Oracle redo log files are not stored in ASM, the operating system user ID under which PowerExchange Express CDC for Oracle log reader runs must have the authority to read the redo logs or copies of the archived redo logs. If the redo log files are stored in Oracle 11g ASM and you plan to connect to the ASM instance to read the logs, Express CDC requires an Oracle user ID with `SYSDBA` or `SYSASM` authority to

connect to the ASM instance to get change data. If you use an earlier Oracle version, PowerExchange requires an ASM user ID that has SYSDBA authority to connect to ASM.

- PowerExchange Express CDC for Oracles reads the active and archived redo logs by default. If you want Express CDC to read only the archived redo logs, set the MODE parameter to ARCHIVEONLY in the Express CDC configuration file. If you do not have the authority to read the active or archived redo logs, you can configure Express CDC to read copies of the archived logs that are located on a file system. In this case, set the MODE parameter to ARCHIVECOPY.
- If a CDC problem occurs, Informatica Global Customer Support might need to request the Oracle archived redo logs from which changes are captured for diagnostic use.
- You must use the PowerExchange Client for PowerCenter (PWXPC) to integrate with PowerCenter. The PowerExchange ODBC driver does not support PowerExchange Express CDC for Oracle.
- If you use Oracle materialized views, PowerExchange can capture change data from the master tables that underlie those views. PowerExchange supports change capture for any type of materialized view. The view and its underlying table have a one-to-one correspondence and share the same name. If you issue DTLDESCRIBE tables from the **Database Row Test** dialog box in the PowerExchange Navigator, the results include a row for the materialized view and a row for the underlying table. The **Type** column indicates which row is for the materialized view and which row is for the table.
- PowerExchange Express CDC for Oracle uses the standard PowerExchange restart and recovery processing for relational sources. The format of PowerExchange Express CDC for Oracle restart tokens is different from that for Oracle CDC with LogMiner sources or for any other data source type. The default restart point for PowerExchange Express CDC for Oracle is one of the following points in the change stream:
  - For CDC sessions that run in batch extraction mode or continuous extraction mode, the beginning of the oldest PowerExchange Logger log file that is recorded in the CDCT file.
  - For CDC sessions that run in real-time extraction mode, the beginning of last log sequence that was archived.
- If you capture change data only from archived redo logs and do not use the PowerExchange Logger, the Express CDC log reader uses one of the following default start points at initialization:
  - For non-RAC instances, the low SCN of the last available archive log.
  - For RAC instances, the highest low SCN of the last log archived across all active nodes.If you capture change data only from archived redo logs and use the PowerExchange Logger, as recommended, the default restart point after a PowerExchange Logger cold start is the end of the last available archived log, also referred to as current end-of log (EOL), unless you set the RESTART\_TOKEN and SEQUENCE\_TOKEN parameters in the pwxcl.cfg file. In a RAC environment, the EOL is considered to be the lowest high SCN of the last log archived across all active nodes.
- PowerExchange Express CDC for Oracle can capture data from Oracle Exadata database machines, subject to the restrictions in [“PowerExchange Express CDC for Oracle Restrictions” on page 147](#).
- For tables that do not have row movement enabled, you can use the PowerExchange-generated DTL\_\_CAPXROWID column in extraction maps and the OPTIONS ROWID=Y statement in the PowerExchange Express CDC for Oracle configuration file to include Oracle physical rowid values in change records. This feature is useful for processing rows in unkeyed tables during CDC extraction sessions.
- PowerExchange Express CDC for Oracle can capture Oracle direct-path operations except for tables that use Oracle Exadata Hybrid Columnar Compression (EHCC). To capture direct-path operations, you must set the SUPPORT\_DIRECT\_PATH\_OPS parameter to Y in the OPTIONS statement of the PowerExchange Express CDC for Oracle configuration file.

- In Oracle 11g Release 2 and later environments, PowerExchange Express CDC for Oracle can capture conventional and direct-path DML changes from tables and table partitions and subpartitions that use Oracle Advanced Compression.
- In Oracle 11g Release 2 and later environments, PowerExchange Express CDC for Oracle can capture conventional DML changes from tables and table partitions and subpartitions that use Oracle Exadata Hybrid Columnar Compression (EHCC). However, Express CDC does *not* capture Oracle direct-path operations for objects that use EHCC.
- PowerExchange Express CDC for Oracle can capture change data from Oracle tablespaces that use Oracle Advanced Security Transparent Data Encryption (TDE). To capture TDE-encrypted change data from tablespaces, you must perform the following tasks:
  - Enter either the TDEWALLETPWD or TDEWALLETEPWD parameter in the pwxorad.cfg configuration file to specify the password for the Oracle TDE wallet.
  - Grant the following privilege to the ORACAPTL user:
 

```
GRANT SELECT ON "PUBLIC"."V$ENCRYPTION_WALLET" TO "ORACAPTL";
```
  - Ensure that the Oracle TDE wallet is on a device that PowerExchange Express CDC for Oracle can access with Read file permissions.
  - Verify that the Oracle TDE wallet is open in the database.
- PowerExchange Express CDC for Oracle can capture changes from Oracle Data Guard logical and physical standby databases. PowerExchange Express CDC supports any configuration of primary and standby databases that Oracle Data Guard supports, including RAC databases in an ASM environment. The number of nodes on the primary and standby systems do not need to match. For example, the primary system can contain a RAC with multiple member instances, and the standby system can contain a single non-RAC instance. To handle database role transitions that involve a physical standby database, you might need to update some statements in the PowerExchange Express CDC configuration file. For more information, see [“Oracle Data Guard Physical Standby Databases as Sources” on page 155](#).
- If an Oracle RESETLOGS event occurs on a source database, PowerExchange Express CDC can continue change capture processing across the RESETLOGS boundary in the archived redo logs. A RESETLOGS event occurs in situations that require you to open the database with the RESETLOGS option, such as after a flashback database operation, incomplete point-in-time recovery, or point-in-time recovery with a backup control file. A RESETLOGS event can also occur transparently in a Data Guard environment with a physical standby database after a failover or after a switchover that is preceded by an incomplete recovery and followed by an ALTER DATABASE ACTIVATE PHYSICAL STANDBY DATABASE operation. A RESETLOGS operation archives the current online redo logs, resets the log sequence number to 1, creates a new database incarnation, creates a new timestamp and SCN for the online redo logs, and updates all the current data files with the new RESETLOGS SCN.
 

The PowerExchange Express CDC restart token includes the resetlogs ID to identify the database incarnation to use for restart processing. At initialization, PowerExchange Express CDC uses the resetlogs ID to check whether the database has undergone a RESETLOGS event. If a RESETLOGS event occurred, PowerExchange Express CDC verifies that the restart information and the last change data that was hardened to the PowerExchange Logger log files are valid and have not been orphaned by the event. PowerExchange Express CDC then continues capture processing.
- PowerExchange Express CDC for Oracle can capture changes from Oracle index-organized tables (IOTs).
- If you use the Oracle Multitenant option, which was introduced in Oracle 12c, PowerExchange Express CDC for Oracle can capture change data from a pluggable database (PDB) within an Oracle multitenant container database (CDB). PowerExchange Express CDC can capture data from only a single PDB at a time. For more information, see [“Oracle Multitenant Pluggable Databases as Sources” on page 157](#).



- PowerExchange Express CDC can capture change from an Oracle database in an Amazon Elastic Compute Cloud (EC2) environment. The Amazon EC2 instances must run on a 64-bit Red Hat Linux server. Express CDC does not support EC2 instances on 64-bit Windows servers. The configuration of PowerExchange, the Oracle database, the PowerCenter Integration Service, and Informatica domain is flexible. All of these applications can run on EC2 instances in the cloud, or some of them can run on premise. Any combination of these applications in the cloud and on premise is supported. To capture change data, PowerExchange Express CDC requires no special configuration tasks. As usual, ensure that Express CDC can access to the Oracle archived redo logs. Also, ensure that the Express CDC user has the same privileges as are required for on-premise CDC processing.

## Performance Considerations

The following considerations pertain to PowerExchange Express CDC for Oracle performance:

- If the memory that PowerExchange Express CDC for Oracle uses to stage change records becomes full, for example, because of large UOWs or a high transaction volume, PowerExchange Express CDC for Oracle can spill additional change records to a temporary spill file on disk. If CDC performance becomes degraded, increase the MEMOPS parameter value in the OPTIONS statement of the PowerExchange Express CDC for Oracle configuration file to a level that is sufficient for the change data volume.
- If you use continuous extraction mode, minimize the size of the CDCT file. The CDCT file contains information about PowerExchange Logger log files. PowerExchange reads the CDCT file each time the interval that is specified in the FILEWAIT parameter of the CAPX CAPI\_CONNECTION statement elapses. If a CDCT file is large, PowerExchange read operations can result in a high level of I/O activity, increased use of system resources, and increased extraction latency. To manage the CDCT file size, use the COND\_CDCT\_RET\_P statement in the pwxcl.cfg configuration file for the PowerExchange Logger.

## RAC Considerations

PowerExchange Express CDC for Oracle can capture changes from online and archived redo log files in an Oracle RAC environment.

The PowerExchange Express CDC capture process and PowerExchange Logger for Linux, UNIX, and Windows can run on a RAC member instance or on a server outside of the RAC. In both cases, PowerExchange must have read access to the Oracle online and archived redo logs. For greater resiliency, run the Express CDC capture process and PowerExchange Logger on a server outside of the RAC. With this configuration, if the RAC member node to which PowerExchange is connected fails, Express CDC continues running and tries to connect again. Oracle directs the connection request to another active RAC member that is referenced by an entry in the an Oracle tnsnames.ora file.

To configure PowerExchange for CDC in a RAC, you must define the RAC statement with the MEMBERS parameter in the PowerExchange Express CDC for Oracle configuration file. For the MEMBERS value, enter the total number of member instances in the RAC, including both active and inactive instances.

If you do not use ASM and you run the PowerExchange Express CDC capture process and PowerExchange Logger on a server outside of the RAC, you might also need to define the DIRSUB statement. Define the DIRSUB statement if the server where PowerExchange Express capture process runs uses a mount point to the directory with the online and archived redo logs that is different from the mount point that the RAC Oracle instance to which PowerExchange connects uses.

If you use the FAILOVER feature of Oracle, you can define a single entry in the tnsnames.ora file that covers multiple RAC member instances. Informatica recommends that you use the FAILOVER feature to prevent

PowerExchange Express CDC for Oracle from failing if a RAC member node stops running. The following sample entry in the `tnsnames.ora` file has the `FAILOVER` option enabled and covers two RAC nodes:

```
ORA1A=
  (DESCRIPTION=
    (FAILOVER=ON)
    (ADDRESS_LIST=
      (ADDRESS=(PROTOCOL=TCP) (HOST=node1.informatica.com) (PORT=1521))
      (ADDRESS=(PROTOCOL=TCP) (HOST=node2.informatica.com) (PORT=1521))
    )
    (CONNECT_DATA=
      (SERVICE_NAME=ORA1A.informatica.com)
    )
  )
```

If you use ASM, use the following equivalent entry in the `tnsnames.ora` file:

```
ASMAAny=
  (DESCRIPTION=
    (FAILOVER=ON)
    (ADDRESS_LIST=
      (ADDRESS=(PROTOCOL=TCP) (HOST=node1.informatica.com) (PORT = 1521))
      (ADDRESS=(PROTOCOL=TCP) (HOST=node2.informatica.com) (PORT = 1521))
    )
    (CONNECT_DATA=
      (SERVICE_NAME=+ASM)
    )
  )
```

**Note:** `SERVICE_NAME` must specify a net service name of `+ASM`.

Also, PowerExchange Express CDC assumes that, for all of the open and closed threads, each is associated with a RAC node. If the number of open and closed threads does not match the Oracle `CLUSTER_DATABASE_INSTANCES` parameter value, Express CDC issues the following error message:

```
PWX-36127 ORAD: Enabled thread count does not match instance count. Enabled threads =
<number_of_threads> : Cluster Instances = <number_of_instances>.
```

If you receive this error, disable the redundant threads by using the following SQL statement:

```
ALTER DATABASE DISABLE THREAD number_of_redundant_threads
```

## ASM Considerations

PowerExchange Express CDC for Oracle can capture changes from online and archived redo log files that are managed by Oracle Automatic Storage Management (ASM).

Configure the PowerExchange Express CDC for Oracle reader to connect to the ASM instance to read the redo logs. PowerExchange establishes connections to both the ASM instance and Oracle database. The PowerExchange Express CDC log reader connects to the ASM instance based on the `ASM_CONNECT_STRING`, `ASM_USERID`, `ASM_EPWD` or `ASM_PASSWORD`, and `ASM_ASSYSASM` parameters that you specify in the `READER` statement. In the `ASM_USERID` parameter, specify a user ID that has `SYSDBA` authority. Or if you use Oracle 11g, you can specify a user ID that has `SYSASM` authority. To use `SYSASM` authority, you must set the `ASM_ASSYSASM` parameter to `Y`. Because of the required authority level, usually the ASM user ID is different from any of the following user IDs:

- If you extract data to PowerExchange Logger for Linux, UNIX, and Windows log files, as recommended:
  - The `USERID` in the `ORAD CAPI_CONNECTION` statement in the `dbmover.cfg` file.
  - The `CAPTURE_NODE_UID` in the PowerExchange Logger configuration file.
- If you extract data directly to PowerCenter:
  - The **User Name** attribute on the `PWX CDC Real Time` application connection for the PowerCenter CDC session.

For optimal CDC performance, run the Express CDC capture process and the PowerExchange Logger for Linux, UNIX, and Windows on the same node as the ASM instance.

To define a bequeth connection, enter statements in the Oracle tnsnames.ora file such as the following sample statements:

```
BEQ_ASM =
  (DESCRIPTION=
    (ADDRESS=
      (PROTOCOL=beq)
      (PROGRAM=/oracle/product/11.1.0.7/asm/bin/oracle)
      (ARGV0=+ASM1)
      (ARGS=' (DESCRIPTION=
        (LOCAL=YES)
        (ADDRESS= (PROTOCOL=beq) )
      ) '
    )
  )
(CONNECT_DATA= (SID=+ASM1) )
```

## Oracle Data Guard Physical Standby Databases as Sources

PowerExchange Express CDC for Oracle can capture change data from Oracle Data Guard physical standby databases. PowerExchange monitors the standby and archived redo logs and the database SCN on the standby system. As long as the database SCN is progressing, PowerExchange captures change data from the logs.

A PowerExchange CDC environment with a physical standby database source has the following characteristics:

- The primary database and the physical standby databases must use Oracle 11g or later.
- You must install and run PowerExchange on the machine with the physical standby database. A PowerExchange installation is not required on the primary database machine.
- The physical standby database can use real-time apply with standby redo logs or apply data directly from the archived logs only.
- The physical standby database can be open with read-only access, or it can be not open, such as when started with the mount option.
- PowerExchange Express CDC supports any configuration of primary and standby databases that Oracle Data Guard supports, including databases in RACs that use ASM. The number of nodes on the primary and standby systems do not need to match.

To configure change capture from a physical standby database, you must complete the following configuration tasks:

- In the PowerExchange Navigator client, create a registration group and capture registrations for the Oracle source instance and tables on the primary system from which the changes originate.
- In the dbmover.cfg configuration file on the standby system, configure the following statements to point to the Oracle instance on the primary system:
  - The ORACOLL parameter in the ORAD CAPI\_CONNECTION statement
  - The DFLTINST parameter in the CAPX CAPI\_CONNECTION statement
  - The first positional parameter, *collection\_id*, in the ORACLEID statement

**Important:** In the ORACLEID statement, do *not* specify the fourth positional parameter *capture\_connect\_string*. This information is provided by the DATABASE or STANDBY statement in the pwxorad.cfg file.

- In the Express CDC `pwxorad.cfg` configuration file on the standby system, define the following statements if the standby database is open for read-only access:
  - Define a `DATABASE` statement that provides connection information for the standby system.
  - If the primary database is in a RAC, define a `RAC` statement. PowerExchange verifies the `RAC MEMBERS` parameter value against the number of threads on the standby or archived redo logs. The `RAC MEMBERS` value must be equal to the number of threads.
- If the standby database is *not* open for read-only access, define the following parameters:
  - Define a `DATABASE` statement that provides connection information for the primary system. PowerExchange requires this connection to access the Oracle data dictionary on the primary system.
  - Define a `STANDBY` statement that provides connection information for the standby system. The user that you specify for the `STANDBY` connection must have `SYSDBA` authority to access the fixed views and monitor the Data Guard apply progress in a database that is not open.
- In the PowerExchange Logger `pwxcl.cfg` configuration file on the standby system, configure the `DBID` parameter to specify the Oracle collection ID that is defined in the registration group for the registered source tables on the primary system.

After CDC is running, the following operational considerations might apply:

- When the standby database is *not* open for read-only access, Oracle might not progress apply processing beyond the tip of the last archived log, even if more recent changes are available in the standby redo logs. To get near-real-time capture processing, add `APPLYACTIVE=Y` in the `STANDBY` statement. PowerExchange will then read change data from the standby redo logs up to the highest of the low SCN values in the standby logs for all threads.
- PowerExchange CDC captures data up to the current database SCN, or apply SCN, for a physical standby database. If Oracle stops applying data for some reason, for example, because of a log gap, PowerExchange change capture processing stalls and waits for the apply process to resume with message `PWX-36098`.
- Before you perform a switchover, stop change data capture.

If PowerExchange Express CDC is configured to access a standby instance that is open for read-only access, you might need to add or edit the `RAC MEMBERS` parameter to properly resume capturing data after a switchover. If the thread count on the original standby database does not match the thread count on the original primary database, set the `RAC MEMBERS` parameter value in the `pwxorad.cfg` file to the number of threads on the original standby database. After the switchover, the new primary database (former standby database) will have the specified number of threads. If the original standby database is not in a RAC, you can omit the `RAC MEMBERS` parameter.

If PowerExchange Express CDC is configured to access a standby instance that is in mount mode (not open), make the following changes in the `pwxorad.cfg` file:

- Remove the `STANDBY` statement.
- Edit the `DATABASE` statement such that PowerExchange connects directly to the new primary instance.
- If the thread count on the original standby database does not match the thread count on the original primary database, set the `RAC MEMBERS` parameter value to the number of threads on the original standby database. After the switchover, the new primary database (former standby database) will have the specified number of threads. If the original standby database is not in a RAC, you can omit the `RAC MEMBERS` parameter.
- PowerExchange Express CDC for Oracle does not support reading Oracle logs across a reset logs boundary. If a `RESETLOGS` operation is performed when a failover to an Oracle Data Guard physical standby database occurs after an unplanned outage, you must special start the PowerExchange Logger for Linux, UNIX, and Windows to minimize data loss. Data that was captured prior to the special start is preserved, but any transactions that were in flight when the outage occurred are lost.

## Oracle Multitenant Pluggable Databases as Sources

Oracle 12c introduced multitenant container databases. If you have an Oracle version that supports the Oracle Multitenant option, you can use PowerExchange Express CDC for Oracle to capture change data for source tables that reside in a pluggable database (PDB) within a multitenant container database (CDB).

The following CDC considerations apply to Oracle multitenant environments:

- PowerExchange Express CDC can capture change data for only a single PDB within a CDB at a time. If you want to capture change data for another PDB in the same CDB, you must configure a separate CDC environment. Informatica recommends that each CDC environment use a unique Oracle capture user ID, pwxorad.cfg file, PowerExchange Logger for Linux, UNIX, and Windows instance and pwxcl.cfg file, and capture registrations.
- If you move or clone the PDB for which PowerExchange Express CDC is capturing change data to another CDB, the PowerExchange Logger for Linux, UNIX, and Windows connection to the Oracle database is lost. You must cold start the PowerExchange Logger to prevent loss of change data.

### Task Flow: Configuring Change Capture from an Oracle Multitenant Pluggable Database

To capture change data from a pluggable database (PDB) in an Oracle 12c container database (CDB), perform the following tasks:

This is the start of the concept.

1. Verify that the PDB is open.
2. Verify that the CDB is running in ARCHIVELOG mode.
3. Add a PDB entry that includes the PDB service name to the tnsnames.ora file, if this entry does not already exist. For example:

```
PDB1234=
  (DESCRIPTION=
    (ADDRESS= (PROTOCOL=TCP) (HOST=host1) (PORT=1521))
    (CONNECT_DATA=
      (SERVER=DEDICATED)
      (SERVICE_NAME=pdb1234.informatica.com)))
```

4. Enable minimal global supplemental logging for the PDB or CDB. Use the following SQL statement in the ora\_orad.sql file:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
COMMIT;
```

5. Recommended. Connect to the CDB\$ROOT container as a user who has the DBA role. Alternatively, log in to the PDB directly.

If you connect to CDB\$ROOT with the DBA role, execute the following SQL statement to switch to the source PDB:

```
ALTER SESSION SET CONTAINER=pdb_name
```

6. Create a unique Oracle user ID for CDC and grant the privileges that are required for change data capture from a PDB to this user.

Use the CREATE USER and GRANT statements in the ora\_orad.sql file. This file uses the user name of "ORACPTL1." You can change this name if you want to use another user name.

Make sure that you issue the following GRANT statements that are required for PowerExchange Express CDC in Oracle multitenant environments:

```
GRANT CREATE SESSION TO "ORACPTL1";
GRANT SELECT on "PUBLIC"."V$PDBS" TO "ORACPTL1";
```

7. In the dbmover.cfg configuration file on the system where change capture occurs, add an ORACLEID statement that points to the name of the database that includes the PDB and the name of the PDB service entry in the tnsnames.ora file. For example:

```
ORACLEID=(PDB1234,ORADBNAM, ,PDB1234)
```

In this example statement:

- The first positional parameter is a *collection\_id* value that identifies the ORACLEID statement. In this case, it is also the PDB name.
- The second positional parameter is the name of the Oracle database that contains the PDB.
- The third positional parameter is empty because it is not used for an Oracle multitenant source.
- The fourth positional parameter is the name of the PDB service entry in the tnsnames.ora file.

Also add an ORAD CAPI\_CONNECTION statement that includes an ORACOLL value that matches the *collection\_id* value in the ORACLEID statement.

8. In the PowerExchange Logger for Linux, UNIX, and Windows configuration file, pwxcl.cfg, ensure that the DBID parameter matches the *collection\_id* value in the ORACLEID statement in the dbmover.cfg file.
9. Optional. If you specify the DATABASE statement in the PowerExchange Express CDC for Oracle configuration file, pwxorad.cfg, ensure that the CONNECT\_STRING parameter points to the name of the PDB service entry in the tnsnames.ora file. Also, ensure that the USERID parameter specifies the Oracle user name that you defined for the change capture user.
10. In the PowerExchange Navigator, create a registration group. Then add a capture registration for each source table in the PDB to the registration group.

**Note:** The **Collection ID** value that you specify for the registration group must match the *collection\_id* value in the ORACLEID statement in the dbmover.cfg file.

## Summary Task Flow

To configure and start a new PowerExchange Express CDC for Oracle environment, complete the following high-level tasks:

1. Configure Oracle for CDC.  
Use the sample SQL in the ora\_orad.sql file that PowerExchange provides. For more information, see [“Configuring Oracle for Express CDC” on page 159](#).
2. Configure PowerExchange for CDC.  
For more information, see [“Configuring PowerExchange for Express CDC” on page 162](#).
3. Configure the PowerExchange Logger for Linux, UNIX, and Windows.  
Use of the PowerExchange Logger is optional but strongly recommended. For information about configuring and starting the PowerExchange Logger, see [“Configuring the PowerExchange Logger” on page 43](#) and [“Starting the PowerExchange Logger” on page 63](#).
4. Configure a CDC restart point.  
For more information, see [“Creating Restart Tokens for Extractions” on page 241](#).
5. Materialize the target tables.
6. Start the PowerExchange Logger.
7. Create PowerCenter CDC workflows that include the Express CDC source tables.  
Use an Oracle CDC application connection. For more information, see the *PowerExchange Interfaces for PowerCenter* publication.
8. Cold start the workflows.

If you need to migrate from PowerExchange Oracle CDC with LogMiner to PowerExchange Express CDC for Oracle, contact Informatica Global Customer Support to discuss the best procedure for your environment. Also review the H2L article "Migrating to PowerExchange Express CDC for Oracle," which is available from the Informatica Knowledge Base at <https://kb.informatica.com>.

## Configuring Oracle for Express CDC

You must perform some configuration tasks in Oracle to prepare for PowerExchange Express CDC for Oracle.

If you have not configured the use of archived logs, enable ARCHIVELOG mode and specify a primary archive log destination. Then use the sample file, ora\_orad.sql, that PowerExchange supplies in its installation directory to perform the other tasks. To run the SQL statements in the ora\_orad.sql file, you must be assigned the role of DBA.

Complete the following Oracle configuration tasks:

1. Specify an archive log destination for PowerExchange Express CDC for Oracle use if one is not already defined.
2. Enable ARCHIVELOG mode if it is not already enabled.
3. If you enable ARCHIVELOG mode, stop and restart the Oracle database.
4. Create an Oracle user and grant user privileges.
5. If you use ASM and plan to connect to an ASM instance to read ASM-managed Oracle redo logs, create an ASM user that has SYSDBA or SYSASM authority.
6. Enable Oracle minimal global supplemental logging.

For configuration tasks specific to Oracle multitenant environments, see [“Task Flow: Configuring Change Capture from an Oracle Multitenant Pluggable Database” on page 157](#).

## Specifying an Archive Log Destination

You must specify the archive log destination from which the Express CDC log reader selects archived redo log files.

Consult with your Oracle DBA. To specify the archived log destination, you can complete one of the following actions:

- Edit the init.ora file to specify the archive log destination and file-name format. For more information about the init.ora file, see your Oracle database administrator's guide.
- Customize the appropriate server parameter file (spfile) to indicate the archived log destination, for example:

```
CONNECT SYS/sys_pwd AS SYSDBA;  
ALTER SYSTEM SET  
log_archive_dest_1 = 'location=/oracle_path/arch'  
SCOPE=SPFILE;
```

## Enabling ARCHIVELOG Mode

PowerExchange Express CDC for Oracle requires Oracle to be running in ARCHIVELOG mode.

By default, ARCHIVELOG mode is not enabled.

To enable ARCHIVELOG mode, issue the following statements:

```
SHUTDOWN IMMEDIATE;
STARTUP MOUNT;
ALTER DATABASE ARCHIVELOG;
ALTER DATABASE OPEN;
SHUTDOWN IMMEDIATE;
STARTUP;
```

**Tip:** Back up your database after each SHUTDOWN command.

## Creating an Oracle User and Granting User Privileges

To capture change data from Oracle redo logs, you must define a CDC user and assign specific Oracle system and object privileges to that user.

You can either use an existing user who has the required authority as the CDC user, or create a user and grant the required privileges to that user. In the ora\_orad.sql file, PowerExchange provides sample SQL statements for creating an Oracle user and for granting the required privileges to that user.

1. Log in to Oracle as a user who has the role of DBA.
2. To create a user named ORACAPTL1, issue the following SQL statement:

```
CREATE USER "ORACAPTL1" PROFILE "DEFAULT"
  IDENTIFIED BY "oracaptl"
  ACCOUNT UNLOCK;
COMMIT;
```

3. To retrieve the information that is required for correct capture processing, grant the following minimum object and system table privileges to the ORACAPTL1 user:

```
GRANT SELECT ON "PUBLIC"."V$ARCHIVED LOG" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$DATABASE" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$LOG" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$LOGFILE" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$TRANSPORTABLE_PLATFORM" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$THREAD" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$DATABASE_INCARNATION" TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$SPPARAMETER" TO "ORACAPTL1";

GRANT SELECT ON "SYS"."ALL TABLES" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."ALL_TAB_PARTITIONS" TO "ORACAPTL1";

GRANT SELECT ON "SYS"."DBA_LOG_GROUPS" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."DBA_LOG_GROUP_COLUMNS" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."DBA_USERS" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."DBA_TABLESPACES" TO "ORACAPTL1";

GRANT SELECT ON "SYS"."OBJ$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."TAB$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."IND$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."COL$" TO "ORACAPTL1";

GRANT SELECT ON "SYS"."PARTOBJ$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."TABPART$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."INDPART$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."TABCOMPART$" TO "ORACAPTL1";
GRANT SELECT ON "SYS"."TABSUBPART$" TO "ORACAPTL1";
COMMIT;
```

4. To extract change data in real time and run PowerExchange Logger tasks, grant the following system privilege to the ORACAPTL1 user:

```
GRANT "CONNECT" TO "ORACAPTL1";
```



5. To capture data from a pluggable database (PDB) in an Oracle multitenant container database (CDB), grant the following privileges to the ORACAPTL1 user:

```
GRANT CREATE SESSION TO "ORACAPTL1";
GRANT SELECT ON "PUBLIC"."V$PDBS" TO "ORACAPTL1";
```

6. To capture data from an Oracle Data Guard physical standby database, grant the following privilege to the ORACAPTL1 user:

```
GRANT SELECT ON "PUBLIC"."V$STANDBY_LOG" TO "ORACAPTL1";
```

7. To capture data from Oracle TDE-encrypted tablespaces, grant the following privilege to the ORACAPTL1 user:

```
GRANT SELECT ON "PUBLIC"."V$ENCRYPTION_WALLET" TO "ORACAPTL1";
```

8. To create capture registrations and perform other tasks in the PowerExchange Navigator, grant the following privileges to the user who creates and manages registrations:

```
GRANT SELECT ON "PUBLIC"."V$PARAMETER" TO "ORACAPTL1";
GRANT SELECT ON table TO "registration_user"; <<-Repeat for each table of CDC
interest.
```

Instead of granting SELECT on each table of interest, you can specify GRANT SELECT ANY TABLE for the *registration user* if your site security rules allow it.

9. To run the SQL for creating supplemental log groups at the end of registration, grant the following system privilege to the user who created the registration:

```
GRANT ALTER ANY TABLE TO "registration_user";
```

If your site security rules do not allow this level of authority to be granted to the registration user, you can give the SQL file to your DBA. The DBA can then use the SQL to create the supplemental log groups.

**Attention:** If you do not use ASM, the operating system user ID under which the PowerExchange Express CDC for Oracle capture process runs must also have the authority to read the Oracle online and archive redo logs. Otherwise, the Express CDC log reader cannot directly read the log files, and the CDC session ends with error message PWX-36140.

## Creating an ASM User (ASM only)

If you use ASM and plan to connect to an ASM instance to retrieve change data from ASM-managed Oracle redo logs, you must configure an ASM login user ID that has SYSDBA or SYSASM authority.

- If you use Oracle 10.2.0.4 or later 10.2 patchset, use an Oracle database user that has SYSDBA authority. By default, Oracle 10.2.x ASM does not allow connections from the Oracle listener on a database server. To prevent ASM connection errors, edit the listener.ora file on each node where you run PowerExchange Express CDC for Oracle to add an entry for the ASM instance. The ASM instance name is usually +ASM or +ASM $n$ , where  $n$  is a RAC node number. For example, to add the "+ASM1" instance for the Oracle instance named "MYORA1," include the following entry in the listener.ora file:

```
SID_LIST_LISTENER_MYORA1 =
(SID_LIST =
(SID_DESC =
(SID_NAME = PLSExtProc)
(ORACLE_HOME = /u01/oracle/asm)
(PROGRAM = extproc)
)
(SID_DESC =
(GLOBAL_DBNAME = MYORA1ASM)
(SID_NAME = +ASM1)
(ORACLE_HOME = /u01/oracle/asm)
)
)
```

- If you use Oracle 11.1.0.7 or later, create an ASM user and grant one of the following authority levels to that user:
  - SYSDBA  
In Oracle 11g, you do not need to edit the listener.ora file, as you do with Oracle 10.2.x ASM, because ASM allows connections from the Oracle listener by default.
  - SYSASM  
To use SYSASM, you must also set the ASM\_ASSYSASM parameter to Y in the READER statement of the PowerExchange Express CDC for Oracle configuration file, pwxorad.

For more information about creating an ASM user, see the *Oracle Automatic Storage Management Administrator's Guide*.

## Enabling Oracle Minimal Global Supplemental Logging

PowerExchange Express CDC for Oracle requires Oracle minimal global supplemental logging to properly handle chained rows.

To enable supplemental logging, log in to the Oracle database and issue the following SQL statement, which is included in the sample ora\_orad.sql script file:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
COMMIT;
```

If you do not know whether minimal global supplemental logging is enabled for your database, you can still execute this ALTER statement. This statement has no effect when minimal supplemental logging is already enabled.

If you execute this statement while the Oracle database is open, Oracle waits for in-flight transactions to complete, which can affect database performance. This issue is likely to occur for databases that have a high level of user activity. To avoid this problem, you can close and re-open the database and then issue the statement manually.

**Note:** You must also define a supplemental log group for each Oracle source table. When you register an Oracle table in the PowerExchange Navigator, PowerExchange generates DDL for adding a supplemental log group for the table. Supplemental log groups cause Oracle to log full before- and after-images of the data that changed. PowerExchange requires these images.

## Configuring PowerExchange for Express CDC

You must complete several tasks to configure PowerExchange for Express CDC for Oracle and to start change capture.

The specific tasks depend on your configuration. The following task flow assumes that you use the PowerExchange Logger, as recommended.

**Important:** Make sure that the system user ID under which the PowerExchange Express CDC for Oracle capture process runs has the authority to read the Oracle redo logs. Otherwise, the Express CDC log reader cannot directly read the log files, and the CDC session ends with error message PWX-36140.

1. Configure the dbmover configuration file on the PowerExchange system that initiates the CAPI to capture change data from the Oracle source database.

Include the following statements:

- Required. CAPT\_PATH and CAPT\_XTRA statements.
- Required. An ORAD CAPI\_CONNECTION statement.
- Required. An ORACLEID statement.
- Recommended. An ORACLE\_CAPTURE\_TYPE statement.
- Recommended. A CAPX CAPI\_CONNECTION statement, if you use the PowerExchange Logger and continuous extraction mode.
- Optional. An ORACLE\_UNHANDLED\_NUMASCHAR statement to treat NUMBER columns that have a precision greater than 28 or an undefined length as variable length strings and to treat FLOAT columns that have a precision greater than 15 as variable length strings. This behavior can prevent data loss. You must define this statement on the PowerExchange instance where the registration and extraction map information will be defined, before you create the registrations.

**Notes:** If you run PowerExchange Express CDC and the PowerExchange Listener, consumer API (CAPI), and PowerExchange Logger on the Oracle system, as in Configuration 1, define the CAPX CAPI\_CONNECTION and ORACLE\_UNHANDLED\_NUMASCHAR statements in this dbmover configuration file.

If you capture change data from an Oracle Data Guard physical standby database, configure the ORACLEID and ORAD CAPI\_CONNECTION in the DBMOVER configuration file on the standby system to point to the source instance on the primary system. If the primary system is a RAC environment, point to a single member in the RAC.

2. Customize the PowerExchange Express CDC for Oracle configuration file, which has the default file name of pwxorad.cfg, on the change capture system.

If you capture changes from RAC members, include the RAC statement.

If you use ASM, include the ASM parameters in the READER statement.

**Tip:** To include rowid values in change records for tables that do not have Oracle row movement enabled, include the OPTIONS ROWID=Y statement. Provided that you use version 9.1.2 or later of the PowerExchange Navigator, PowerExchange can write rowid values to the generated DTL\_\_CAPXROWID column in captured change records.

3. Customize the PowerExchange Logger configuration file, which has the default file name of pwxcl.cfg, on the system where the PowerExchange Logger runs.

For more information, see ["Customizing the PowerExchange Logger Configuration File" on page 44](#).

4. Configure a dbmover configuration file on any other PowerExchange instance in the CDC environment that needs to read capture registrations or extraction maps from disk.

If you use Configuration 2, include the following statements in the dbmover configuration file on the system with the separate PowerExchange Listener:

- Required. An ORACLE\_CAPTURE\_TYPE statement.
- Required. An ORACLEID statement.
- Recommended. A CAPX CAPI\_CONNECTION statement if you run the PowerExchange Logger on the change capture system and use continuous extraction mode.
- Optional. An ORACLE\_UNHANDLED\_NUMASCHAR statement if you need PowerExchange to treat NUMBER columns that have a precision greater than 28 or an undefined length as variable length strings and to treat FLOAT columns that have a precision greater than 15 as variable length strings.

5. In the PowerExchange Navigator, create a registration group and a capture registration for each Oracle source table.

**Note:** If you are migrating from PowerExchange Oracle CDC with LogMiner to PowerExchange Express CDC for Oracle, you can use the same registrations.

Make sure that you enter values in the following fields:

- In the **Type** list, select **ORACLE**.
- In the **Collection ID** box, enter a user-defined name for the Oracle instance. This value must match the collection ID in the ORACLEID statement in the dbmover.cfg file that is on the node specified in the **Location** field. If you use the PowerExchange Logger, this value also must match the DBID value in the PowerExchange Logger configuration file.
- Set the **Status** option to **Active**.
- In the **Condense** list, select **Part**. Informatica recommends this setting even if you do not use the PowerExchange Logger, because it allows you to implement the PowerExchange Logger later, if necessary, without changing the registration.
- In the **Supplement Log Group Name** box, enter a name for the Oracle supplemental log group.

**Tip:** If you want PowerExchange to run the DDL that it generates to create the supplemental log group at registration completion, select **Execute DDL Now**. However, Informatica recommends that you save the generated DDL to a file and give it to your Oracle DBA instead. The DBA can then use this DDL when migrating PowerExchange from a test or QA environment to the production environment.

After you click **Finish**, PowerExchange generates corresponding extraction maps.

For more information about registrations and extraction maps, see the *PowerExchange Navigator User Guide*.

6. Perform a database row test on each extraction map.  
For more information, see the *PowerExchange Navigator User Guide*.
7. Configure a restart point.  
For more information, see [“Creating Restart Tokens for Extractions” on page 241](#).
8. Materialize the target.
9. Start the PowerExchange Logger.

Next, configure PowerCenter CDC sessions. You can use batch extraction mode, continuous extraction mode, or real-time extraction mode. For more information, see *PowerExchange Interfaces for PowerCenter*.

## RELATED TOPICS:

- [“Configuring the dbmover Configuration File” on page 164](#)
- [“Customizing the PowerExchange Express CDC for Oracle Configuration File” on page 173](#)

# Configuring the dbmover Configuration File

In the dbmover configuration file, define the statements that are required for PowerExchange Express CDC for Oracle. Include the optional statements as needed.

Define the following statements in the dbmover configuration file on the system where the capture registrations and CDC control files are stored, which is specified as the **Location** node in the registration group:

### CAPT\_PATH

Required. Path to the local directory on the system that contains the control files for CDC. These files are: the CCT file for capture registrations, the CDEP file for application names that are used for ODBC extractions, and the CDCT file for the PowerExchange Logger for Linux, UNIX, and Windows.

### CAPT\_XTRA

Required. Path to the local directory on a system that stores extraction maps for CDC.

## ORACLE\_CAPTURE\_TYPE

Recommended. Indicates whether PowerExchange Express CDC for Oracle or PowerExchange Oracle CDC with LogMiner is in use for the PowerExchange installation. This setting must be consistent with the CAPI\_CONNECTION type.

If multiple systems are involved in CDC processing, for example, because you run a separate PowerExchange Listener to handle registrations and extraction maps, you must also define the ORACLE\_CAPTURE\_TYPE statement in the dbmover.cfg file on each system.

## ORACLEID

Required. The Oracle source instance, database, and connection information that is used for CDC.

If you capture change data from a pluggable database (PDB) in an Oracle multitenant environment, the second positional parameter, *oracle\_db*, specifies the name of the database that contains the PDB and the fourth positional parameter, *capture\_connect\_string*, specifies the name of the PDB service entry in the tnsnames.ora file.

## ORACLE\_UNHANDLED\_NUMASCHAR

Optional. Controls whether capture registration processing for PowerExchange Express CDC or PowerExchange Oracle CDC with LogMiner sources handles NUMBER columns that have a precision greater than 28 or an undefined length as variable length strings and handles FLOAT columns that have a precision greater than 15 as variable length strings.

## ORAD CAPI\_CONNECTION

Required. A named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and control PowerExchange Express CDC for Oracle processing for Oracle data sources. In this statement, you can include the PARMFILE parameter to point to the separate configuration file that contains statements and parameters specifically for PowerExchange Express CDC for Oracle. Include the PARMFILE parameter if you want to override the default file name of pwxorad.cfg.

**Note:** Do not also specify ORCL CAPI\_CONNECTION statements in the same dbmover configuration file.

For more information about the ORACLE\_CAPTURE\_TYPE, ORACLE\_UNHANDLED\_NUMASCHAR, and ORAD CAPI\_CONNECTION statements, see the following detailed descriptions. For more information about all other DBMOVER statements, see the *PowerExchange Reference Manual*.

## RELATED TOPICS:

- [“CAPI\\_CONNECTION - ORAD Statement” on page 168](#)
- [“CAPI\\_CONNECTION - CAPX Statement” on page 57](#)
- [“ORACLEID Statement” on page 123](#)
- [“ORACLE\\_UNHANDLED\\_NUMASCHAR Statement” on page 170](#)
- [“ORACLE\\_CAPTURE\\_TYPE Statement” on page 168](#)
- [“CAPT\\_PATH Statement” on page 30](#)
- [“CAPT\\_XTRA Statement” on page 30](#)

## ORACLEID Statement

The ORACLEID statement specifies the Oracle source instance, database, and connection information for CDC.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Oracle CDC sources

**Related Statements:** CAPI\_CONNECTION - ORCL and CAPI\_CONNECTION - ORAD

**Required:** Yes, for PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle

**Syntax:**

```
ORACLEID=(collection_id
          ,oracle_db
          [,source_connect_string]
          [,capture_connect_string]
          [,fifth_positional_parameter]
          [,USEDATABASE])
```

**Parameters:**

***collection\_id***

Required. User-defined identifier for this ORACLEID statement. This value must match the ORACOLL parameter value in the ORCL CAPI\_CONNECTION or ORAD CAPI\_CONNECTION statement, the collection ID in the registration group defined for the source tables, and the DBID value in the PowerExchange Logger pwxcl configuration file.

Maximum length is eight characters.

***oracle\_db***

Required. Name of the Oracle database that contains the source tables that are registered for change data capture. If you use PowerExchange Express CDC for Oracle to capture change data from a pluggable database (PDB) in an Oracle multitenant environment, this value is the name of the database that contains the PDB.

***source\_connect\_string***

Optional. Oracle connection string, defined in TNS, that is used to connect to the Oracle database that contains the source tables. This connection string must be defined in the Oracle Client tnsnames.ora file on the system with the source database.

For PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle, the source connection string is used only for PowerExchange Navigator access to the Oracle source database. Enter this parameter in the dbmover configuration file on the machine from which the PowerExchange Listener retrieves data for PowerExchange Navigator requests. If you plan to run a database row test on extraction maps for the source tables, also specify the *capture\_connect\_string* parameter.

**Note:** The source connection string is not used to transfer change data.

If this value is null and the Oracle source is *not* a PDB in a multitenant environment, the value of the ORACLE\_SID environment variable is used by default. If you use PowerExchange Express CDC for Oracle to capture change data from a PDBs, you must enter a value in this parameter.

***capture\_connect\_string***

Optional. Oracle connection string, defined in TNS, that the PowerExchange Logger uses to connect to the Oracle database with the source tables for PowerExchange Oracle CDC with LogMiner or PowerExchange Express CDC with LogMiner. This connection string must be specified in the Oracle Client tnsnames.ora file that is used for connection to the Oracle source database. If you use PowerExchange Express CDC to capture change data from a PDB in an Oracle multitenant environment, specify the name of the PDB service entry in the tnsnames.ora file.

If this value is null and the Oracle source is *not* a PDB in a multitenant environment, the value of the ORACLE\_SID environment variable is used by default.

If this value is null and the Oracle source is a PDB, PowerExchange cannot capture change data for the source. If you use PowerExchange Express CDC for Oracle to capture change data from PDBs, you must enter a value in this parameter.

Also, for PowerExchange Oracle CDC with LogMiner or Express CDC for LogMiner, if you have multiple Oracle databases and capture changes from a database other than the default database, you must specify both the *source\_connect\_string* and *capture\_connect\_string* parameters.

**Tip:** If possible, bypass the use of SQL\*Net to improve PowerExchange Logger performance, even if the PowerExchange Logger is running on the same machine as the Oracle source database. Set the following environment variables, whenever possible, to enable connection to the appropriate Oracle database without using the *capture\_connect\_string* parameter and SQL\*Net:

- ORACLE\_HOME
- ORACLE\_SID
- PATH
- On Linux or UNIX, one of the following:
  - LD\_LIBRARY\_PATH
  - LIBPATH
  - SHLIB\_PATH

#### ***fifth\_positional\_parameter***

Not used. Add a comma as a placeholder if you specify the USEDDBNAME positional parameter, for example:

```
ORACLEID=(collection_id,oracle_db,src_connect_string,capture_connect_string,,USEDDBNAME)
```

#### **USEDDBNAME**

Optional. Specify this parameter only under all of the following conditions:

- You upgrade to PowerExchange 9.1.0 or later from an earlier release.
- You use Oracle 11g or later.
- You run the following SQL query on the V\$DATABASE view and the query returns different values for the NAME and DB\_UNIQUE\_DATABASE fields, including values that vary in case only such as ORAABC1 and oraabc1:

```
select name, db_unique_name from v$database;
```

In this situation, the USEDDBNAME parameter can prevent potential restart errors that are caused by the difference in the NAME and DB\_UNIQUE\_DATABASE values.

**Tip:** Alternatively, you can specify the DB\_UNIQUE\_NAME value in the second positional parameter, *oracle\_db*.

#### **Usage Notes:**

- PowerExchange requires an ORACLEID statement for each Oracle database for which you want to capture and extract change data. You can define a maximum of 20 ORACLEID statements in a single dbmover configuration file.
- Define the ORACLEID statement in the dbmover configuration file on the system where the PowerExchange Logger runs, or if you plan to perform Oracle CDC without the PowerExchange Logger, on the system where your PowerExchange extractions run.

## ORACLE\_CAPTURE\_TYPE Statement

The ORACLE\_CAPTURE\_TYPE statement specifies which Oracle CDC solution is in use for a PowerExchange installation: PowerExchange Express CDC for Oracle or PowerExchange Oracle CDC with LogMiner.

The ORACLE\_CAPTURE\_TYPE value must be consistent with the type of Oracle CAPI\_CONNECTION statement that is defined on the system that initiates the connection to the Oracle system for change capture.

To ensure consistent behavior, define the ORACLE\_CAPTURE\_TYPE statement on all systems that are involved in Oracle CDC processing, including the system where the CAPI\_CONNECTION statements are defined.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Oracle

**Required:** No

**Syntax:**

```
ORACLE_CAPTURE_TYPE={D|L}
```

**Valid Values:**

- **D.** PowerExchange uses Express CDC for Oracle and ORAD CAPI\_CONNECTION statements.
- **L.** PowerExchange uses Oracle CDC with LogMiner and ORCL CAPI\_CONNECTION statements.

No default value is available.

**Usage Notes:**

- PowerExchange Express CDC for Oracle handles all character columns as variable length columns, whereas PowerExchange Oracle CDC with LogMiner handles character columns as fixed length or variable length columns. This behavioral difference affects PowerExchange column-level processing and the view of the CDC extraction map that is imported into PowerCenter. As a result, all systems that are involved in Oracle change capture processing must be aware of the Oracle CDC type that is in use.

On the PowerExchange system where the ORAD or ORCL CAPI\_CONNECTION statements are defined, the CAPI\_CONNECTION type implicitly defines the CDC type, and the ORACLE\_CAPTURE\_TYPE statement is optional. However, if other systems are involved in CDC processing, for example, because you run a separate PowerExchange Listener or use offload processing, you must define the ORACLE\_CAPTURE\_TYPE statement in the dbmover.cfg file on each system to explicitly define the CDC type.

- You can use only one type of Oracle CAPI\_CONNECTION in a dbmover.cfg file. The ORACLE\_CAPTURE\_TYPE value must be consistent with this CAPI\_CONNECTION type. Otherwise, PowerExchange issues an error message and ends abnormally. The following settings are consistent:
  - If you use ORAD CAPI\_CONNECTION statements, set ORACLE\_CAPTURE\_TYPE to D.
  - If you use ORCL CAPI\_CONNECTION statements, set ORACLE\_CAPTURE\_TYPE to L.

## CAPI\_CONNECTION - ORAD Statement

The ORAD CAPI\_CONNECTION statement specifies a named set of parameters that the Consumer API (CAPI) uses to connect to the change stream and control PowerExchange Express CDC for Oracle processing for Oracle data sources.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Oracle

**Related Statements:** ORACLEID, ORACLE\_CAPTURE\_TYPE



**Required:** Yes, for PowerExchange Express CDC for Oracle

**Syntax:**

```
CAPI_CONNECTION= ([DLLTRACE=trace_id
                  ,NAME=capi_connection_name
                  ,TYPE=(ORAD
                        [,EPWD=database_encrypted_password]
                        ,ORACOLL=collection_id
                        [,PARMFILE=express_cdc_configuration_file]
                        [,PASSWORD=database_password]
                        [,USERID=database_user_id]
                  )
                )
```

**Parameters:**

**DLLTRACE=*trace\_id***

Optional. User-defined name of the TRACE statement that activates internal DLL tracing for this CAPI.

Specify this parameter only at the direction of Informatica Global Customer Support.

**NAME=*capi\_connection\_name***

Required. Unique user-defined name for this CAPI\_CONNECTION statement.

Maximum length is eight alphanumeric characters.

**TYPE=(ORAD, ... )**

Required. Type of CAPI\_CONNECTION statement. For PowerExchange Express CDC for Oracle sources, this value must be ORAD.

**EPWD=*database\_encrypted\_password***

Optional. An encrypted password that PowerExchange uses to connect to the Oracle source database for PowerExchange Express CDC for Oracle. If you specify the USERID parameter in this statement, you must also specify either the EPWD or PASSWORD parameter.

This encrypted password overrides the EPWD parameter value in the DATABASE statement of the PowerExchange Express CDC for Oracle configuration file and the **Password** value in the CDC session connection attributes. To use one of these other encrypted passwords, do not include the EPWD parameter in the ORAD CAPI\_CONNECTION statement.

**ORACOLL=*collection\_id***

Required. The collection identifier for the Oracle instance. This value must match the collection ID in the first positional parameter of an ORACLEID statement in the same dbmover.cfg file.

Usually, this value also matches the collection ID that you specify in the registration group for the Oracle instance. If you specify a different collection ID in the registration group, the registration collection ID overrides this ORACOLL value.

**PARMFILE=*path\_and\_filename***

Optional. The path and file name for the PowerExchange Express CDC for Oracle configuration file, relative to the current working directory. You can use this parameter to override the default path and file name or to remind PowerExchange users of the default path and file name.

If this parameter is not specified, PowerExchange uses *pwd\_home\_directory\pwxorad.cfg* by default. The default path is the path in the PWD\_HOME environment variable, or if this environment variable is not defined, the default path is the path to the PowerExchange bin directory. If the pwxorad.cfg file does not exist at the default location and a PARMFILE override is not defined, PowerExchange issues error messages PWX-09951 and PWX-00268 and change capture fails.

**PASSWORD=database\_password**

Optional. A clear text password that PowerExchange uses to connect to the Oracle source database for PowerExchange Express CDC for Oracle. If you specify the USERID parameter in this statement, you must specify either the EPWD or PASSWORD parameter.

This clear-text password overrides the PASSWORD parameter in the DATABASE statement of the PowerExchange Express CDC for Oracle configuration file and the **Password** value in the CDC session connection attributes. To use one of these other passwords, do not include the PASSWORD parameter in the ORAD CAPI\_CONNECTION statement.

**USERID=database\_user\_id**

Optional. A user ID that PowerExchange uses to connect to the Oracle source database for PowerExchange Express CDC for Oracle. If you specify the USERID parameter in this statement, you must also specify either the EPWD or PASSWORD parameter.

This user ID overrides the USERID parameter in the DATABASE statement of the PowerExchange Express CDC for Oracle configuration file and the **User Name** value in the CDC session connection attributes. To use one of these other user IDs, do not include the USERID parameter in the ORAD CAPI\_CONNECTION statement.

**Usage Notes:**

- You can specify multiple ORAD CAPI\_CONNECTION statements in the dbmover.cfg file to capture change data from more than one Oracle instance or to use different parameter settings for the same Oracle instance.
- Define the ORAD CAPI\_CONNECTION and ORACLEID statements on the PowerExchange system that must connect to the Oracle source database for change data capture. Usually, the PowerExchange Logger for Linux, UNIX, and Windows runs on this system.
- You cannot define both ORAD CAPI\_CONNECTION and ORCL CAPI\_CONNECTION statements in the same dbmover.cfg file. If you use PowerExchange Express CDC for Oracle, define ORAD CAPI\_CONNECTION statements. If you use PowerExchange Oracle CDC with LogMiner, define ORCL CAPI\_CONNECTION statements. If you need to run both PowerExchange Express CDC for Oracle and PowerExchange Oracle CDC with LogMiner sessions against the same Oracle instance, use separate dbmover.cfg files, PowerExchange Listeners, and PowerExchange Loggers.
- The database user ID and password or encrypted password can be specified in multiple locations. If you do so, PowerExchange uses the following order of precedence:
  1. The USERID value and the EPWD or PASSWORD value that are specified in the ORAD CAPI\_CONNECTION statement in the dbmover configuration file
  2. The USERID value and EPWD or PASSWORD value that are specified in the DATABASE statement in the PowerExchange Express CDC for Oracle configuration file, pwxorad.cfg.
  3. The **User Name** and **Password** values that are specified in the Oracle application connection attributes for the PowerCenter CDC session

**ORACLE\_UNHANDLED\_NUMASCHAR Statement**

The ORACLE\_UNHANDLED\_NUMASCHAR statement controls how PowerExchange handles some numeric Oracle source columns.

If you enter Y, PowerExchange converts the following Oracle numeric datatypes:

- NUMBER columns that have a precision greater than 28 or an undefined length are treated as variable-length strings instead of double-precision floating-point numbers.

- FLOAT columns that have a precision greater than 15 significant digits are treated as variable-length strings.

PowerExchange uses the ORACLE\_UNHANDLED\_NUMASCHAR setting when creating capture registrations.

This statement applies to PowerExchange Express CDC for Oracle and PowerExchange Oracle CDC with LogMiner sources. You can use this statement to override PowerExchange default processing of numeric data to prevent data loss in certain circumstances. To override default processing, you must specify this statement prior to creating capture registrations.

**Operating Systems:** Linux, UNIX, and Windows

**Data Sources:** Oracle

**Required:** No

**Syntax:**

```
ORACLE_UNHANDLED_NUMASCHAR={ Y | N }
```

**Valid Values:**

- **N.** PowerExchange uses its default processing of Oracle NUMBER data. If you have NUMBER columns that have a precision greater than 28 or an undefined length, or if you have FLOAT columns with a precision greater than 15, change data loss might occur.
- **Y.** PowerExchange handles NUMBER and FLOAT data in a manner that prevents data loss.

Default is N.

**Usage Notes:**

- Enter this parameter with a value of Y before you create capture registrations for the Oracle source tables that contain the NUMBER or FLOAT columns. If you enter this parameter after the registrations exist, you must set the status of the registrations to History and then create the registrations again. Otherwise, change data loss might occur.
- Oracle allows columns that have the NUMBER datatype to have their precision and scale determined by the numeric data that is written to the columns. Oracle supports a maximum precision of 38 and an exponent of +/-127.

If you do not explicitly define the precision and scale for NUMBER columns from which change data is captured, the following default PowerExchange and PowerCenter processing of change data can result in loss of precision and change data:

- PowerExchange handles data in NUMBER columns that have an undefined length or a length greater than 100 bytes as double-precision floating-point numbers.
- PowerCenter allows a maximum precision of 28 for decimal numbers.

To prevent change data loss with this type of data, enter Y for this statement and then create your capture registrations. PowerExchange registration processing can then handle numbers that have a precision greater than 28 as variable-length strings.

If you are writing the data to an Oracle target and want to maintain the precision as a variable-length string, edit the target definition to modify the column datatype. Within a PowerCenter mapping, you can convert a variable-length string to a number either implicitly by connecting to a numeric port or explicitly by using expressions. To avoid loss of precision in implicit conversions, you might need to edit the mapping to pass the data as a string from source to target.

- PowerExchange supports the BINARY\_DOUBLE and BINARY\_FLOAT numeric datatypes by treating them as internal DOUBLE or FLOAT datatypes. However, PowerCenter converts BINARY\_DOUBLE and BINARY\_FLOAT datatypes to Oracle NUMBER(15) datatypes, which can result in arithmetic overflow and data loss.

- In columns with an Oracle numeric datatype, PowerExchange treats the value of infinity as 0.

## Example dbmover Configuration File for the Oracle Change Capture System

This example dbmover configuration file contains the basic statements that are required on the Oracle system where PowerExchange Express CDC for Oracle initiates CAPI connections to Oracle for change capture.

The following dbmover configuration file corresponds to the example configuration, in which the Express CDC capture process and PowerExchange Logger run on the same system as the Oracle database:

```

LISTENER=(pwxlst1,TCPIP,2480)
NODE=(local,TCPIP,192.168.6.220,2480)
NODE=(pwxnode1,TCPIP,192.168.6.220,2480)
NODE=(oranode1,TCPIP,192.168.6.220,2480)
APPBUFFSIZE=256000
COLON=:
COMPRESS=N
CONSOLE_TRACE=N
DECPOINT=.
DEFAULTCHAR=*
DEFAULTDATE=19800101
MAXTASKS=60
MSGPREFIX=PWX
NEGSIGN=-
PIPE=|
POLLTIME=1000
TIMEOUTS=(300,600,600)
CAPT_PATH=/Informatica/PowerExchangeVR/capture
CAPT_XTRA=/Informatica/PowerExchangeVR/capture/extmaps
LOGPATH=/Informatica/PowerExchangeVR/capture/logs
CODEPAGE=(utf-8,utf-8,utf-8)
/*
/* Define the ORACLE_CAPTURE_TYPE statement to explicitly define
/* the CDC type.
ORACLE_CAPTURE_TYPE=D
/*
/* Define an ORACLEID statement for each Oracle instance involved
/* in CDC.
ORACLEID=(ORAD1DB,ORAD1,ORAD1DB,ORAD1DB)
/*
/* An ORAD CAPI_CONNECTION statement is required for Oracle
/* Express CDC. Do not also specify ORCL CAPI_CONNECTION
/* statements in this file.
CAPI_CONNECTION=(NAME=CAPORA3,TYPE=(ORAD,ORACOLL=ORAD1DB,
PARMFILE=/Informatica/PowerExchangeVR/capture/pwxorad.cfg))
/*
/* A CAPX CAPI_CONNECTION statement is required if you use
/* the PowerExchange Logger and continuous extraction mode.
CAPI_CONNECTION=(NAME=oralcpx,TYPE=(CAPX,DFLTINST=ORAD1DB))

```

The ORAD CAPI\_CONNECTION is required on the PowerExchange system that initiates the CAPI connection to Oracle. Include the PARMFILE parameter to point to the PowerExchange Express CDC for Oracle configuration file when you do not use the default file name or location.

The ORACLE\_CAPTURE\_TYPE statement is not required on this system because an ORAD CAPI\_CONNECTION statement is present. However, Informatica recommends that you include ORACLE\_CAPTURE\_TYPE for consistency. The ORACLE\_CAPTURE\_TYPE setting must be **D** to be consistent with the CAPI\_CONNECTION type of ORAD.

## RELATED TOPICS:

- [“PowerExchange Express CDC for Oracle Architecture” on page 141](#)

# Customizing the PowerExchange Express CDC for Oracle Configuration File

You specify PowerExchange Express CDC for Oracle parameters in a configuration file that is separate from the dbmover.cfg file.

PowerExchange supplies the sample pwxorad.cfg file in the directory that is specified in the PWX\_HOME environment variable, or if that variable is not defined, in the PowerExchange bin directory. The sample file contains comments that describe the required and optional statements and keywords.

Copy the sample file and customize the copy. If you copy the file under another name or to another directory, you must include the PARMFILE parameter in the ORAD CAPI\_CONNECTION statement in the dbmover.cfg file to point to the customized copy.

The following table summarizes the statements that you can define in the PowerExchange Express CDC for Oracle configuration file:

Statement	Required or Optional	Description
DATABASE	Optional	Provides overrides for database connection information, including the capture connection string and database user ID and password. Also includes the TDE encryption password, which is required if you capture changes from TDE-encrypted tablespaces.
DICTIONARY	Required	Indicates where PowerExchange gets the Oracle data dictionary information that it uses to interpret the redo logs and how PowerExchange reacts to DDL changes that might occur after the dictionary is read into memory.
DIRSUB	Optional	Specifies a path prefix that PowerExchange Express CDC for Oracle substitutes for the original path prefix that the Oracle server uses to access active and archived redo logs. This substitute path is required when the PowerExchange Express CDC log reader runs on a system other than the Oracle server and uses a different path to access the redo log files.
OPTIONS	Optional	Provides options for CDC processing, including the number of log records staged in memory, maximum spill file size, and capture of Oracle rowid values.
RAC	Optional	Specifies the number of member instances in an Oracle RAC from which you capture changes.
READER	Required	Provides options for reading the redo logs.
STANDBY	Optional	Defines a connection to an Oracle physical standby database when the database is not open for read-only access. This statement applies only when the CDC source is an Oracle physical standby database.

In each statement that you include, you must specify at least one valid keyword. The end of an entire statement is indicated by a semicolon (;).

## RELATED TOPICS:

- [“Example PowerExchange Express CDC for Oracle Configuration File” on page 184](#)
- [“DATABASE Statement” on page 174](#)
- [“DICTIONARY Statement” on page 176](#)
- [“DIRSUB Statement” on page 177](#)
- [“OPTIONS Statement” on page 177](#)
- [“RAC Statement” on page 180](#)
- [“READER Statement” on page 181](#)
- [“STANDBY Statement” on page 183](#)

## DATABASE Statement

The DATABASE statement specifies overrides for the capture connection string and database user ID and password.

The DATABASE statement is optional. If you specify it, include at least one of its optional parameters.

If you want to capture change data from an Oracle Data Guard physical standby database that is open for read-only access, configure the DATABASE statement to connect to the standby database instance. If the standby database is *not* open for read-only access, configure the DATABASE statement to connect to the primary database instance and also define the STANDBY statement.

### Syntax:

```
DATABASE
  [CONNECT_STRING=capture_connect_string]
  [EPWD=database_encrypted_password|PASSWORD=database_password]
  [TDEWALLETDIR=TDE_wallet_path]
  [TDEWALLETEPWD=encrypted_password_for_TDE_wallet]
  [TDEWALLETPWD=cleartext_password_for_TDE_wallet]
  [USERID=database_user_id]
;
```

### Parameters:

#### CONNECT\_STRING

An Oracle connection string, defined in TNS, that the PowerExchange Express CDC for Oracle uses to connect to the Oracle database.

This connection string is overridden by the **Connect String** value, if specified, in the CDC session connection attributes.

#### EPWD

An encrypted password that PowerExchange uses to connect to the Oracle source database for PowerExchange Express CDC for Oracle. If you specify the USERID parameter in this statement, you must also specify either the EPWD or PASSWORD parameter.

This encrypted password overrides the **Password** value, if specified, in the CDC session connection attributes.

This encrypted password is overridden by the EPWD parameter value in the ORAD CAPI\_CONNECTION statement, if specified.

#### PASSWORD

A clear text password that PowerExchange uses to connect to the Oracle source database for PowerExchange Express CDC for Oracle. If you specify the USERID parameter in this statement, you must also specify either the EPWD or PASSWORD parameter.

This password overrides the **Password** value, if specified, in the CDC session connection attributes.

This password is overridden by the PASSWORD parameter value in the ORAD CAPI\_CONNECTION statement, if specified.

#### **TDEWALLETDIR**

The fully qualified path and file name for the Oracle wallet file that is used for Oracle Transparent Data Encryption (TDE). Specify this parameter only if you capture change data from TDE-encrypted tablespaces and either the Oracle wallet is not available to the database or the database is running on a server that is remote from Oracle redo logs from which Express CDC reads changes.

#### **TDEWALLETPWD**

An encrypted password that PowerExchange requires to access the Oracle TDE wallet and get the master key that is required for reading and decrypting data from Oracle TDE-encrypted tablespaces. If you capture change data from TDE-encrypted tablespaces, you must specify either this parameter or the TDEWALLETPWD parameter. Do not specify both parameters.

**Note:** If you need to change the encryption password, first stop the PowerExchange Logger for Linux, UNIX, and Windows and the CDC session. Then edit the password, restart the PowerExchange Logger, and restart the CDC session.

#### **TDEWALLETPWD**

A clear text password that PowerExchange requires to access the Oracle TDE wallet and get the master key that is required for reading and decrypting data from Oracle TDE-encrypted tablespaces. If you capture change data from TDE-encrypted tablespaces, you must specify either this parameter or the TDEWALLETPWD parameter. Do not specify both parameters.

**Note:** If you need to change this encryption password, first stop the PowerExchange Logger for Linux, UNIX, and Windows and the CDC session. Then edit the password, restart the PowerExchange Logger, and restart the CDC session.

#### **USERID**

A user ID that PowerExchange uses to connect to the Oracle source database for PowerExchange Express CDC for Oracle. If you specify the USERID parameter in this statement, you must also specify either the EPWD or PASSWORD parameter.

This user ID overrides the **User Name** value, if specified, in the CDC session connection attributes.

This user ID is overridden by the USERID parameter value in the ORAD CAPI\_CONNECTION statement, if specified.

**Usage Notes:** You can specify the connection string in multiple locations. If you do so, PowerExchange uses the following order of precedence:

1. The **Connect String** value, if specified, in the CDC session connection attributes
2. The fourth positional parameter of the ORACLEID statement in the dbmover.cfg file
3. The CONNECT\_STRING parameter value in the DATABASE statement of the PowerExchange Express CDC for Oracle configuration file
4. The ORACLE\_SID environment variable

If you specify none of these values, PowerExchange passes Null values in the OCI call.

## DICTIONARY Statement

The **DICTIONARY** statement indicates where PowerExchange Express CDC for Oracle gets the data dictionary information that it uses to interpret redo logs. This statement also controls how PowerExchange Express CDC for Oracle reacts when it encounters DDL changes in the redo log records.

The **DICTIONARY** statement is required.

### Syntax:

```
DICTIONARY
  MODE=STATIC
  SOURCE=ONLINE
  [EXCEPTIONS={FAIL|WARN}]
;
```

### Parameters:

#### MODE

Required. Indicates whether PowerExchange Express CDC for Oracle expects the data dictionary to remain the same or change after reading it into memory. The only valid value is **STATIC**. In **STATIC** mode, PowerExchange expects the data dictionary to remain the same. If structural changes to source tables occur, the **EXCEPTIONS** parameter determines whether PowerExchange fails or issues a warning. However, PowerExchange Express CDC always tolerates the following DDL changes and continues CDC processing, regardless of the **EXCEPTIONS** setting:

- **ALTER TABLE ADD** statements for adding one or more columns to a table
- **ALTER TABLE ADD PARTITION** statements
- **ALTER TABLE ADD CONSTRAINT** statements
- **CREATE USER** statements
- **ALTER USER** statements
- **DROP USER** statements

#### SOURCE

Required. Indicates where PowerExchange Express CDC for Oracle gets its data dictionary information. The only valid value is **ONLINE**, which indicates that PowerExchange Express CDC for Oracle gets the data dictionary information from the Oracle online system when the change capture process initializes.

#### EXCEPTIONS

Optional. When **MODE=STATIC**, controls whether PowerExchange Express CDC for Oracle fails or continues with a warning message when it detects that a structural change was made to an Oracle table from which change data is captured.

This parameter does not apply to **ADD PARTITION** changes. PowerExchange Express CDC for Oracle tolerates **ADD PARTITION** changes.

Valid values are:

- **FAIL**. CDC processing ends abnormally with error message **PWX-36355**.
- **WARN**. PowerExchange prints the message **PWX-36354**, and CDC processing continues. This setting can result in loss of change data. Use this option only at the direction of Informatica Global Customer Support.

Default is **FAIL**.



## DIRSUB Statement

The DIRSUB statement specifies a path prefix that PowerExchange Express CDC substitutes for the original path prefix that the Oracle server uses to access active and archived redo logs.

This substitute path is required when the PowerExchange Express CDC log reader runs on a system other than the Oracle server and uses a different mapping to access the redo log files. The redo logs might reside on shared disk, or they might have been copied to the system where Express CDC runs.

Use this statement only under the following conditions:

- You do not use Oracle Automatic Storage Management (ASM) to manage the redo logs that PowerExchange Express CDC needs to access.
- PowerExchange Express CDC runs on a system that has the same operating system and architecture as the Oracle server.

You can specify multiple DIRSUB statements. Each one must have a unique value for the original path prefix that Oracle server uses. PowerExchange compares the original path prefix against redo log directories and file names, starting with the longest string for a file. For the matching log files, the PowerExchange Express CDC log reader uses the corresponding substitute path prefix to access the redo logs.

### Syntax:

```
DIRSUB SERVER="original_path_prefix",LOCAL="substitute_path_prefix";
```

If you want to include a trailing backslash in a Windows path, you must enter a double backslash (\\).

### Parameters:

#### SERVER

Required. The original path prefix for the redo logs that the Oracle server uses. This value must be unique in each DIRSUB statement.

#### LOCAL

Required. The substitute path prefix that the PowerExchange Express CDC log reader uses to access redo logs. This value does not have to be unique across multiple DIRSUB statements.

**Example:** The Oracle server and PowerExchange Express CDC run on separate Linux systems. The redo logs reside on shared disk. In this case, the PowerExchange Express CDC configuration file contains the following DIRSUB statement:

```
DIRSUB SERVER="/ora01/oraarchlogs/ORAB11",LOCAL="/oracle/oralogs/orab";
```

## OPTIONS Statement

The OPTIONS statement specifies parameters for controlling CDC processing. The parameters control memory use, spill file size, capture of rowid values, long outstanding UOWs, and the wait period for advancing restart tokens when no change capture activity is occurring.

The OPTIONS statement is optional. If you specify it, include at least one of its optional parameters.

### Syntax:

```
OPTIONS
  [AGEOUTPERIOD=minutes]
  [CONNRETRYMAX=n]
  [CONNRETRYWAIT=seconds]
  [LARGEOPS=number_of_operations]
  [MEMOPS=number_of_log_records]
  [MONITOR INTERVAL=minutes]
  [ROWID={Y|N}]
  [ROW_MOVEMENT_FAIL={Y|N}]
  [RSTRADV=seconds]
  [SPILLMAX=kilobytes]
```

```
[TIME_STAMP_MODE={LOGTIME|COMMITTIME|BEGINTIME}
[TRUNCINVALIDCHARS={Y|N}
;
```

#### Parameters:

##### **AGEOUTPERIOD=minutes**

The age, in number of minutes, that an outstanding UOW that has no change records of CDC interest must reach before it is removed from the calculation of the CDC restart point. The age is calculated as the time difference between the start of the outstanding UOW and the start of the most recent UOW. This age-out processing occurs during the monitoring interval.

Use this parameter to prevent CDC failures that can occur if you shut down and then restart capture processing while the transaction is outstanding. After the restart, the archived redo log in which the outstanding UOW started might not be available, causing the Express CDC log reader to fail.

Valid values are 60 to 43200. By default, no value is specified and this parameter is disabled.

**Note:** Oracle stores all time values in the log in local time. As a result, at the beginning or end of daylight saving time, a UOW might age out an hour late or an hour early.

##### **CONNRETRYMAX**

Optional. The maximum number of times that PowerExchange tries to reconnect to an Oracle source database or ASM instance. Use this parameter in conjunction with the CONNRETRYWAIT parameter only if you get multiple PWX-36086 messages followed by an Oracle error that indicates a dropped connection. This parameter can help improve connection resiliency. Valid values are 0 to 86400. A value of 0 results in no retries. Default is 12.

##### **CONNRETRYWAIT**

Optional. The number seconds that PowerExchange waits between attempts to reconnect to an Oracle source database or ASM instance. Use this parameter in conjunction with CONNRETRYMAX parameter only if you get multiple PWX-36086 messages followed by an Oracle error that indicates a dropped connection. This parameter can help improve connection resiliency. Valid values are 1 to 300. Default is 5 seconds.

##### **LARGEOPS**

Optional. Overrides the default value that PowerExchange uses to identify transactions as large transactions for reporting purposes. Enter the number of log records that a transaction must process to be considered a large transaction. A log record might contain multiple DML operations or part of a single operation.

PowerExchange issues status messages for large transactions that meet this criteria. If PowerExchange issues too many messages, you can increase this value to limit the number of messages.

Valid values are 1 through 2147483 (1000 through 2,147,483,000 operations). The default value is one fifth of the MEMOPS value, rounded to the nearest thousand. For example, if the MEMOPS value is 5120, the default LARGEOPS value is 1000 (1,000,000 operations).

##### **MEMOPS**

Optional. The maximum number of redo log records that contain DML operations that PowerExchange can stage in memory while processing Oracle transactions.

Valid values are 1000 through 1048576. Default is 5120.

##### **MONITOR\_INTERVAL**

Optional. The time interval, in minutes, at which PowerExchange checks transaction activity for long outstanding transactions and large transactions. A long outstanding transaction is one that remains

active for two monitoring intervals, and a large transaction is one that meets the LARGEOPS criteria. When this interval elapses, PowerExchange issues messages that identify the large transactions and long outstanding transactions and report their processing activity. PowerExchange also issues messages that identify the current position in the change stream. Valid values are 0 through 720. A value of 0 disables monitoring. Default is 5.

#### **ROWID**

Controls whether Oracle physical rowid values are included in captured change records for tables that do not have Oracle row movement enabled. PowerExchange writes the rowid values to the PowerExchange-generated DTL\_\_CAPXROWID column.

For example, you might want to use this parameter if you have unkeyed source tables on which you need to perform some processing that requires a unique row ID when extraction sessions run.

Valid values are:

- **N.** Do not capture rowid values. The DTL\_\_CAPXROWID column contains null values.
- **Y.** Capture rowid values. The DTL\_\_CAPXROWID column contains rowid values for tables that do not have row movement enabled.

Default is N.

#### **ROW\_MOVEMENT\_FAIL**

When ROWID=Y, controls whether PowerExchange Express CDC for Oracle processing fails or continues if PowerExchange detects that row movement is enabled for a source table.

You might want to continue processing if you do not need to capture rowid values for tables that have row movement enabled.

Valid values are:

- **Y.** CDC processing fails if PowerExchange encounters a table with row movement enabled.
- **N.** CDC processing continues. For any table that has row movement enabled, the DTL\_\_CAPXROWID column contains null values.

Default is Y.

#### **RSTRADV**

The number of seconds that PowerExchange waits before advancing restart and sequence tokens for a registered data source during periods when UOWs do not include any changes of interest for the data source. When the wait interval expires, PowerExchange returns the next committed "empty UOW," which includes only updated restart information.

PowerExchange resets the wait interval to 0 when one of the following events occur:

- PowerExchange completes processing a UOW that includes changes of interest.
- PowerExchange returns an empty UOW because the wait interval expired without PowerExchange receiving any changes of interest.

Valid values are 0 through 86400. No default is provided.

If you do not specify RSTRADV, PowerExchange does not advance restart and sequence tokens for a registered source during periods when PowerExchange receives no changes of interest. In this case, when PowerExchange warm starts, it reads all changes, including those not of interest for CDC, from the restart point.

**Note:** A value of 0 can degrade performance. In addition to the UOWs that contain changes for registered sources of interest, PowerExchange returns an empty UOW for every UOW that does not contain changes for the registered sources of interest.

#### SPILLMAX

The maximum size, in kilobytes, of a single PowerExchange Express CDC for Oracle spill file on UNIX. PowerExchange writes transactions to spill files when it does not have sufficient memory to store them. After PowerExchange processes all of the transactions in a spill file, the spill file is freed.

A large spill file retains disk space longer before being freed than a smaller-sized spill file.

Valid values are 1 through 2097151. Default is 10240.

#### TIME\_STAMP\_MODE

The type of timestamp that PowerExchange records in the generated DTL\_\_CAPXTIMESTAMP column of each change record for a transaction. Usually, you specify this parameter only if you want to display the Oracle commit timestamp instead of the Oracle log timestamp.

Options are:

- **LOGTIME.** The timestamp that Oracle periodically writes to the Oracle archived logs.
- **COMMITTIME.** The timestamp of the transaction commit on the source database. Specify this option if you use the timestamp to calculate latency.
- **BEGINTIME.** The timestamp of the begin UOW log record.

Default is LOGTIME.

#### TRUNCINVALIDCHARS

Indicates whether to detect and remove invalid multibyte characters that appear at the end of a character field in an Oracle source table. The characters are invalid because they have been truncated. If you allow these invalid characters to be passed to a PowerCenter workflow that has an Oracle target, PowerCenter might corrupt subsequent columns in the target table when running in Unicode mode.

Valid values are:

- **Y.** Remove invalid multibyte characters from the source data. The invalid characters are not passed to PowerCenter or applied to the target database.
- **N.** Capture the invalid multibyte characters and pass them to PowerCenter. When the PowerCenter workflow tries to write the invalid characters to the target, Oracle issues an ORA-02290 error that reports a check constraint violation on the target database. Also, PowerCenter might corrupt data in the subsequent columns in the target table.

Default is N. If you receive the ORA-02290 error in this situation, set this parameter to Y.

## RAC Statement

The RAC statement specifies the number of member instances in the Oracle Real Application Cluster (RAC) from which you capture Oracle changes.

You must define this statement if you use PowerExchange Express CDC in an Oracle RAC environment.

#### Syntax:

```
RAC MEMBERS=number_of_instances;
```

#### Parameter:

## MEMBERS

Required. The number of all member instances in the RAC, including both active and inactive member instances. Valid values are 1 to 100. Default is 1.

## READER Statement

The READER statement provides parameters that the PowerExchange Express CDC log reader uses to read redo log files, issue log-read status messages, connect to an Oracle Automatic Storage Management (ASM) instance, and check the status of the Oracle server during periods of inactivity.

The READER statement is required.

### Syntax:

```
READER
  MODE={ACTIVE | ARCHIVEONLY}
  [ACTIVELOGMASK=mask]
  [ARCHIVEDEST1=location]
  [ARCHIVEDEST2=location]
  [ASM_ASSYSASM={Y|N}]
  [ASM_CONNECT_STRING=tns_connect_string]
  [ASM_EPWD=encrypted_password]
  [ASM_PASSWORD=password]
  [ASM_USERID=user_id]
  [READBUFFSIZE=kilobytes]
  [STATUSCHECKINTERVAL=hundredths_of_seconds]
  [STATUSREPORTINTERVAL=seconds]
;
```

### Parameters:

#### MODE

Required. An option that indicates the source of and types of redo logs that the PowerExchange Express CDC log reader reads. Valid options are:

- **ACTIVE.** The PowerExchange Express CDC for Oracle log reader reads active and archived redo logs from the Oracle online system. With this option, you can run the PowerExchange for Linux, UNIX, and Windows in either continuous or batch mode.
- **ARCHIVEONLY.** The PowerExchange Express CDC log reader reads archived redo logs only. After reading all of the available archived logs, the log reader checks for additional archived logs to read based on the STATUSCHECKINTERVAL parameter.

In this mode, the log reader determines the current end of log (EOL) at initialization by using one of the following values:

- For non-RAC instances, the high SCN of the last log archived.
- For RAC instances, the lowest high SCN of the last log archived across all active nodes.

This option is a suitable choice when the PowerExchange Logger for Linux, UNIX, and Windows runs in batch mode and shuts down at EOL. In this case, you should coordinate the PowerExchange Logger runs with Oracle log file switches.

**Note:** If the PowerExchange Logger runs in continuous mode, this option can increase CDC latency because the log reader is idle for long periods while waiting for archived logs to become available.

Default is ACTIVE.

#### ACTIVELOGMASK

Optional. When the MODE parameter is set to ACTIVE, this parameter specifies a mask for selecting active redo logs for the PowerExchange Express CDC log reader when the Oracle instance uses multiplexing of redo logs. The log reader compares the mask against the member names in an active

redo log group to determine which log to read. In the mask, you can use the asterisk (\*) wildcard to represent zero or more characters.

The mask can be up to 128 characters in length. The mask is case-sensitive on Linux or UNIX systems but not on Windows systems.

#### **ARCHIVEDEST1**

Optional. The archived redo log destination from which the PowerExchange Express CDC log reader first attempts to select archived logs. Enter an integer that corresponds to the number *n* in a LOG\_ARCHIVE\_DEST\_*n* initialization parameter in Oracle, where *n* can be a number from 1 to 10.

By default, the log reader uses 1 for the LOG\_ARCHIVE\_DEST\_1 location. If you want to specify another archived log destination as the primary destination for the log reader to use, define this parameter.

#### **ARCHIVEDEST2**

Optional. The secondary archived redo log destination from which the PowerExchange Express CDC log reader selects archived logs when the logs cannot be read from the primary archived log destination. For example, the archived logs from the primary archive log destination might have been damaged or deleted. Enter an integer that corresponds to the number *n* in a LOG\_ARCHIVE\_DEST\_*n* initialization parameter in Oracle. Usually, this value is a number other than 1. No default is provided.

#### **ASM\_ASSYSASM**

Optional. If you use Oracle 11g ASM or later and want the PowerExchange Express CDC log reader to use a user ID that has SYSASM authority to connect to the ASM instance, set this parameter to Y. Then specify a user ID that has SYSASM authority in the ASM\_USERID parameter. If you want to use a user ID that has SYSDBA authority, set this parameter to N. Default value is N.

#### **ASM\_CONNECT\_STRING**

Optional. In an Oracle ASM environment, the Oracle connection string, defined in TNS, that the PowerExchange Express CDC log reader uses to connect to the ASM instance that manages storage of active and archived redo logs for the source database.

#### **ASM\_EPWD**

Optional. In an Oracle ASM environment, an encrypted password for the user that is specified in the ASM\_USERID parameter. The PowerExchange Express CDC log reader uses this password and the ASM user ID to connect to the ASM instance that manages storage of active and archived redo logs for the source database. Define ASM\_EPWD or ASM\_PASSWORD but not both.

#### **ASM\_PASSWORD**

Optional. In an Oracle ASM environment, a clear text password for the user that is specified in the ASM\_USERID parameter. The PowerExchange Express CDC log reader uses this password and the ASM user ID to connect to the ASM instance that manages storage of active and archived redo logs for the source database. Define ASM\_EPWD or ASM\_PASSWORD but not both.

#### **ASM\_USERID**

Optional. In an Oracle ASM environment, an Oracle user ID that the PowerExchange Express CDC log reader uses to connect to the ASM instance that manages storage of active and archived redo logs for the source database. This user ID must have SYSDBA or SYSASM authority. To use SYSASM authority, set the ASM\_ASSYSASM parameter to Y.

#### **READBUFFSIZE**

Optional. The default buffer size, in kilobytes, that the PowerExchange Express CDC log reader uses to read a redo log. PowerExchange Express CDC for Oracle can automatically expand this buffer size, if necessary.

Valid values are 1 through 262144. Default is 10240.

#### STATUSCHECKINTERVAL

Optional. The time interval, in hundredths of seconds, that the PowerExchange Express CDC log reader waits at EOL when no additional data is available to read, before checking with Oracle on the following items:

- In ACTIVE mode, whether the active redo log file is still valid and available.
- In ARCHIVEONLY mode, if additional Oracle archive logs are available to read.
- Whether the number of data blocks that the log reader read matches the number of data blocks that Oracle wrote.

Valid values are 1 through 8640000. Default is 200.

**Note:** If PowerExchange issues message PWX-36171, the number of blocks read did not match the number of blocks written. In this case, try increasing the STATUSCHECKINTERVAL value so that any stale NFS read buffer has time to refresh and accept new data for the log reader to process.

#### STATUSREPORTINTERVAL

Optional. The frequency, in seconds, at which the PowerExchange Express CDC log reader issues message PWX-36151 to report log read progress.

Valid values are 1 through 86400. Default is 120.

## STANDBY Statement

The STANDBY statement defines a connection to an Oracle physical standby database for change data capture when the database is *not* open for read-only access.

The STANDBY statement is optional. Use it only when the database is not open for read-only access. To access a database that is not open, you must have SYSDBA authority.

#### Syntax:

```
STANDBY
CONNECT STRING=capture_connect_string
[APPLYACTIVE={N|Y}]
[EPWD=database_encrypted_password|PASSWORD=database_password]
[USERID=database_user_id]
;
```

#### Parameters:

##### APPLYACTIVE

Indicates whether PowerExchange Express CDC can process the standby redo logs on the standby system up to the highest low SCN when the Oracle does not progress apply processing beyond the tip of the last archived log, even though more recent changes are available in the standby logs. Set this parameter to Y if you want to perform near-real-time capture in this situation. Default is N.

##### CONNECT\_STRING

An Oracle connection string, defined in TNS, that PowerExchange Express CDC for Oracle uses to connect to the Oracle physical standby database for change capture when the database is not open for read-only access.

##### EPWD

An encrypted password that PowerExchange Express CDC for Oracle uses to connect to the Oracle physical standby database for change capture.

You must specify either the EPWD or PASSWORD parameter, but do not specify both.

#### **PASSWORD**

A clear text password that PowerExchange Express CDC for Oracle uses to connect to the Oracle physical standby database for change capture.

You must specify either the PASSWORD or EPWD parameter, but do not specify both.

#### **USERID**

A user ID that PowerExchange Express CDC for Oracle uses to connect to the Oracle physical standby database for change capture. This user ID must have SYSDBA authority.

#### **Usage Notes:**

- To access the Oracle data dictionary on the primary system, PowerExchange Express CDC for Oracle uses the DATABASE statement, which points to the primary system when the database is not open for read-only access to the logs.

### Example PowerExchange Express CDC for Oracle Configuration File

This example PowerExchange Express CDC for Oracle configuration file contains the required statements only.

```
DICTIONARY
MODE=STATIC
SOURCE=ONLINE;
READER
MODE=ACTIVE;
```

## Managing PowerExchange Express CDC for Oracle

After CDC is running, you might need to perform some tasks to maintain and manage your PowerExchange Express CDC for Oracle environment occasionally.

These tasks include:

- Monitoring CDC processing
- Adding a capture registration
- Stopping CDC processing for a table
- Changing the structure of a table

### Monitoring PowerExchange Express CDC for Oracle

To determine the progress of the Express CDC log reader in reading the redo logs, look for key messages in the PowerExchange message log file.

When a redo log switch occurs, PowerExchange Express CDC for Oracle issues messages such as the following example messages:

```
110324 125031 WIN32 2216 PWX-36145 ORAD Info: Low SCN 0x0000.05c2b16e.0000[96645486].
Low SCN Time 03/24/2011 01:34:58.
110324 125031 WIN32 2216 PWX-36146 ORAD Info: Next SCN 0x0000.05c32597.0000[96675223].
Next SCN Time 03/24/2011 07:11:09.

110324 125051 PWX-36144 ORAD Info: Reader processing active Log file \\s160020\fdrive
```



```
\ORACLE\PRODUCT\ORADATA\ORCL\REDO03.LOG, SEQ 1797.  
110324 125051 PWX-36145 ORAD Info: Low SCN 0x0000.05c32597.0000[96675223]. Low SCN Time  
03/24/2011 07:11:09.
```

To determine how far behind the Express CDC log reader is in processing the redo logs, look at each PWX-36145 message and compare the time when the message was issued to the Low SCN Time in the message. Then compare this time interval across multiple PWX-36145 messages. The Low SCN represents the lowest SCN in the redo log, and the Next SCN is the SCN at the time the log was closed.

Use the example messages to perform the analysis:

- In the first PWX-36145 message, compare the message time of 12:50:31 to the Low SCN Time of 1:34:58. The difference is about 11 hours 15 minutes.
- In the second PWX-36145 message, compare the message time of 12:50:51 to the Low SCN Time of 7:11:09. The difference is about 5 hours 40 minutes.

In this case, the large reduction in the time interval indicates that the Express CDC log reader is catching up on reading change records from the redo logs.

Also, the PWX-36151 message, which is issued periodically, gives an indication of PowerExchange Express CDC for Oracle progress. When the Express CDC log reader is either catching up or falling behind, this message indicates that the log reader is reading an archived log, an active log, or the current log. After the log reader has caught up, this message indicates that the log reader is reading at the tip of the current log sequence. The following example PWX-36151 message indicates that the log reader has caught up with reading changes in the current log:

```
110324 145352 WIN32 7844 PWX-36151 ORAD Info: Reading at the tip of the current log  
sequence 1798, block 28475. Low SCN 0x0000.05c3bee9.0000[96714473], Low SCN time  
03/24/2011 14:53:48.
```

In this message, the Low SCN Time is the earliest time that the Oracle log writer recorded in the last span of the log that was read by the Express CDC log reader. To determine change capture latency, you can compare the Low SCN Time to the time when the message was issued, in a manner similar to the preceding analysis.

## Adding a Capture Registration

After PowerExchange Express CDC for Oracle is running, you might need to add a capture registration for another Oracle table.

1. In PowerExchange Navigator, create the capture registration.  
Make sure that you include the following settings:
  - In the **Condense** list, select **Part**.
  - In the **Supplement Log Group Name** box, enter a name for the supplemental log group for the table. The PowerExchange Navigator generates DDL for creating the supplemental log group. If you select **Execute DDL now**, the PowerExchange Navigator runs the DDL when you complete the registration. If you do not have the authority to run the DDL, ask your DBA to run it.
  - In the **Status** list, select **Active**.
2. Run the DDL for creating a supplemental log group for the table, if you did not enable the PowerExchange Navigator to do so at registration completion.
3. If you use the PowerExchange Logger for Linux, UNIX, and Windows, shut down the PowerExchange Logger and then warm start it.  
The PowerExchange Logger begins capturing change data for the additional table.
4. In PowerCenter, import the extraction map for the table to create a source definition.
5. Add the source definition to a new mapping or to an existing mapping.

If you add the source definition to an existing mapping, you must stop the workflow first.

6. Start the workflow that processes the added source table.

## Stopping CDC Processing for a Table

If you no longer need to capture change data for a table, you can stop PowerExchange Express CDC for Oracle processing for the table.

1. In PowerExchange Navigator, open the capture registration and set the **Status** option to **History**.  
A capture registration that has a status of **History** cannot be activated again.
2. If you use the PowerExchange Logger for Linux, UNIX, and Windows, shut down the PowerExchange Logger and then warm start it.
3. Drop the supplemental log group for the table by using the following SQL:

```
ALTER TABLE schema.table_name DROP SUPPLEMENTAL LOG GROUP
```

Oracle stops logging full before-images and after-images of column data.

**Note:** If you were to add the supplemental log group for the table again because you had to reinstate change capture, you would need to rematerialize the target.

## Stopping CDC Processing Temporarily

You might need to temporarily stop PowerExchange Express CDC for Oracle processing to troubleshoot issues or to perform a maintenance task on the target database.

If you use the PowerExchange Logger for Linux, UNIX, and Windows, shut down the PowerExchange Logger to stop CDC processing for all source tables. Later, you can warm start the PowerExchange Logger to resume change capture processing without any change data loss. This method is preferable.

## Changing the Structure of an Oracle Source Table

Occasionally, you might need to make DDL changes to a registered Oracle source table that add, alter, or drop columns from which PowerExchange Express CDC for Oracle captures changes. Learn how to switch to the new table definition in a manner that preserves access to previously captured data.

You do not need to do this task in the following situations:

- You selectively capture change data for a subset of columns, and the DDL changes do not affect any of these columns.
  - You need to stop change data extraction processing for a column. In this case, remove the column from the extraction map and do not edit the capture registration. PowerExchange still captures change data for the column but does not extract it when CDC sessions run.
1. Stop data change activity (inserts, updates, and deletes) on the table.
  2. Verify that any change data that was captured under the current table definition has completed extraction processing. Then stop all PowerCenter workflows that extract change data for the table.
  3. If you use the PowerExchange Logger for Linux, UNIX, and Windows, shut down the Logger.
  4. In the PowerExchange Navigator, open the original capture registration and set its status to **History**.  
PowerExchange does not capture change data based on capture registrations that have a status of **History** or **Inactive**.

**Tip:** If you no longer need to capture change data from a column, you can remove the column from the extraction map without changing the capture registration. Change data for the column is still captured but is not extracted.

5. Drop the supplemental log group for the table.
6. Make the DDL changes to the table.
7. In the PowerExchange Navigator, create a new capture registration for the table that reflects the DDL changes.

Make sure that you include these settings:

- In the **Condense** list, select **Part**.
  - In the **Supplement Log Group Name** box, enter a name for the supplemental log group that must be created for the table. The PowerExchange Navigator generates DDL for creating the supplemental log group. If you select **Execute DDL now**, the PowerExchange Navigator runs the DDL when you complete the registration. If you do not have the authority to run the DDL, you can ask your DBA to run it.
  - In the **Status** list, select **Active**.
8. If you shut down the PowerExchange Logger, warm start it.  
The PowerExchange Logger begins capturing changes based on the new capture registration.
  9. Change the target table definition to reflect the source table changes, if necessary.
  10. In PowerCenter Designer, import the new extraction map for the altered source table to create a new source definition. Also, if you changed the target table, edit or re-create the target definition. Then, edit the mapping, if necessary.
  11. If necessary, rematerialize the target tables and then create new restart tokens.
  12. Allow change activity on the table to resume.
  13. Start the workflows again.

Extraction processing resumes.

## CHAPTER 8

# Remote Logging of Data

This chapter includes the following topics:

- [Remote Logging Overview, 188](#)
- [Requirements for Capture Registrations, 193](#)
- [Security Settings for Data from z/OS Sources, 193](#)
- [Configuration Tasks for Remote Logging, 194](#)
- [Example of Remote Logging from a z/OS Data Source, 200](#)
- [Example of Remote Logging from a DB2 for i5/OS Data Source, 202](#)
- [Example of Remote Logging from an Oracle CDC with LogMiner Data Source, 204](#)

## Remote Logging Overview

You can log change data from any supported data source to PowerExchange Logger for Linux, UNIX, and Windows log files on another system.

You can log change data from data sources on i5/OS or z/OS to PowerExchange Logger log files on a Linux, UNIX, or Windows system. The PowerExchange Logger for Linux, UNIX, and Windows reads change data from PowerExchange on the source and logs the data to its log files. CDC sessions that run in continuous extraction mode can then extract the change data from the PowerExchange Logger log files instead of from the source.

The benefits of logging or relogging change data off of the source system depend on the source type and CDC environment. You can use remote logging to reduce resource consumption on the source system, move some resource-intensive CDC processing to the remote system, and reduce the network overhead of data transfer.

### RELATED TOPICS:

- [“Requirements for Capture Registrations” on page 193](#)
- [“Configuration Tasks for Remote Logging” on page 194](#)
- [“Customizing the PowerExchange Logger Configuration File for Remote Logging” on page 194](#)
- [“Customizing the dbmover Configuration File on the System to Which Data Is Logged” on page 198](#)
- [“Customizing the dbmover Configuration File on the PowerCenter Integration Service System” on page 199](#)
- [“Example of Remote Logging from an Oracle CDC with LogMiner Data Source” on page 204](#)
- [“Logging Data from Remote z/OS or i5/OS Sources” on page 43](#)

## Remote Logging of Data from i5/OS or z/OS Sources

You can use the PowerExchange Logger for Linux, UNIX, and Windows to extract change data for data sources on i5/OS and z/OS and relog that data to a less costly Linux, UNIX, or Windows system. Multiple PowerCenter CDC sessions can then retrieve the change data from the local PowerExchange Logger for Linux, UNIX, and Windows log files.

For i5/OS and z/OS sources, the remote logging of data to a Linux, UNIX, or Windows system has the following benefits:

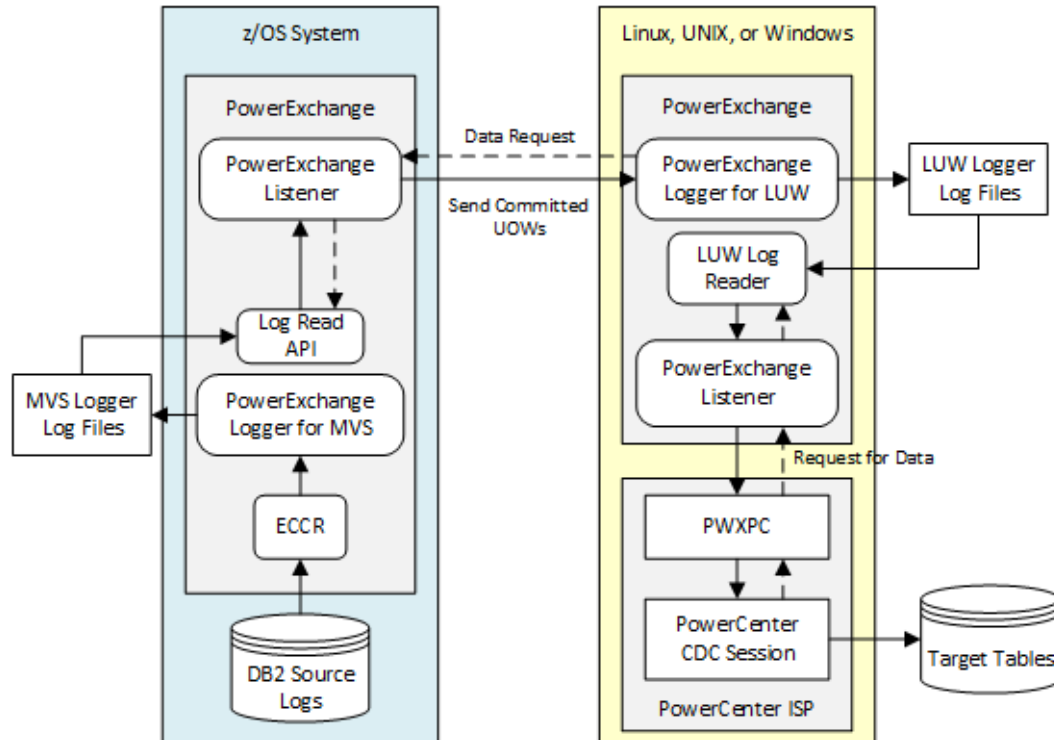
- Moves resource-intensive, column-level processing and UOW Cleanser processing off of the i5/OS or z/OS system onto the Linux, UNIX, or Windows system where the PowerExchange Logger for Linux, UNIX, and Windows runs.
- Extracts change data from the DB2 for i5/OS journal receivers or PowerExchange Logger for MVS log files on z/OS in a single pass and transmits that data over the network to the PowerExchange Logger for Linux, UNIX, and Windows. The data is then available locally for PowerCenter CDC sessions to process. This single-pass processing reduces network traffic and avoids the overhead of multiple data extraction reads.
- Reduces costly CPU usage, disk space, and CDC processing time on the i5/OS or z/OS source system.

To configure this remote logging scenario, you must specify the **CAPTURE\_NODE** statement in the PowerExchange Logger for Linux, UNIX, and Windows configuration file, `pwxccl.cfg`, on the system where the Logger for Linux, UNIX, and Windows runs. The **CAPTURE\_NODE** statement specifies the node name of the PowerExchange Listener that runs on the source system. When you create the registration group in the PowerExchange Navigator, enter the node name of the PowerExchange Listener that runs on the source system in the **Location** field. In PowerCenter, configure a PWX CDC Real Time connection for the PowerCenter CDC sessions that process change data from the source. In the connection attributes, set the **Location** attribute to the node name of the PowerExchange Listener that runs on the system where the PowerExchange Logger log files reside and set the **Mapping Location** attribute to the node name of the PowerExchange Listener that runs on the source system where the extraction maps reside.

**Note:** When the PowerExchange Logger for Linux, UNIX, and Windows runs on the PowerCenter Integration Service Platform (ISP) machine, you can use a Local connection rather than run a PowerExchange Listener on this machine. However, Informatica recommends that you run a PowerExchange Listener on the PowerCenter ISP machine so that you can issue commands to display information about the active PowerExchange Listener tasks, print PowerExchange Listener monitoring statistics, and stop the PowerExchange Listener task, if necessary.

For example, you can configure the PowerExchange Logger for Linux, UNIX, and Windows to extract DB2 for z/OS change data from PowerExchange Logger for MVS logs files on a z/OS system and then relog that data

to PowerExchange Logger for Linux, UNIX, and Windows log files on the PowerCenter ISP machine. The following image shows this remote logging configuration:



In this scenario, set the PowerExchange Logger CAPTURE\_NODE statement to point to the node name of the PowerExchange Listener on the z/OS system with the DB2 logs. Set the PowerCenter **Location** connection attribute to the node name of the PowerExchange Listener on the PowerCenter ISP machine where the PowerExchange Logger for Linux, UNIX, and Windows runs. Set the **Map Location** connection attribute to point to the node name of the PowerExchange Listener on the z/OS system.

The PowerExchange Logger for Linux, UNIX, and Windows sends a request for change data to the PowerExchange Listener on z/OS. This PowerExchange Listener contacts the Log Read API (LRAPI) to read captured change data from the PowerExchange Logger for MVS log files. The PowerExchange Listener on z/OS transmits the change data in a single stream over the network to the PowerExchange Logger for Linux, UNIX, and Windows. The UOW Cleanser runs on the Powercenter ISP machine to cleanse the data, and then the PowerExchange Logger for Linux, UNIX, and Windows relogs the data in its local log files. When a Powercenter CDC session runs and requests change data for the tables of CDC interest, the PowerExchange Client for PowerCenter (PWXP) requests change data from the PowerExchange Listener on the system with the LUW Logger log files. The PowerExchange Listener contacts the local PowerExchange Logger Log Reader to read change data from the Logger log files. PWXP makes the data available to the PowerCenter CDC session. Multiple PowerCenter CDC sessions can extract change data from the local PowerExchange Logger log files.

## Remote Logging of Data from Linux, UNIX, or Windows Sources

You can use the PowerExchange Logger for Linux, UNIX, and Windows to extract change data from sources on a Linux, UNIX, or Windows system and log that data to another Linux, UNIX, or Windows system.

PowerCenter CDC sessions can then extract the change data from the log files on the PowerExchange Logger system.


For Linux, UNIX, and Windows sources, the remote logging of data has the advantage of maximizing the performance of the capture process while minimizing the impact on the database server. Remote logging has the following benefits:

- Reduces disk space and CPU usage on the source database server by moving PowerExchange Logger processing and log files to another Linux, UNIX, or Windows system.
- Retains local access to the source database and database logs, which eliminates the latency of accessing the source across a network and avoids having to configure NFS to export the database logs.
- Retains object filtering and any internal UOW cleansing processing on the source system. Only the data of interest is sent across the network.
- Allows multiple PowerCenter sessions to access the same logged data without increasing overhead on the database server and, if the PowerExchange Logger runs on the same server as PowerCenter, without affecting network latency.

To configure this remote logging scenario, you must specify the **CAPTURE\_NODE** statement in the PowerExchange Logger for Linux, UNIX, and Windows configuration file, `pwxccl.cfg`, on the system where the Logger runs. The **CAPTURE\_NODE** statement specifies the node name of the PowerExchange Listener that runs on the source system. When you create the registration group in the PowerExchange Navigator, enter the node name of the PowerExchange Listener that runs on the source system in the **Location** field. In PowerCenter, configure a PWX CDC Real Time connection for the PowerCenter CDC sessions that process change data from the source. In the connection attributes, set the **Location** attribute to the node name of the PowerExchange Listener that runs on the system where the PowerExchange Logger log files reside, and set the **Mapping Location** attribute to the node name of the PowerExchange Listener that runs on the source system where the extraction maps reside.

**Note:** When the PowerExchange Logger for Linux, UNIX, and Windows runs on the PowerCenter ISP machine, you can use a Local connection rather than run a PowerExchange Listener on this machine. However, Informatica recommends that you run a PowerExchange Listener on the PowerCenter ISP machine so that you can issue commands to display information about the active PowerExchange Listener tasks, print PowerExchange Listener monitoring statistics, and stop the PowerExchange Listener task, if necessary.

For example, configure the PowerExchange Logger for Linux, UNIX, and Windows to extract DB2 change data from DB2 database logs on a Linux system and then log that data to PowerExchange Logger log files on the PowerCenter ISP system. The following image shows this type of configuration:



In this scenario, set the PowerExchange Logger CAPTURE\_NODE statement to the node name of the PowerExchange Listener on the DB2 source system. Set the PowerCenter **Location** connection attribute to the node name of the PowerExchange Listener on the PowerCenter ISP machine where the PowerExchange Logger runs. Set the **Map Location** connection attribute to the node name of the PowerExchange Listener on the DB2 source system.

When the PowerExchange Logger sends a change data capture request, the PowerExchange Listener on the source system communicates with the CAPI to retrieve change data from the DB2 database logs by means of the DB2 LogRead API. The PowerExchange Listener sends only the committed UOWs for the objects of CDC interest to the PowerExchange Logger on the remote system. The PowerExchange Logger logs the data in its log files. When a PowerCenter CDC session requests change data for the tables of CDC interest, the PowerExchange Client for PowerCenter (PWXPC) communicates with the PowerExchange Listener on the PowerExchange Logger system to get the change data from the local PowerExchange Logger log files.



# Requirements for Capture Registrations

For the PowerExchange Logger for Linux, UNIX, and Windows to log change data from a remote source, verify that the capture registrations are compatible with the following requirements:

- To use the PowerExchange Logger for Linux, UNIX, and Windows, you must configure capture registrations for partial condense processing. In the PowerExchange Navigator, select **Part** in the **Condense** list for each registration. If you have remote i5/OS or z/OS data sources with capture registrations that specify **Full** for the **Condense** option, the PowerExchange Logger for Linux, UNIX, and Windows ignores these registrations. The PowerExchange Logger also ignores any capture registration that specify **None** for the **Condense** option.
- A PowerExchange Logger for Linux, UNIX, and Windows process must be able to read all of the capture registrations that it uses from a single CCT file on the source system.
- For the remote data sources, you cannot use capture registrations that were created from data maps that use any of the following features:
  - User access methods
  - User-defined fields that invoke programs by using the CALLPROG function
  - Record-level exits

## Security Settings for Data from z/OS Sources

For the highest level of security for data from z/OS data sources, set the SECURITY option to 2 in the z/OS DBMOVER configuration member where the extraction maps are located. With this setting, PowerCenter CDC sessions are permitted to extract z/OS data from PowerExchange Logger for Linux, UNIX, and Windows log files only if their user credentials pass z/OS security checking.

When defining a PWXPC connection for the CDC sessions that extract data from the PowerExchange Logger log files, enter a valid z/OS user ID and password in the **Map Location User** and **Map Location Password** connection attributes. If the location of the log files is not local, enter the z/OS user ID and password in the **User Name** and **Password** connection attributes for use by the PowerExchange Listener on the Linux, UNIX, or Windows system where the log files reside.

For data extraction, these z/OS user credentials must have the following permissions:

- READ access to the PowerExchange data set that is defined in the DTLCAMAP DD statement of the PowerExchange Listener JCL
- READ access to CAPX.CND.\* resource profiles in the FACILITY class, which are managed by your z/OS security product.

For more information about security, see the *PowerExchange Reference Manual*.

# Configuration Tasks for Remote Logging

To log change data to remote PowerExchange Logger for Linux, UNIX, and Windows log files and have PowerCenter CDC sessions extract data from those log files, complete the following configuration tasks:

1. Install PowerExchange on the system where the PowerExchange Logger log files will be located.
2. Customize the `pwxccl.cfg` configuration file on the system with the PowerExchange Logger log files.
3. Customize the `dbmover` configuration file on the system with the PowerExchange Logger log files. Copy the source-specific `CAPL_CONNECTION` statements from the source system to the `dbmover` file on the system with the PowerExchange Logger log files.

**Note:** Each PowerExchange Logger must have a unique `pwxccl.cfg` configuration file and a unique `dbmover` configuration file.

4. Configure a `dbmover` configuration file for the PowerExchange Listener on the system with the PowerExchange Logger log files.  
You can use the same `dbmover` file for the PowerExchange Logger and the PowerExchange Listener. If you use different `dbmover` files, both files must specify the same `CAPT_PATH` value.  
If the PowerExchange Logger log files are on the PowerCenter Integration Service machine, you can use a local connection instead of the PowerExchange Listener for change data extractions.
5. If you are not using a "local" connection, start the PowerExchange Listener on the system with the PowerExchange Logger log files.
6. Start the PowerExchange Logger on the system with the PowerExchange Logger log files.
7. Customize the `dbmover` configuration file on the PowerCenter Integration Service machine.
8. Configure capture registrations for PowerExchange Logger use.
9. Configure PWX CDC Real Time connection attributes for the CDC session to extract change data from the PowerExchange Logger log files.

## Customizing the PowerExchange Logger Configuration File for Remote Logging

For the PowerExchange Logger for Linux, UNIX, and Windows to log data from a remote source, you must customize the PowerExchange Logger configuration file on the system where the PowerExchange Logger log files will reside.

PowerExchange provides a sample configuration file, named `pwxccl.cfg`, in the PowerExchange installation directory. You can copy this file and customize the copy.

For a complete list of PowerExchange Logger configuration parameters, see the PowerExchange Logger for Linux, UNIX, and Windows chapter in the *PowerExchange CDC Guide for Linux, UNIX, and Windows*.

The following table describes the parameters that are needed for remote logging:

Parameter	Description
CAPTURE_NODE	Required for remote logging. The node name that the PowerExchange Logger uses to retrieve capture registrations and change data from the source system. This name must be defined in a NODE statement in the dbmover configuration file on the system where the PowerExchange Logger runs. The PowerExchange Logger uses this node name to connect to the PowerExchange Listener on the source system. This name should correspond to the node name in the LISTENER statement on the source system.
CAPTURE_NODE_EPWD or CAPTURE_NODE_PWD	Optional. An encrypted password (EPWD) or clear text password (PWD) that is associated with the user ID specified in the CAPTURE_NODE_UID parameter. If you specify CAPTURE_NODE_UID, you must specify either CAPTURE_NODE_EPWD or CAPTURE_NODE_PWD. However, do not specify both CAPTURE_NODE_EPWD and CAPTURE_NODE_PWD.
CAPTURE_NODE_UID	A user ID that controls PowerExchange Logger read access to capture registrations and change data on the remote node that is specified in the CAPTURE_NODE parameter. Whether this parameter is required depends on the operating system of the remote node and the SECURITY setting in the DBMOVER configuration file for the PowerExchange Listener on that node.  If CAPTURE_NODE specifies a z/OS or i5/OS node that has a SECURITY setting of 0, do not specify this parameter. PowerExchange uses the user ID under which the PowerExchange Listener job runs to control access to capture registrations and change data.  If CAPTURE_NODE specifies a z/OS or i5/OS node that has a SECURITY setting of 1, you must enter a valid operating system user ID for this parameter. Otherwise, error message PWX-00231 is issued, indicating a signon failure. However, PowerExchange uses the user ID under which the PowerExchange Listener job runs to control access to capture registrations and change data.  If CAPTURE_NODE specifies a z/OS or i5/OS node that has a SECURITY setting of 2, you must enter a valid operating system user ID for this parameter. Otherwise, error message PWX-00231 is issued, indicating a signon failure. PowerExchange uses this user ID to control access to capture registrations and change data. If the specified user ID does not have the authority that is required to read capture registrations or change data, access fails.  If CAPTURE_NODE specifies a Linux, UNIX, or Windows node, enter a user ID that is valid for the data source type: <ul style="list-style-type: none"> <li>- For DB2 for Linux, UNIX, or Windows sources, enter a valid operating system user ID that has DB2 DBADM or SYSADM authority.</li> <li>- For Microsoft SQL Server instances that use SQL Server Authentication, enter a database user ID that permits access to the SQL Server distribution database. For a SQL Server instances that use Windows Authentication, PowerExchange uses the user ID under which the PowerExchange Listener was started. In this case, do not specify this parameter unless you want to specify another user.</li> <li>- For Oracle sources, if you use PowerExchange Oracle CDC with LogMiner, enter the ORACAPT user ID that you defined, which permits access to Oracle archive logs and to Oracle LogMiner. If you use PowerExchange Express CDC for Oracle, enter the ORACAPTL user ID that you defined, which permits access to the Oracle online and archive redo logs.</li> </ul>
CONDENSENAME	Optional. A name for the command-handling service for a PowerExchange Logger for Linux, UNIX, and Windows process to which pwxcmd commands are issued.  This service name must match the service name in the associated SVCNODE statement in the dbmover configuration file.

Parameter	Description
CONN_OVR	<p>Recommended. The name of the override CAPI_CONNECTION statement to use for the PowerExchange Logger. If you do not enter CONN_OVR, the PowerExchange Logger uses the default CAPI_CONNECTION in the dbmover configuration file, if specified.</p> <p>Enter a valid CAPI_CONNECTION name for the source type.</p> <p>Informatica recommends that you specify CONN_OVR because it is the only type of override that the PowerExchange Logger can use.</p>
DB_TYPE	<p>Required. The source database type. Options are:</p> <ul style="list-style-type: none"> <li>- <b>ADA</b>. For Adabas sources.</li> <li>- <b>AS4</b>. For DB2 for i5/OS sources.</li> <li>- <b>DB2</b>. For DB2 for z/OS sources.</li> <li>- <b>DCM</b>. For Datacom sources.</li> <li>- <b>IDL</b>. For IDMS log-based CDC sources.</li> <li>- <b>IMS</b>. For IMS sources.</li> <li>- <b>MSS</b>. For Microsoft SQL Server sources.</li> <li>- <b>ORA</b>. For Oracle sources.</li> <li>- <b>UDB</b>. For DB2 for Linux, UNIX, and Windows sources.</li> <li>- <b>VSM</b>. For VSAM sources.</li> </ul>

Parameter	Description
DBID	<p>Required. A source identifier, sometimes called the <i>instance</i> name, that is defined in capture registrations. When used with DB_TYPE, it defines selection criteria for capture registrations in the CCT file.</p> <p>This value must match the instance or database name that is displayed in the Resource Inspector of the PowerExchange Navigator for the registration group that contains the capture registrations.</p> <p>Enter one of the following values based on the source type:</p> <ul style="list-style-type: none"> <li>- For Adabas, enter the <b>Instance</b> name that is displayed for the registration group.</li> <li>- For Datacom, enter the <b>MUF Name</b> value that is displayed for the registration group. Alternatively, if you use Datacom synchronous CDC, enter the value of the MUF parameter in the DTLINPUT data set specified in the MUF JCL. Or, if you use Datacom table-based CDC, enter the value of REG_MUF parameter in the ECCRDCMP member of the RUNLIB library.</li> <li>- For DB2 for i5/OS, enter the <b>Instance</b> name that is displayed for the registration group. This name should match the INST parameter value in the AS4J CAPI_CONNECTION statement in the DBMOVER member of the CFG file.</li> <li>- For DB2 for Linux, UNIX, and Windows, enter the <b>Database</b> name that is displayed for the registration group.</li> <li>- For DB2 for z/OS, enter the <b>Instance</b> name that is displayed for the registration group. This name should match the RN parameter value in the DB2 statement in the RUNLIB(REPDB2OP) member.</li> <li>- For IDMS Log-based CDC, enter the <b>Logsid</b> value that is displayed for the registration group. This value should match the LOGSID parameter value in the RUNLIB(ECCRIDLP) member.</li> <li>- For IMS, enter the <b>IMSID</b> value that is displayed for the registration group. For IMS log-based CDC, this value should match the first parameter value in the IMSID statement in the RUNLIB(CAPTIMS) member.</li> <li>- For Microsoft SQL Server, this value depends on whether you also specify the optional DISTSRV and DISTDB parameters in the PowerExchange Logger configuration file: <ul style="list-style-type: none"> <li>- If you specify the optional DISTSRV and DISTDB parameters, enter a name that serves as the collection identifier for all of the registrations. This name must be one to eight characters in length and start with a letter. This name overrides the instance name that is associated with the individual registrations.</li> <li>- If you do not specify the DISTSRV and DISTDB parameters, enter the value that the PowerExchange Navigator generates and displays in the <b>Instance</b> field of the <b>Resource Inspector</b> for the registration group.</li> </ul> </li> <li>- For Oracle, enter the <b>Instance</b> name that is displayed for the registration group. This value also should match the first positional parameter of the ORACLEID statement in the dbmover configuration file.</li> <li>- For VSAM, enter the <b>Instance</b> name that is displayed for the registration group.</li> </ul>
EXT_CAPT_MASK	<p>Required. An existing directory path and a unique prefix to be used for generating the PowerExchange Logger log files.</p>

## Customizing the dbmover Configuration File on the System to Which Data Is Logged

For the PowerExchange Logger for Linux, UNIX, and Windows to log data from a remote Linux, UNIX, Windows, i5/OS, or z/OS source, you must customize the dbmover configuration file on the system where the PowerExchange Logger log files will reside.

PowerExchange provides a sample dbmover file in the PowerExchange installation directory. You can copy this file and customize the copy. For a complete list of all dbmover configuration statements, see the *PowerExchange Reference Manual*.

The following table describes the dbmover statements that are needed for remote logging:

Statement	Description
CAPT_PATH	Required. The path to the directory on the Linux, UNIX, or Windows system where the PowerExchange Logger CDCT file resides. The PowerExchange Logger stores information about its log files in the CDCT file. Each PowerExchange Logger that captures change data requires its own CDCT file.
CAPX CAPI_CONNECTION	Required. Parameters that the Consumer API (CAPI) uses for continuous extraction of change data from PowerExchange Logger for Linux, UNIX, and Windows log files. The DFLTINST parameter value in this statement must match the DBID value in the PowerExchange Logger configuration file, pwxcl.cfg.
LOGPATH	Optional. A unique path and directory for PowerExchange message log files on the Linux, UNIX, or Windows system where the PowerExchange Logger logs data in its log files.
NODE	Required. Information that PowerExchange uses to connect to the PowerExchange Listener on the source system from which change data is captured. This information includes a unique user-defined node name, the TCP/IP host name, and the port number. The node name that you enter in this statement must match the CAPTURE_NODE parameter value in the PowerExchange Logger configuration file.
Source-specific CAPI_CONNECTION	Required. A named set of parameters that the CAPI uses to connect to the change stream for a source type and control CDC processing. Copy the source-specific CAPI_CONNECTION statements from the DBMOVER configuration file on the source system. Use one of the following the statement types, as appropriate for the source from which you are remotely logging data: <ul style="list-style-type: none"><li>- For DB2 for i5/OS sources, use AS4J and UOWC CAPI_CONNECTION statements.</li><li>- For DB2 for Linux, UNIX, and Windows sources, use UDB CAPI_CONNECTION statements.</li><li>- For Microsoft SQL Server sources, use MSQLE CAPI_CONNECTION statements.</li><li>- For Oracle LogMiner sources, use ORCL and UOWC CAPI_CONNECTION statements.</li><li>- For Express CDC for Oracle sources, use ORAD CAPI_CONNECTION statements.</li><li>- For z/OS sources, use LRA and UOWC CAPI_CONNECTION statements. Remove the z/OS-specific parameters from the UOWC statement.</li></ul>
SVCNODE	Optional. The TCP/IP port on which a command-handling service for a PowerExchange process, such as a PowerExchange Logger for Linux, UNIX, and Windows process, listens for pwxcmd commands.
TRACING	Optional. Enables PowerExchange alternative logging and specifies attributes for the alternative log files. PowerExchange uses the alternative log files instead of the default PowerExchange message log file to store messages.

## Customizing the dbmover Configuration File on the PowerCenter Integration Service System

If you log change data in PowerExchange Logger for Linux, UNIX, and Windows log files on a system other than the source system, customize the dbmover configuration file on the PowerCenter Integration Service system, where the CDC sessions run, to identify the source and PowerExchange Logger nodes.

Add NODE statements for the PowerExchange Listeners that run on the following systems:

- The source system where the capture registrations reside and from which the PowerExchange Logger for Linux, UNIX, and Windows reads change data.
- The remote system where the PowerExchange Logger logs change data in its log files.

## Configuring Capture Registrations for the PowerExchange Logger

For the PowerExchange Logger for Linux, UNIX, and Windows to extract change data from a remote source, the capture registrations for the source tables must specify **Part** for the **Condense** option.

**Note:** This requirement is not specific to remote logging. It also applies to PowerExchange Logger for Linux, UNIX, and Windows use on a source system.

If the capture registrations do not specify **Part** for the **Condense** option, you can edit the **Condense** setting. This change does not increment the registration version. You can continue to use the same registration and extraction map.

**Tip:** Do not add DTL\_BI or DTL\_CI columns to the extraction maps if you set the CAPT\_IMAGE parameter to AI in the pwxcl.cfg configuration file. With the AI setting, the PowerExchange Logger stores after images only. Consequently, you cannot use before images of the data in extraction processing. Also, CDC sessions that reference any CI fields fail.

## Configuring PowerCenter Connection Attributes for Extracting Data from the Log Files

For CDC sessions to extract change data from PowerExchange Logger for Linux, UNIX, and Windows log files on a system other than the source system, you must configure certain attributes on the PWX CDC Real Time connection.

The following table describes these connection attributes:

Connection Attribute	Value
Location	Enter the node name for the PowerExchange Listener that runs on the system where the PowerExchange Logger log files reside. If the log files are on the PowerCenter Integration Service machine, you can enter "local."
Map Location	Enter the node name for the location where the PowerExchange Listener on the source system stores the extraction maps. Usually, this node is the source system node.

Connection Attribute	Value
Map Location User and Map Location Password	<p>Enter a user ID and password that can access the extraction maps.</p> <p>If the PowerExchange Listener runs on a source system with PowerExchange security enabled, the user ID and password depends on the SECURITY statement settings in the DBMOVER configuration file.</p> <p>If the first parameter in the SECURITY statement is 2 and you are extracting z/OS data from the log files, enter a valid z/OS user ID and password in these fields. Also ensure that these z/OS user credentials have the following permissions:</p> <ul style="list-style-type: none"> <li>- READ access to the PowerExchange data set that is defined in the DTLCAMAP DD statement of the PowerExchange Listener JCL</li> <li>- READ access to CAPX.CND.* resource profiles in the FACILITY class, which are managed by your z/OS security product</li> </ul>
CAPI Connection Name Override	Enter the name of the CAPX CAPI_CONNECTION statement that is used by the PowerExchange Listener on the system where the PowerExchange Logger for Linux, UNIX, and Windows log files reside.

For more information about PWX CDC Real Time application connections, see *PowerExchange Interfaces for PowerCenter*.

## Example of Remote Logging from a z/OS Data Source

In this example, you use a PowerExchange Logger for Linux, UNIX, and Windows instance on a UNIX system to relog change data for DB2 for z/OS data sources. The system where the PowerExchange Logger runs is separate from the PowerCenter Integration Service system where you run the CDC sessions.

The PowerExchange Logger for MVS captures change data for registered DB2 for z/OS tables and logs that data to its log files on the z/OS system. The PowerExchange Logger for Linux, UNIX, and Windows reads data from the PowerExchange Logger for MVS log files and relogs that data on the UNIX system. PowerCenter CDC sessions then extract change data from the PowerExchange Logger for Linux, UNIX, and Windows log files rather than from the log files on the z/OS source system.

You need the PowerExchange Logger for Linux, UNIX, and Windows to read change data for registered tables in the DB2 instance DSN9 and then relog that data to its log files on the remote UNIX system. To do so, you must customize a PowerExchange Logger for Linux, UNIX, and Windows configuration file on the UNIX system and dbmover configuration files on both the z/OS and UNIX systems. Also, for the PowerCenter CDC sessions to extract change data from the PowerExchange Logger log files on UNIX, you must add NODE statements for the source and PowerExchange Logger systems to the dbmover configuration file on the Integration Service system and configure some PWXPC connection attributes.

First install PowerExchange on all three systems. You must run a PowerExchange Listener on the source system and on the PowerExchange Logger system. A PowerExchange Listener is not required on the PowerCenter Integration Service system.

1. On the z/OS source system, ensure that the DBMOVER member in the RUNLIB library includes the following CAPI\_CONNECTION statements:

```
LISTENER=(MVS02,TCPIP,2480)
/* UOW Cleanser
```



```

CAPI_CONNECTION=(NAME=MV2UOWC,TYPE=(UOWC,CAPINAME=M2_LRAP,RSTRADV=600,MEMCACHE=20480,
DATACLAS=UOWC))
/* Log Read API Connection
CAPI_CONNECTION=(NAME=MV2_LRAP,TYPE=(LRAP,LOG=MV2L,AGENT=MV2A))

```

2. On the UNIX system with the PowerExchange Logger for Linux, UNIX, and Windows log files, ensure that the dbmover configuration file includes the following statements:

```

/*
/* dbmover
/*
LISTENER=(unix1,TCPIP,2480)
NODE=(MVS02,TCPIP,prod mvs2,2480)
...
LOGPATH=/pwx/logs/mvscond
CAPT_XTRA=/pwx/capture/mvscond/camaps
CAPT_PATH=/pwx/capture/mvscond
/*
/* Source-specific CAPI Connection
CAPI_CONNECTION=(NAME=MV2UOWC,TYPE=(UOWC,CAPINAME=M2_LRAP,RSTRADV=600,MEMCACHE=20480)
)
CAPI_CONNECTION=(NAME=MV2_LRAP,TYPE=(LRAP,LOG=MV2L,AGENT=MV2A))
/*
/* CAPX CAPI Connection for continuous extraction
CAPI_CONNECTION=(NAME=CAPXDSN9,TYPE=(CAPX,DFLTINST=DSN9,FILEWAIT=60,RSTRADV=600))

```

**Note:** In the CAPX CAPI\_CONNECTION, the DFLTINST value is the name that is displayed in the **Instance** field for the registration group in the PowerExchange Navigator.

3. On the UNIX system with the PowerExchange Logger system log files, customize the PowerExchange Logger for Linux, UNIX, and Windows configuration file, pwxcl.cfg. For this example, include the following statements:

```

/*
/* pwxcl
/*
DBID=DSN9
DB_TYPE=DB2
CONN_OVR=MV2UOWC
CAPTURE_NODE=MVS02
PROMPT=Y
EXT_CAPT_MASK=/pwx/capture/mvscond/condense
COND_CDCT_RET P=50
LOGGER_DELETES_EXPIRED_CDCT_RECORDS=Y
COLL_END_LOG=0
NO_DATA_WAIT=0
NO_DATA_WAIT2=10
FILE_SWITCH_VAL=20000
FILE_SWITCH_CRIT=R
CAPT_IMAGE=BA

```

**Note:** The CAPTURE\_NODE parameter points to the source system node where the PowerExchange Listener processes capture requests.

4. Start the PowerExchange Listener and PowerExchange Logger for Linux, UNIX, and Windows on the UNIX system. Verify that the PowerExchange Listener is also running on the z/OS system.
5. On the PowerCenter Integration Service system, add the following NODE statements to the dbmover file:
  - A NODE statement that points to the PowerExchange Listener on the source system
  - A NODE statement that points to the PowerExchange Listener on the UNIX system with the PowerExchange Logger log files, if you not using a "local" connection

This example uses the following NODE statements in the dbmover file on the PowerCenter Integration Service machine:

```

NODE=(unix1,TCPIP,unix1,2480)
NODE=(MVS02,TCPIP,prod mvs2,2480)

```

6. Create a PowerCenter mapping, session, and workflow.

7. Configure a PWX DB2zOS CDC Real Time application connection for the CDC sessions that extract change data from the PowerExchange Logger log files on the UNIX system.

For this example, set the following connection attributes:

- For the **Location** attribute, enter unix2 to point to the node where the PowerExchange Logger for Linux, UNIX, and Windows log files reside. CDC sessions will read data from this location.
- For the **Map Location** attribute, enter MVS02 to point to the location of the extraction maps, which the z/OS source system node.
- For the **Map Location User** attribute, enter a valid user ID for the map location.
- For the **Map Location Password** attribute, enter the password for the map location user.
- For the **CAPI Connection Name** attribute, enter CAPXDSN9 to indicate the CAPX CAPI\_CONNECTION statement to use.

8. Cold start the CDC session.

The session begins extracting change data from the PowerExchange Logger log files on the UNIX system.

## Example of Remote Logging from a DB2 for i5/OS Data Source

In this example, you use a PowerExchange Logger for Linux, UNIX, and Windows instance on a UNIX system to capture change data from DB2 for i5/OS journals. The system where the PowerExchange Logger runs is separate from the PowerCenter Integration Service system where you run CDC sessions.

You need the PowerExchange Logger to capture change data for registered tables from DB2 journals in the DB2 instance PROD2 and then log that data to its log files on the remote UNIX system. To do so, you must customize a PowerExchange Logger configuration file on the UNIX system and dbmover configuration files on both the i5/OS and UNIX systems. Also, for the PowerCenter CDC sessions to extract change data from the PowerExchange Logger log files on UNIX, you must add NODE statements for the source and PowerExchange Logger systems to the dbmover configuration file on the Integration Service system and configure some PWXPC connection attributes.

First install PowerExchange on all three systems. You must run a PowerExchange Listener on the source system and on the PowerExchange Logger system. A PowerExchange Listener is not required on the PowerCenter Integration Service system.

1. On the i5/OS source system, ensure that the DBMOVER member in the *datalib*/CFG library includes the following CAPI\_CONNECTION statements:

```
LISTENER=(i50S1,TCPIP,2480)
/* UOW Cleanser CAPI Connection
CAPI_CONNECTION=(NAME=i5UOWC,TYPE=(UOWC,CAPINAME=i5_AS4J,RSTRADV=600,MEMCACHE=20480))
/* DB2 for i5/OS CAPI Connection
CAPI_CONNECTION=(NAME=i5_AS4J,TYPE=(AS4J,JOURNAL=PRODDATA/
PRODJRN,INST=PROD2,EOF=N,STOPIT=(CONT=5),LIBASUSER=Y))
```

**Note:** In the AS4J CAPI\_CONNECTION statement, the INST parameter value must match the **Instance** name that is displayed for the registration group in the PowerExchange Navigator.

2. On the UNIX system with the PowerExchange Logger log files, ensure that the dbmover configuration file includes the following statements:

```

/*
/* dbmover
/*
LISTENER=(unix2,TCPIP,2480)
NODE=(unix1,TCPIP,prod2,2480)
...
LOGPATH=/pwx/logs/i5oscond
CAPT_XTRA=/pwx/capture/i5oscond/camaps
CAPT_PATH=/pwx/capture/i5oscond
/*
/* Source-specific CAPI Connection
CAPI_CONNECTION=(NAME=i5UOWC,TYPE=(UOWC,CAPINAME=i5_AS4J,RSTRADV=600,MEMCACHE=20480))
CAPI_CONNECTION=(NAME=i5_AS4J,TYPE=(AS4J,JOURNAL=PRODDATA/
PRODURN,INST=PROD2,EOF=N,STOPIT=(CONT=5),LIBASUSER=Y))
/*
/* CAPX CAPI Connection for continuous extraction
CAPI_CONNECTION=(NAME=CAPXPROD,TYPE=(CAPX,DFLTINST=PROD2,FILEWAIT=60,RSTRADV=600))

```

**Note:** In the CAPX CAPI\_CONNECTION, the DFLTINST value is the name that is displayed in the **Instance** field for the registration group in the PowerExchange Navigator.

3. On the UNIX system with the PowerExchange Logger system log files, customize the PowerExchange Logger configuration file, pwxccl.cfg. For this example, include the following statements:

```

/*
/* pwxccl
/*
DBID=PROD2
DB_TYPE=AS4
CONN_OVR=i5UOWC
CAPTURE_NODE=i5OS1
CAPTURE_NODE_UID=db2user
CAPTURE_NODE_EPWD=encrypted_password
PROMPT=Y
EXT_CAPT_MASK=/pwx/capture/i5oscond/condense
COND_CDCT_RET_P=50
LOGGER_DELETES_EXPIRED_CDCT_RECORDS=Y
COLL_END_LOG=0
NO_DATA_WAIT=0
NO_DATA_WAIT2=10
FILE_SWITCH_VAL=20000
FILE_SWITCH_CRIT=R
CAPT_IMAGE=BA

```

**Note:** The CAPTURE\_NODE parameter points to the source system node where the PowerExchange Listener processes capture requests.

4. Start the PowerExchange Listener and PowerExchange Logger for Linux, UNIX, and Windows on the UNIX system. Verify that the PowerExchange Listener is also running on the i5/OS system.
5. On the PowerCenter Integration Service system, add the following NODE statements to the dbmover file:
  - A NODE statement that points to the PowerExchange Listener on the source system
  - A NODE statement that points to the UNIX system with the PowerExchange Logger log files

This example uses the following NODE statements in the dbmover file on the PowerCenter Integration Service system:

```

NODE=(i5OS1,TCPIP,i5OS1,2480)
NODE=(unix2,TCPIP,prod2,2480)

```

6. Create a PowerCenter mapping, session, and workflow.
7. Configure a PWX DB2i5OS CDC Real Time application connection for CDC sessions that extract change data from the PowerExchange Logger log files on the UNIX system.

For this example, set the following connection attributes:

- For the **Location** attribute, enter unix2 to point to the node where the PowerExchange Logger for Linux, UNIX, and Windows log files reside. CDC sessions will read data from this location.
- For the **Map Location** attribute, enter i5OS1 to point to the location of the extraction maps, which is the source system node.
- For the **Map Location User** attribute, enter a valid user ID for the map location.
- For the **Map Location Password** attribute, enter the password for the map location user.
- For the **CAPI Connection Name** attribute, enter CAPXPROD to indicate the CAPX CAPI\_CONNECTION statement to use.

8. Cold start the CDC session.

The session begins extracting change data from the PowerExchange Logger log files on the UNIX system.

## Example of Remote Logging from an Oracle CDC with LogMiner Data Source

In this example, you use a PowerExchange Logger for Linux, UNIX, and Windows instance on a UNIX system to capture change data from an Oracle instance on another UNIX system. The system where the PowerExchange Logger runs is separate from the PowerCenter Integration Service system where you run CDC sessions.

You need the PowerExchange Logger to capture change data from registered tables in the Oracle instance PRODORA on the source system and then log that data to its log files on the remote UNIX system. To do so, you must customize a PowerExchange Logger configuration file on the remote system and dbmover configuration files on both systems. Also, for the PowerCenter CDC sessions to extract change data from the PowerExchange Logger log files, you must add NODE statements for the source and PowerExchange Logger systems to the dbmover configuration file on the Integration Service system and configure some PWXPC connection attributes.

First install PowerExchange on all three systems. You must run a PowerExchange Listener on the source system and on the PowerExchange Logger system. A PowerExchange Listener is not required on the PowerCenter Integration Service system.

1. On the Oracle source system, ensure that the dbmover configuration file includes the following statements:

```
LISTENER=(unix1,TCPIP,2480)
/* UOW Cleanser
CAPI_CONNECTION=(NAME=UOWCORA,TYPE=(UOWC,CAPINAME=CAPIORA,RSTRADV=600))
/* Oracle Connection
CAPI_CONNECTION=(NAME=CAPIORA,TYPE=(ORCL,catint=120,ORACOLL=PRODORA))
ORACLEID=(PRODORA,ORAINST2,ORAINST2,ORAINST2)
```

**Note:** The value of the first positional parameter, *collection\_id*, in the ORACLEID statement must match the value of the DBID parameter in the PowerExchange Logger configuration file on the PowerCenter Integration Service system.

2. On the PowerExchange Logger system with the log files, ensure that the dbmover configuration file includes the following statements:

```
/*
/* dbmover
```

```

/*
LISTENER=(unix2,TCPIP,2480)
NODE=(unix1,TCPIP,prodora,2480)
...
LOGPATH=/pwx/logs/oracond
CAPT_XTRA=/pwx/capture/oracond/camaps
CAPT_PATH=/pwx/capture/oracond
/*
/* Source-specific CAPI Connection
CAPI_CONNECTION=(NAME=UOWCORA,TYPE=(UOWC,CAPINAME=CAPIORA,RSTRADV=600))
CAPI_CONNECTION=(NAME=CAPIORA,TYPE=(ORCL,CATINT=120,ORACOLL=PRODORA))
/*
/* CAPX CAPI Connection for continuous extraction
CAPI_CONNECTION=(NAME=CAPXORA,TYPE=(CAPX,DFLTINST=PRODORA,FILEWAIT=60,RSTRADV=600))

```

3. On the PowerExchange Logger system with the log files, customize the PowerExchange Logger configuration file, pwxcl.cfg. For this example, include the following statements:

```

/*
/* pwxcl
/*
DBID=PRODORA
DB_TYPE=ORA
CONN_OVR=UOWCORA
CAPTURE_NODE=unix1
CAPTURE_NODE_UID=orauser
CAPTURE_NODE_EPWD=encrypted_password
PROMPT=Y
EXT_CAPT_MASK=/capture/condense0
COND_CDCT_RET P=60
LOGGER_DELETES_EXPIRED_CDCT_RECORDS=Y
COLL_END_LOG=0
NO_DATA_WAIT=0
NO_DATA_WAIT2=10
FILE_SWITCH_VAL=20
FILE_SWITCH_CRIT=M
CAPT_IMAGE=BA

```

**Note:** The CAPTURE\_NODE parameter points to the source system node where the PowerExchange Listener processes capture requests.

4. Start the PowerExchange Listener and PowerExchange Logger on the remote system. Verify that the PowerExchange Listener is also running on the source system.
5. On the PowerCenter Integration Service system, add the following NODE statements to the dbmover configuration file:
  - A NODE statement that points to the PowerExchange Listener on the source system.
  - A NODE statement that points to the PowerExchange Logger system.

This example uses the following NODE statements in the dbmover file on the PowerCenter Integration Service system:

```

NODE=(unix1,TCPIP,unix1,2480)
NODE=(unix2,TCPIP,unix2,2480)

```

6. Create a PowerCenter mapping, session, and workflow.
7. Configure a PWX Oracle CDC Real Time application connection for CDC sessions that extract change data from the PowerExchange Logger log files.

For this example, set the following connection attributes:

- For the **Location** attribute, enter unix2 to point to the node where the PowerExchange Logger log files reside. CDC sessions will read data from this location.
- For the **Map Location** attribute, enter unix1 to point to the location of the extraction maps, which is the source system node.

- For the **Map Location User** attribute, enter a valid Oracle user ID.
  - For the **Map Location Password** attribute, enter the password for the Oracle user ID.
  - For the **CAPI Connection Name** attribute, enter CAPXORA to indicate the CAPX CAPI\_CONNECTION statement to use.
8. Cold start the CDC session.
- The session begins extracting change data from the PowerExchange Logger log files.

# Part IV: Change Data Extraction

This part contains the following chapters:

- [Introduction to Change Data Extraction, 208](#)
- [Extracting Change Data, 220](#)
- [Managing Change Data Extractions, 246](#)

## CHAPTER 9

# Introduction to Change Data Extraction

This chapter includes the following topics:

- [Change Data Extraction Overview, 208](#)
- [Extraction Modes, 209](#)
- [PowerExchange-Generated Columns in Extraction Maps, 210](#)
- [Uses of BI and CI Fields in Extraction Maps, 214](#)
- [Restart Tokens and the Restart Token File, 215](#)
- [Multiple-Source Processing in CDC Sessions, 216](#)
- [Commit Processing with PWXPC, 218](#)
- [Tuning Options, 219](#)

## Change Data Extraction Overview

PowerExchange works in conjunction with PWXPC and PowerCenter to extract captured change data and write it to one or more targets. Learn key concepts about extraction processing so that you can configure CDC sessions for efficient extraction of data and proper restart and recovery.

To extract change data that PowerExchange captured, import the metadata for the capture source into PowerCenter Designer. Use one of the following methods:

- For relational data sources, import either the extraction maps from PowerExchange or the source metadata from the database. If you import source metadata, you might need to modify the source definition in Designer to add PowerExchange-defined CDC columns or to remove any columns that are not included in the extraction map. If you import extraction maps, you do not need to manually add or remove these columns from the PowerCenter source definition.
- For nonrelational data sources, import the extraction maps from PowerExchange.

After you import the metadata, you can use the source definitions in PowerCenter to create mappings, sessions, and workflows for extracting change data from PowerExchange.

### RELATED TOPICS:

- [“Extraction Modes” on page 209](#)
- [“Restart Tokens and the Restart Token File” on page 215](#)
- [“Multiple-Source Processing in CDC Sessions” on page 216](#)



- [“Commit Processing with PWXPC” on page 218](#)
- [“Tuning Options” on page 219](#)

## Extraction Modes

You can extract the change data that PowerExchange captured in near real time or as a batch process.

You indicate the extraction mode by setting the PowerCenter connection type and certain PowerExchange CDC configuration parameters. Some extraction modes are available only if you use PowerExchange Condense or the PowerExchange Logger for Linux, UNIX, and Windows.

Based on your extraction requirements, use one of the following extractions modes:

### Real-time extraction mode

Continuously extracts change data in near real time from the change stream. Extraction processing continues until the CDC session stops or is interrupted.

To implement this mode, configure a PWX CDC Real Time application connection in PowerCenter for your data source type.

### Batch extraction mode

Extracts change data from PowerExchange Condense condense files on z/OS or i5/OS, or from PowerExchange Logger for Linux, UNIX, and Windows log files. Data is extracted only from the files that are closed at the time the CDC session runs. The CDC session ends after it completes processing the files.

To implement this mode, configure the following items:

- In the PowerExchange Navigator, set the **Condense** option to **Part** or **Full** in the capture registrations.
- In PowerCenter, configure a PWX CDC Change application connection for your data source type.

### Continuous extraction mode.

Continuously extracts change data from open and closed PowerExchange Logger for Linux, UNIX, and Windows log files in near real time.

For z/OS or i5/OS data sources, this extraction mode is available only if you log data to a remote PowerExchange Logger for Linux, UNIX, and Windows on another system.

To implement this mode, configure the following items:

- In the PowerExchange Navigator, set the **Condense** option to **Part** in the capture registrations.
- In PowerCenter, configure a PWX CDC Real Time application connection for your data source type.
- Configure a CAPX CAPI\_CONNECTION statement in the DBMOVER configuration file.
- If you remote logging of data from z/OS or i5/OS data sources to a PowerExchange Logger for Linux, UNIX, and Windows, configure the remote PowerExchange Logger to log change data from the source system.

# PowerExchange-Generated Columns in Extraction Maps

Besides the table columns that are defined in capture registrations, extraction maps include columns that PowerExchange generates.

These PowerExchange-generated columns contain CDC-related information, such as the type of SQL change and time stamp.

When you import an extraction map in Designer, PWXPC includes the PowerExchange-generated columns in the source definition.

When you run a database row test on an extraction map, the PowerExchange Navigator displays the PowerExchange-generated columns in the results. By default, the PowerExchange Navigator hides these columns from view when you open the extraction map. To display these columns, open the extraction map, right-click anywhere within the **Extract Definition** window, and select **Show Auto Generated Columns**.

**Note:** By default, all columns except the DTL\_\_columnname\_CNT and DTL\_\_columnname\_IND columns are selected in an extraction map. To select these columns, you must edit the extraction map.

The following table describes the columns that PowerExchange generates for each change record:

Column	Description	Datatype	Length
DTL__CAPXRESTART1	<p>A binary value that represents the position of the end of the UOW for that change record followed by the position of the change record itself.</p> <p>The length of a sequence token varies by data source type, except on z/OS where sequence tokens for all data source types have the same length.</p> <p>The value of DTL__CAPXRESTART1 is also known as the <i>sequence token</i>, which when combined with the <i>restart token</i> comprises the restart token pair.</p> <p>A sequence token for a change record is a strictly ascending and repeatable value.</p>	VARBIN	255
DTL__CAPXRESTART2	<p>A binary value that represents a position in the change stream that can be used to reconstruct the UOW state for the change record, with the following exceptions:</p> <ul style="list-style-type: none"><li>- Microsoft SQL Server CDC. A binary value that contains the DBID of the distribution database and the name of the distribution server.</li><li>- Change data extracted from full condense files on z/OS or i5/OS. A binary value that contains the instance name from the registration group of the capture registration.</li></ul> <p>The length of a restart token varies by data source type. On z/OS, restart tokens for all data source types have the same length, except for change data extracted from full condense files.</p> <p>The value of DTL__CAPXRESTART2 is also known as the <i>restart token</i>, which when combined with the <i>sequence token</i> comprises the restart token pair.</p>	VARBIN	255

Column	Description	Datatype	Length
DTL__CAPXROWID	<p>For PowerExchange Oracle CDC with LogMiner and Express CDC for Oracle, provides the physical rowid value. PowerExchange can include rowid values in change records for Oracle tables only if the tables do not have row movement enabled.</p> <p>To enable the capture of rowid values, you must configure one of the following parameters:</p> <ul style="list-style-type: none"> <li>- For PowerExchange Oracle CDC with LogMiner, set the ROWID parameter in the ORCL CAPI_CONNECTION statement to Y or ALLOW.</li> <li>- For PowerExchange Express CDC for Oracle, include the OPTIONS ROWID=Y statement in the Express CDC configuration file.</li> </ul> <p>The rowid is useful for processing rows in unkeyed tables during CDC extraction sessions.</p>	CHAR	18
DTL__CAPXRRN	For DB2 on i5/OS only, the relative record number.	DECIMAL	10
DTL__CAPXUOW	A binary value that represents the position in the change stream of the start of the UOW for the change record.	VARBIN	255
DTL__CAPXUSER	<p>The user ID of the user who made the change to the data source, with the following exceptions:</p> <ul style="list-style-type: none"> <li>- For Adabas 8.3 CDC sources, this value is the Security User-id (SECUID) of the user if the Adabas File Definition includes the system field SY=SECUID.</li> <li>- For Datacom table-based CDC sources, this value is the MUF name.</li> <li>- For DB2 for i5/OS sources, this value depends on the LIBASUSER parameter in the AS4J CAPI_CONNECTION statement. If LIBASUSER=Y, this value is the library name and file name of the file where the change was made. If LIBASUSER=M, this value is the library name, file name, and data member name of the file where the change was made. If LIBASUSER=N, this value is the user ID of the user who made the change.</li> <li>- For DB2 for z/OS sources, this value depends on the UIDFMT parameter in the LRAP CAPI_CONNECTION. Depending on the parameter setting, this value can be a DB2 connection identifier, correlation identifier, connection type, plan name, user ID, or all of these values in the format UID:PLAN:CORR:CONN:CTYPE. If you do not specify the UIDFMT parameter, this value is the user ID of the user who made the change.</li> <li>- For IDMS sources, this value is the value that the user program puts in the program name field of the application subschema control block. Usually, this value is the user program name.</li> <li>- For Microsoft SQL Server sources, this value depends on the UIDFMT parameter in the MSQL CAPI_CONNECTION statement. If UIDFMT=DBNAME, this value is the SQL Server publication database name. If UIDFMT=NONE, this value is a null.</li> <li>- For Oracle sources, this value is a user ID that PowerExchange gets from Oracle, if available. Otherwise, this value is null. This information applies to both PowerExchange Oracle CDC with LogMiner and PowerExchange Express CDC for Oracle.</li> </ul>	VARCHAR	255

Column	Description	Datatype	Length
DTL__CAPXTIMESTAMP	<p>The time stamp that the source DBMS records for a change on the source database.</p> <p>This value can be either the time stamp that the source DBMS writes to the change record in the database logs or the time stamp of the transaction commit on the source database.</p> <p>The time stamp type depends on the source type and certain parameters:</p> <ul style="list-style-type: none"> <li>- For DB2 for Linux, UNIX, and Windows sources, the transaction commit time stamp.</li> <li>- For Microsoft SQL Server sources, the time at which the change was written to the distribution database.</li> <li>- For PowerExchange Express CDC for Oracle sources, the time stamp type is controlled by the TIME_STAMP_MODE parameter in the OPTIONS statement of the Express CDC configuration file.</li> <li>- For all sources that require a UOWC CAPI_CONNECTION statement, the time stamp type is controlled by the TIMESTAMP parameter in the UOWC CAPI_CONNECTION statement in the DBMOVER file.</li> </ul> <p>For more detailed information about time stamps for each source type, see <a href="#">Appendix A, "DTL__CAPXTIMESTAMP Time Stamps" on page 279</a>.</p> <p>The time stamp format is:</p> <p>YYYYMMDDhhmmssnnnnnn</p> <p>Where:</p> <ul style="list-style-type: none"> <li>- YYYY is the four-digit year.</li> <li>- MM is the month.</li> <li>- DD is the day.</li> <li>- hhmmssnnnnnn is hours, minutes, seconds, and microseconds.</li> </ul> <p><b>Note:</b> DB2 for Linux, UNIX, and Windows and Oracle do not support microseconds in the time stamp.</p>	CHAR	20

Column	Description	Datatype	Length
DTL__CAPXACTION	<p>A single character that indicates the type of change record that PowerExchange passes to the target during extraction processing. A DTL__CAPXACTION value corresponds to the type of change operation on the source database.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> <li>- I. Insert.</li> <li>- D. Delete.</li> <li>- U. After image of an update.</li> <li>- T. Before image of an update. (ODBC connections only)</li> </ul> <p>If you specify an <b>Image Type</b> of <b>BA</b> on the connection for a CDC session, PowerExchange generates a delete record followed by an insert record for a source update. In the delete record, the DTL__CAPXACTION column contains the value D. In the insert record, the DTL__CAPXACTION column contains the value I.</p> <p>If you specify an <b>Image Type</b> of <b>AI</b> on the connection for a CDC session, PowerExchange generates one record for an update. In this record, the DTL__CAPXACTION column contains the U value.</p> <p>If you use an ODBC connection to write change data to a staging table and either set the ODBC driver CAPXIMAGETYPE parameter to TU or enter the SQL escape sequence DTLIMTYPE=TU in PowerCenter, this column can contain a value of T or U. For each source update, PowerExchange delivers two records to the staging table: one for the before image and another for the after image. In the before image record, the DTL__CAPXACTION column contains the T value. In the after image record, The DTL__CAPXACTION column contains the U value.</p>	CHAR	1
DTL__CAPXCASDELIND	<p>For DB2 for z/OS sources only, a single character that indicates whether DB2 has deleted the row because the table specifies the ON DELETE CASCADE clause. Valid values are:</p> <ul style="list-style-type: none"> <li>- Y. Indicates that DB2 deleted this row because of a cascade delete rule.</li> <li>- N. Indicates that DB2 did not delete this row because of a cascade delete rule.</li> </ul>	CHAR	1
DTL__BI_<columnname>	For UPDATE operations, the value of the before image of the selected column in the change record.	Datatype of the source column	Length of the source column
DTL__CI_<columnname>	<p>For UPDATE operations, a single character that indicates whether the selected column was changed. Valid values are:</p> <ul style="list-style-type: none"> <li>- Y. Indicates that the column changed.</li> <li>- N. Indicates that the column did not changed.</li> <li>- Null value. Indicates an INSERT or DELETE operation.</li> </ul>	CHAR	1

Column	Description	Datatype	Length
DTL__columnname_CNT	Binary count column. PowerExchange generates this column for variable length columns of types VARCHAR and VARBIN to determine the length of the column during change data extraction processing. <b>Note:</b> By default, binary count columns are not selected in an extraction map. You must edit an extraction map to select these columns.	NUM32U	0
DTL__columnname_IND	Null indicator column. PowerExchange generates this column for nullable columns to indicate the nullable value for the column. <b>Note:</b> By default, null indicator columns are not selected in an extraction map. You must edit an extraction map to select these columns.	BIN	1

## Uses of BI and CI Fields in Extraction Maps

PowerExchange captures both before images and after images of data for all SQL UPDATE operations on source columns. To access before image data to process change data in some way during CDC sessions, add before image (BI) and change indicator (CI) fields to extraction maps.

For example, you can use the BI and CI fields for the following purposes:

- To filter captured data for extraction and apply processing.
- To update primary keys on the target based on whether primary keys on the source changed.

### Case 1. Filtering Change Data for Extraction and Apply Processing

If you add CI fields for one or more data columns in an extraction map, PowerExchange compares before and after images of the data captured for these columns. If an UPDATE occurred, PowerExchange sets the generated DTL\_\_CI\_column\_name value to Y.

You can use a DTL\_\_CI\_column\_name in WHERE clause filters for CDC sessions to filter the change stream during extraction processing. In PowerCenter, define the filters in the **Filter Override** attribute of the session properties. By using these filters, you can reduce the amount of data that PowerCenter processes.

During extraction processing, PWXPC creates SQL SELECT statements that include the WHERE clause filters. PWXPC passes these statements to PowerExchange. PowerExchange selects and returns the data that matches the WHERE conditions. PWXPC then makes this data available to the CDC sessions. Additional manipulation of the data might occur in PowerCenter, based on how you define the mappings.

To filter change data for extraction and apply processing:

1. In the PowerExchange Navigator, edit the extraction map that you plan to import as the source definition for the CDC session. For each column that you want to filter on, add a CI field.  
PowerExchange generates CI fields that have names in the format DTL\_\_CI\_column\_name.  
For more information about adding CI fields to extraction maps, see the *PowerExchange Navigator User Guide*.

2. In PowerCenter, define WHERE clause filters in the **Filter Override** attribute of the CDC session properties.  
For the filters, enter `DTL__CI_column_name` conditions. For example, enter `DTL__CI_ACCOUNT='Y'`, where 'Y' indicates an Update occurred.

For more information about filter overrides on CDC sessions, see *PowerExchange Interfaces for PowerCenter*.

When the CDC session runs, PWXPC provides only the change data that matches the WHERE filter to PowerCenter for extraction and apply processing.

**Note:** Using many filters with CI fields might noticeably increase CPU overhead.

### Case 2. Updating Primary Key Fields on the Target

If the target primary key does not match the source primary key, or if the source database allows updates to primary key fields, CDC sessions cannot apply updates to target keys based on after image data only.

To prevent this problem, you can select the **BA** option for the **Image Type** attribute on PWX CDC application connections. This option causes PWXPC to generate two transactions for each source UPDATE: a DELETE followed by an INSERT. The DELETE deletes the old row based on the before image. The INSERT inserts a row based on the after image.

Alternatively, to avoid the overhead of generating two transactions for every source UPDATE, select the **AI** option for the **Image Type** attribute. Also use CI and BI columns in combination with a PowerCenter Flexible Target Key Custom transformation. With this configuration, PowerCenter generates an INSERT or UPDATE transaction only when a source UPDATE results in changes to primary key fields on the target. Complete the following steps to implement this solution.

To update primary key fields on the target using BI and CI fields:

1. In the PowerExchange Navigator, edit the extraction map that you plan to import as the source definition for the CDC session. Add both BI and CI fields for one or more of the primary key columns on the source.
2. Verify that the **Image Type** attribute on the PWX CDC application connection for the CDC session is **AI**. This setting causes PWXPC to pass Updates to the CDC session as Updates. Because you added BI and CI fields for key columns in the extraction map, Update rows for these columns include both before and after images.
3. In PowerCenter, define a Flexible Target Key Custom transformation.  
The transformation uses the `DTL__CI` indicator for the source key columns to detect when Updates to primary key columns on the target are needed.
4. Add the transformation to the mapping for the CDC session.

For more information about Flexible Target Key Custom transformations, see *PowerExchange Interfaces for PowerCenter*.

## Restart Tokens and the Restart Token File

PowerExchange uses a pair of token values, called a restart token pair, to determine where to begin extracting change data in the change stream for each source in a CDC session. A restart token pair matches the position of a specific change record in the change stream.

You can specify restart token pairs in the restart token file. PWXPC also stores restart tokens for CDC sessions that have run in a state table or file. The token values in the restart token file override those in the state table or file.

Specify restart tokens in the restart token file in the following situations:

- For a new CDC session, specify restart token pairs for the sources in the session. You can define a unique restart token pair for each source, or use the special override statement to specify a restart token pair that pertains to all or multiple data sources. The restart tokens should represent the point-in-time in the change stream when you materialized the corresponding targets.
- If you add a data source to a CDC session, specify a restart token pair for that source.
- If you need to override token values for one or more data sources in a CDC session, use override statements in the restart token file.

A restart token pair is composed of the following token types:

#### Sequence token

A binary value that represents, for each change record that is read, the change stream position of the end of the UOW followed by the position of the change record. A sequence token is a strictly ascending and repeatable value.

#### Restart token

A binary value that represents, for each change record that is read, a change stream position that PowerExchange can use to reconstruct the UOW state for the change record.

In some cases, the restart token might contain the position of the oldest open UOW. An open UOW is a UOW for which PowerExchange has read the beginning of the UOW from the change stream but has not yet read the commit record, or end-UOW.

When a CDC session runs, PWXPC reads the token values for each source from the state table or file and also reads the restart token file. PowerExchange uses the appropriate restart token values to determine the point from which to start reading change data from the change stream for each source in the CDC session. After determining the start point, PowerExchange starts reading and passing change data to PWXPC. PWXPC uses the sequence token for a source to determine the point at which to start providing the change data for the source.

#### RELATED TOPICS:

- [“Recovery and Restart Processing for CDC Sessions” on page 235](#)
- [“Configuring the Restart Token File” on page 242](#)
- [“Creating Restart Tokens for Extractions” on page 241](#)

## Multiple-Source Processing in CDC Sessions

When you use PWX CDC application connections to extract change data, PowerExchange reads the change stream in a single pass for all source definitions in the mapping. The sources must be of the same type and use the same change stream.

To create source definitions in Designer, import source metadata in one of the following ways:

- Import a PowerExchange extraction map by using the **Import from PowerExchange** dialog box.
- Import table definitions from a relational database by using the **Import from PowerExchange** dialog box or the **Import from Database** dialog box.



Informatica recommends that you import extraction maps. It makes creating mappings and sessions easier for the following reasons:

- The source definition contains the extraction map name. You do not need to provide this name when you configure the session.
- The source definition contains the PowerExchange-generated CDC columns, such as the DTL\_\_CAPX columns. You do not need to add these columns to the source definition.


During change data extraction, PowerExchange processes all source definitions in the mapping that have the same source type. Do not include multiple data source types in the mapping. Otherwise, the CDC session fails with message PWXPC\_10080.

For example, do not run a CDC session that contains a mapping with both DB2 and Oracle source definitions, even if changes for those sources are in the same change stream. Instead, create a mapping and session for all of the DB2 sources and a separate, unique mapping and session for all of the Oracle sources.

PowerExchange reads the change stream twice: once for the session with the DB2 sources, and once for the session with the Oracle sources.

The following figure shows an example mapping in PowerCenter Designer that includes three DB2 sources:

*Mapping Designer*



If you include this mapping in a session that uses a PWX DB2LUW CDC application connection, PowerExchange reads the change stream and extracts changes for all three source tables in a single pass. PowerExchange extracts change data in chronological order, based on when the UOWs completed. PowerExchange passes the change data to PWXPC, and PWXPC provides the changes to the appropriate source qualifier.

If you create a workflow that contains multiple CDC sessions, PowerExchange uses a connection for each session, even if the sessions extract change data from the same change stream. For example, the sessions might extract change data from the same PowerExchange Logger log files.

**Note:** Because the example mapping uses source definitions created from extraction maps, it cannot be used for bulk data movement operations. However, mappings that use source definitions created from database relational metadata can be used for either change data extraction or bulk data movement.

# Commit Processing with PWXPC

The PowerCenter Integration Service, in conjunction with PWXPC, commits data to the target based on the **Commit Type** session property and the commitment control attributes specified on PWX CDC Change or Real Time application connections.

By default, the **Commit Type** session property specifies **Target** for target-based commit processing. However, the PowerCenter Integration Service always uses source-based commit processing for CDC sessions. Change the commit type to **Source**. If you retain the default value and run a CDC session, the PowerCenter Integration Service automatically uses source-based commit processing and writes message WRT\_8226 in the session log. You do not need to set the **Commit Interval** session property because PWXPC ignores it.

To control when commits occur, configure commitment control attributes on the PWX CDC Change and Real Time application connections.

The following table describes these connection attributes:

Connection Attribute	PWX Real Time or Change Connections	Description
Maximum Rows Per commit	Both	Maximum number of change records that PWXPC processes before it flushes the data buffer to commit the change data to the targets. If necessary, PWXPC continues to process change records across UOW boundaries until this maximum rows limit is met. PWXPC does not wait for a UOW boundary to commit the change data. Default is 0, which causes PWXPC to not use this maximum rows limit.
Minimum Rows Per commit	Real Time	Minimum number of change records that PowerExchange reads from the change stream before it passes any commit records in the change stream to PWXPC. Before reaching this minimum, PowerExchange skips commit records and passes only the change records to PWXPC. Default is 0, which causes PowerExchange to not use this minimum rows limit.
Real-time Flush Latency in milli-seconds	Real Time	Number of milliseconds that must elapse before PWXPC flushes the data buffer to commit change data to the targets. When this latency period expires, PWXPC continues to read the changes in the current UOW until it reaches the end of the UOW. Then, PWXPC flushes the data buffer to commit the change data to the targets. Default is 0, which causes PWXPC to use 2,000 milliseconds.
UOW Count	Both	Number of UOWs that PWXPC must process before flushing the data buffer to commit the change data to the targets. Default is 1.

PWXPC flushes the data buffer to commit change data to the targets when one of the following thresholds is met, whichever one is first:

- **Maximum Rows Per commit**
- **Real-Time Flush Latency in milli-seconds**
- **UOW Count**

If you specify **Minimum Rows Per commit**, this threshold must also be met before a commit occur.

After PWXPC commits the change data, it resets the UOW count, the maximum and minimum rows per commit, and the real-time flush latency timer. PWXPC continues to read change data. Whenever one of the commitment control thresholds is met, PWXPC commits change data to the targets. Commit processing continues until the CDC session is stopped, ends, or terminates abnormally. When the PWXPC CDC reader ends normally, PWXPC issues a final commit to flush all complete, buffered UOWs and their final restart tokens to the targets. Prior to ending, the PWXPC CDC reader writes the following message to the session log:

```
PWXPC_12075 [INFO] [CDCRestart] Session complete. Next session will restart at: Restart
1 [restart1_token] : Restart 2 [restart2_token]
```

## RELATED TOPICS:

- [“Commitment Control Attributes” on page 230](#)
- [“Examples of Controlling Commit Processing” on page 233](#)

# Tuning Options

PowerExchange provides flexible tuning options that you can use to reduce CPU usage on a source system that has constrained CPU resources. These options can also potentially improve throughput for CDC sessions.

The tuning options move some extraction processing to another machine such as the PowerCenter Integration Service machine. If the machine to which processing is offloaded has sufficient resources, the performance of CDC sessions might improve.

The following tuning options can help you take maximum advantage of the system resources that are available and maximize throughput for CDC sessions:

- **Offload processing.** Use offload processing to transfer column-level extraction processing from the PowerExchange Listener on the source system to the PowerExchange client on the PowerCenter Integration Service machine. Also, if the data source type requires use of the UOW Cleanser (UOWC), offloading transfers UOWC processing to the Integration Service machine. Use offloading to help increase throughput when resources available for the PowerExchange Listener are constrained on the source system.
- **Remote logging of change data.** Configure a PowerExchange Logger for Linux, UNIX, and Windows instance on a system other than the source system. The PowerExchange Logger reads change data from the source and writes the data to its local log files. CDC sessions extract the change data from the PowerExchange Logger log files. This configuration moves resource-intensive, column-level processing from the source system to the PowerExchange Logger system. Use remote logging to help improve throughput for CDC sessions when resources on the source system are constrained.
- **Multithreading.** Enable the use of multiple worker threads for resource-intensive, column-level extraction processing. You can use multithreading on the source system to process data from Linux, UNIX, or Windows data sources, or on another system where the extraction processing runs. Enable multithreading only if extractions appear to be CPU bound. You can use multithreading with the offloading feature or remote logging.

## RELATED TOPICS:

- [“CDC Offload Processing” on page 275](#)
- [“Multithreaded Processing” on page 277](#)
- [“Remote Logging Overview” on page 188](#)

## CHAPTER 10

# Extracting Change Data

This chapter includes the following topics:

- [Overview of Extracting Change Data, 220](#)
- [Security Considerations for Extracting z/OS Data, 221](#)
- [Task Flow for Extracting Change Data, 222](#)
- [Testing an Extraction Map, 222](#)
- [Configuring PowerCenter CDC Sessions, 224](#)
- [Recovery and Restart Processing for CDC Sessions, 235](#)
- [Creating Restart Tokens for Extractions, 241](#)
- [Displaying Restart Tokens, 241](#)
- [Configuring the Restart Token File, 242](#)

## Overview of Extracting Change Data

Use PowerExchange in conjunction with PWXPC and PowerCenter to extract captured change data and write the data to one or more targets.

To extract the change data that PowerExchange captures, in Designer, import metadata for the CDC sources and targets and create a mapping. Then, in Workflow Manager, create an application connection, a session, and a workflow. You can create multiple mappings, sessions, and workflows based on the same source and target definitions, if appropriate.

For relational data sources, you can import the metadata from either database definitions or PowerExchange extraction maps. For nonrelational sources, you must import the metadata from PowerExchange extraction maps.

**Tip:** Informatica recommends that you import the metadata from PowerExchange extraction maps. When you use extraction maps, the source definitions contain all of the PowerExchange-generated CDC columns, including any before image (BI) and change indicator (CI) columns you added. Also, you do not need to specify the extraction map name for each source in the session properties because PWXPC can derive the extraction map name from the source definition.

Before starting a CDC session for the first time, create restart tokens to define the extraction start point in the change stream. You might also need to create restart tokens to resume extraction processing in a recovery scenario.

Optionally, you can configure event table processing to stop a CDC session that uses real-time extraction mode based on user-defined events.

Also, you can use the following tuning options to help take maximum advantage of the available system resources and maximize throughput for CDC sessions:

- Offload processing. Use offload processing to transfer column-level extraction processing from the PowerExchange Listener on the source system to the PowerExchange client on the PowerCenter Integration Service machine.
- Remote logging of change data. Configure a PowerExchange Logger for Linux, UNIX, and Windows instance on a system other than the source system. The PowerExchange Logger reads change data from the source and logs it in the PowerExchange Logger log files on the other system. CDC sessions can then extract change data from the PowerExchange Logger log files.
- Multithreading. Enable the use of multiple worker threads to use multithreading for resource-intensive, column-level extraction processing. You can use multithreading on the source system if you are processing data from Linux, UNIX, or Windows data sources, or on another system where the extraction processing runs.

#### RELATED TOPICS:

- [“Task Flow for Extracting Change Data” on page 222](#)
- [“Configuring PowerCenter CDC Sessions” on page 224](#)
- [“Recovery and Restart Processing for CDC Sessions” on page 235](#)
- [“Creating Restart Tokens for Extractions” on page 241](#)
- [“Configuring the Restart Token File” on page 242](#)

## Security Considerations for Extracting z/OS Data

For the highest level of security for data from z/OS data sources, set the SECURITY option to 2 in the z/OS DBMOVER configuration member where the extraction maps are located. With this setting, a PowerCenter CDC session is permitted to extract data only if its user credentials pass z/OS security checking.

When you define a PWXPC connection for CDC sessions that extract data for a z/OS source from the z/OS system, you must enter a valid z/OS user ID and password in the **User Name** and **Password** connection attributes.

If you log captured z/OS data to remote PowerExchange Logger for Linux, UNIX, and Windows log files, when defining a PWXPC connection for CDC sessions that extract data from these log files, enter the z/OS user ID and password in the **Map Location User** and **Map Location Password** connection attributes. If the location of the PowerExchange Logger log files is not local, enter the z/OS user ID and password in the **User Name** and **Password** connection attributes for use by the PowerExchange Listener on the Linux, UNIX, or Windows system where the log files reside.

In all cases, the z/OS user credentials must be a valid z/OS user ID and password combination and have READ access to the PowerExchange data set that is defined in the DTLCAMAP DD statement of the PowerExchange Listener JCL.

To extract z/OS data from PowerExchange Logger from Linux, UNIX, and Windows log files, the z/OS user credentials also must have READ access to CAPX.CND.\* resource profiles in the FACILITY class, which are managed by your z/OS security product.

For more information, see the *PowerExchange Reference Manual*.

# Task Flow for Extracting Change Data

Use this task flow to identify the tasks that you need to complete to configure and start extraction processing. You do these tasks in the PowerExchange Navigator, PowerCenter Designer, and PowerCenter Workflow Manager.

Before you begin, complete configuration of the data source and PowerExchange, and create capture registrations in the PowerExchange Navigator.

1. Edit the extraction map if necessary.

You can make the following changes:

- Deselect any column for which you do not want to extract change data. PowerExchange still captures change data for these columns.
- Add change indicator (CI) and before image (BI) columns.

2. To test the extraction map, do a database row test on the extraction map in the PowerExchange Navigator.
3. In Designer, import metadata for the sources and targets.
4. In Designer, configure a mapping to extract and process change data.
5. In Workflow Manager, configure a connection and session.
6. Create restart tokens for the CDC session.
7. Configure the restart token file.
8. If you want to stop extraction processing based on user-defined events, implement event table processing.
9. To offload column-level extraction processing and UOW Cleanser processing from the source system to the PowerCenter Integration Service machine, configure offload processing. You can also use offload processing to offload change data to a remote PowerExchange Logger for Linux, UNIX, and Windows process on another machine.  
  
If you configure offload processing for real-time extractions, you can also configure multithreaded processing to help improve throughput.
10. Start the CDC session.

## Testing an Extraction Map

In the PowerExchange Navigator, perform a database row test to verify that PowerExchange can retrieve change data from a registered source based on an extraction map.

A database row test enables you to:

- Preview change data that PowerExchange captured for the registered data source.
  - Preview change data that either PowerExchange Condense on i5/OS or z/OS or the PowerExchange Logger for Linux, UNIX, and Windows captured for registered source.
  - Verify that the extraction map properly maps the captured change data.
1. In the PowerExchange Navigator, open the extraction group and the extraction map.
  2. Select the extraction map and click **File > Database Row Test**.
  3. In the **Database Row Test** dialog box, enter information in the following fields:

**DB Type**

An option that indicates the extraction mode:

- **CAPXRT**. Real-time extraction mode or continuous extraction mode.
- **CAPX**. Batch extraction mode.

**Location**

Node name for the location of the system on which the captured change data resides. This name must be defined in a NODE statement in the dbmover.cfg configuration file on the Windows machine from which you run the database row test.

**UserID and Password**

Optional. A user ID and password that provides access to the source data.

**Fetch**

To preview data, select **Data**.

**Application**

An application name. For a row test, an application name is not required. However, you must enter at least one character in this field. PowerExchange does not retain this value.

**SQL Statement**

A SQL SELECT statement that PowerExchange generates for the fields in the extraction map. You can edit this statement, if necessary.

In the statement, a table is identified as follows:

*Schema.RegName\_TableName*

Where:

- *Schema* is a schema name for the extraction map.
- *RegName* is the name of the capture registration that corresponds to the extraction map.
- *TableName* is the table name of the data source.

**Note:** If you enter **CAPX** in the **DB Type** field, you can extract change data only after PowerExchange Condense or the PowerExchange Logger for Linux, UNIX, and Windows closes at least one condense file or log file. Otherwise, PowerExchange does not display change data and writes message PWX-04520 to the PowerExchange message log. PowerExchange also writes this message if no change data for the source has been captured, condensed, or logged.

4. Click **Advanced**.
5. Complete the fields in the **CAPX Advanced Parameters** dialog box or **CAPXRT Advanced Parameters** dialog box.
  - If you use continuous extraction mode, enter the CAPX CAPI\_CONNECTION name in the **CAPI Connection Name** field.
  - If you offload change data to PowerExchange Logger for Linux, UNIX, and Windows log files on a system that is remote from the source, enter the location of the extraction maps in the **Location** field.
6. Click **OK**.
7. Click **Go**.

The database row test returns each change from the extraction start point, by column. The results include the PowerExchange-generated CDC columns, which provide information such as the change type, timestamp, and user ID.

# Configuring PowerCenter CDC Sessions

After you import metadata for CDC data sources and targets into PowerCenter, you can create a mapping, connection, and a CDC session for extracting change data. You must configure many session and connection attributes.

## Changing Default Values for Session and Connection Attributes

Some PowerCenter session and application connection attributes have default values that are appropriate only for bulk data movement operations. You must edit these attributes for CDC sessions.

The following table describes the session and connection attributes that you need to set for CDC, including the recommended values:

Attribute Name	Attribute Location	Recommended Value for CDC	Description
Commit Type	Properties Tab for the session	Source	Default value is <b>Target</b> . If you accept the default, the PowerCenter Integration Service automatically overrides the default to use source-based commit processing. However, you should change this attribute to <b>Source</b> so that you can disable the <b>Commit On End Of File</b> attribute.
Commit On End Of File	Properties Tab for the session	Disabled	By default, this attribute is enabled. If you accept the default, the PowerCenter Integration Service commits the change data in the buffer to the targets when the session ends. The final commit occurs after the PWXPC CDC reader has committed all complete UOWs in the buffer, along with their restart tokens, to the targets. This timing can cause the restart tokens and target data to be out of sync. The final restart tokens might represent a point in the change stream that is earlier than final change data that the PowerCenter Integration Service commits to the targets. As a result, duplicate data might occur when the CDC session restarts. To prevent potential duplicate data, disable this attribute.
Recovery Strategy	Properties Tab for the session	Resume from last checkpoint	Default value is <b>Fail task and continue workflow</b> . To properly restart CDC session, PowerExchange CDC and PWXPC require that this option is set to <b>Resume from last checkpoint</b> .
Stop on errors	Config Object Tab for the session	1	Default value is 0. By default, the PowerCenter Integration Service does not consider errors when writing to targets as fatal. The following types of error are non-fatal: <ul style="list-style-type: none"><li>- Key constraint violations</li><li>- Loading nulls into a not null field</li><li>- Database trigger responses</li></ul> If write errors occur, change data loss might occur because PWXPC has advanced the restart tokens values. To maintain target data and restart token integrity, set this option to 1.



Attribute Name	Attribute Location	Recommended Value for CDC	Description
Application Name	Application Connection	Enter a unique name for each CDC session.	Default is the first 20 characters of the workFlow name. <b>Attention:</b> Because the default might not result in a unique name, enter a unique name.
RestartToken File Folder	Application Connection	Default value	The default is \$PMRootDir/Restart. This default is acceptable for CDC.
RestartToken File Name	Application Connection	Enter a unique name for each CDC session.	If you enter an <b>Application Name</b> value, the default is that application name. If you do not enter an <b>Application Name</b> value, the default is the workflow name. <b>Attention:</b> Because a default might not result in a unique name, enter a unique restart token file name.
Number of Runs to Keep RestartToken File	Application Connection	1 or greater	Default is 0. PWXPC keeps only one backup copy of the restart token initialization and termination files. Enter a value greater than 0 to make history available for recovery purposes.

## Configuring Application Connection Attributes

To extract change data, you must configure certain application connection attributes. For a complete list of all PWX CDC application connection attributes, see *PowerExchange Interfaces for PowerCenter*.

### Image Type

Use the **Image Type** attribute to indicate how PWXPC passes captured Updates to CDC sessions that extract and apply the updates to the target.

Enter one of the following options for this attribute:

- **AI.** Process Updates as Update operations. PWXPC passes each Update as a single Update record. An Update record includes after images of the data only, unless you add before image (BI) and change indicator (CI) fields to the extraction map that you import for the source definition for the CDC session.
- **BA.** Process Updates as Deletes followed by Inserts. PWXPC passes each Update as a Delete record followed by an Insert record. The Delete record contains the before image of the data, and the Insert record contains the after image.

Default is **BA**.

If you use **BA**, PWXPC generates, for each captured Update operation, a Delete record that contains the before image of the data and an Insert record that contains the after image. If you also define BI and CI fields for some columns in the extraction map that you import for the source definition, PWXPC populates the BI and CI fields with data in both the generated Delete and Insert records. However, for any Insert and Delete operations captured from the source, the BI and CI fields in the generated Delete and Insert records contain Null values.

If you specify **AI**, you can still use before images of the data, if available, in extraction processing. PWXPC can embed before-image data and after-image data in the same Update row. To embed before-image data, you must complete the following configuration tasks:

- In the PowerExchange Navigator, add BI and CI fields to the extraction map that you plan to import for the source definition in PowerCenter.

- If you use batch or continuous extraction mode, enter BA for the CAPT\_IMAGE parameter in the PowerExchange Condense or PowerExchange Logger for Linux, UNIX, and Windows configuration file. This setting causes both before and after images to be stored in the PowerExchange Logger log files or PowerExchange Condense condense files. When CDC sessions run, they extract data from these files.

Informatica recommends that you use the **AI** setting if you want to process before images of data. CDC sessions can process a single Update record more efficiently than separate Delete and Insert records to get the before image data.

For example, embed before-image data and after-image data in the same Update row to handle changes to primary keys. Relational databases that allow changes to primary keys, such as DB2 for z/OS, treat these Updates as equivalent to deleting the row and readding it with a new key value. To enable PowerExchange to detect primary key changes, include BI and CI fields for the primary key columns in the extraction map for the source definition. Then, in PowerCenter, define a Flexible Target Key Custom transformation to apply the changes to the target as a Delete followed by an Insert. Include the transformation in the mapping for the CDC session. If a target relational database does not allow changes to primary keys, updates to primary keys fail.

**Note:** To use a Flexible Target Key Custom transformation, you must set the **Image Type** attribute to **AI** and configure BI and CI fields in the PowerExchange extraction map for the source.

For more information about adding BI and CI columns, see the *PowerExchange Navigator User Guide*.

## CAPI Connection Name Override

If you define multiple CAPI\_CONNECTION statements in the DBMOVER configuration file, you can use the **CAPI Connection Name Override** connection attribute to select one of the statements for a CDC session.

PowerExchange allows a maximum of eight CAPI\_CONNECTION statements in the DBMOVER configuration file. You might want to use multiple CAPI\_CONNECTION statements to extract changes for multiple source types with a single PowerExchange Listener on a single machine. For example, you can extract changes for Oracle and DB2 sources through a single PowerExchange Listener by specifying multiple CAPI\_CONNECTION statements.

If you use CDC offload processing, you must define the CAPI\_CONNECTION statements in the dbmover.cfg file on the PowerCenter Integration Service machine. If you do not use CDC offload processing, you must define the CAPI\_CONNECTION statements on the system where the change data resides.

To specify the CAPI\_CONNECTION statement to use for a specific CDC session, enter the name of the CAPI\_CONNECTION statement in the **CAPI Connection Name Override** connection attribute. By using the override instead of a default CAPI\_CONNECTION statement, you clearly indicate which statement to use for a session.

## Idle Time

Use the **Idle Time** connection attribute to indicate whether a CDC session that uses real-time or continuous extraction mode runs continuously or shuts down after it reaches the end-of-log (EOL).

You can specify that PowerExchange wait for a certain period without change activity before shutting down.

Enter one of the following values:

- -1. The CDC session runs continuously. PowerExchange returns an end-of-file (EOF) only when you manually stop the CDC session.
- 0. After reaching the EOL, PowerExchange returns an EOF and the CDC session ends.  
If you want a CDC session to end periodically on an active system that is rarely idle, enter 0.

- *n*. After reaching the EOL, PowerExchange waits the specified number of seconds, *n*. If PowerExchange receives no change data of interest during this time interval, PowerExchange sends an EOF to the PowerCenter Integration Service and the CDC session ends successfully.  
If you enter a low value, such as 1, the CDC session might end before PowerExchange has read all available data in the change stream.

Default is -1.

PowerExchange determines the EOL by using the current end of the change stream at the point that PowerExchange started to read the change stream. PowerExchange uses the concept of EOL because the change stream is usually not static. The actual EOL is continually moving forward. After PowerExchange reaches the EOL, it writes message PWX-09967 in the PowerExchange message log.

Often, CDC sessions that run in real-time or continuous extraction mode use the default value of -1. You can manually stop a long-running CDC session by using the PowerCenter Workflow Monitor, pmcmd commands, or the PowerExchange STOPTASK command.

If you set the **Idle Time** attribute to 0, when PowerExchange reaches the EOL, it returns an EOF to PWXPC. PWXPC and the PowerCenter Integration Service then perform the following processing:

1. PWXPC flushes all buffered UOWs and ending restart tokens to the targets.
2. The CDC reader ends.
3. After the PowerCenter Integration Service finishes writing the flushed data to the targets, the writer ends.
4. After any post-session commands and tasks run, the CDC session ends.

If you set the **Idle Time** attribute to a positive number, the following processing occurs:

1. PowerExchange reads the change stream until it reaches EOL and then the **Idle Time** wait interval begins.
2. If more data is in the change stream after the EOL, PowerExchange continues to read the change stream, looking for change data of interest to the CDC session, as follows:
  - If the idle time expires before PowerExchange reads a change record of interest for the CDC session, PowerExchange stops reading the change stream.
  - If PowerExchange reads a change record of interest to the CDC session, PowerExchange restarts the timer, passes the change data to PWXPC, and continues to read the change stream. This processing continues until the idle time expires.
3. After the idle time expires, PowerExchange passes an EOF to PWXPC.
4. PWXPC and the PowerCenter Integration Service perform the same processing as when the **Idle Time** value is 0 and the CDC session ends.

When a CDC session ends because the idle time elapsed or a PowerExchange STOPTASK command was issued, PWXPC writes the following message in the session log:

```
[PWXPC_10072] [INFO] [CDCDispatcher] session ended after waiting for [idle_time]
seconds. Idle Time limit is reached
```

If you stop a continual CDC session with the PowerExchange STOPTASK command, PWXPC substitutes 86400 for the *idle\_time* variable in the PWXPC\_10072 message.

**Note:** If you specify both the **Reader Time Limit** and **Idle Time** attributes, the PowerCenter Integration Service stops reading data from the source when one of these attribute conditions is met, whichever one is first. Because the reader time limit does not result in normal termination of a CDC session, Informatica recommends that you use only the idle time limit.

## Restart Control Attributes

Use PWXPC restart control attributes to identify restart information to use for a CDC session. The restart information determines the point from which PowerExchange starts reading change data for the session.

Specify restart control attributes in the following situations:

- When you create CDC session.
- When you add a source to an existing CDC session and need to specify restart information for that source.
- When you want to override some restart information that is in the state table or file for a CDC session.

The following table describes the restart control attributes that you can enter on a PWX CDC application connection:

Connection Attribute	Description
Application Name	A unique application name for the CDC session. The application name is case sensitive and cannot exceed 20 characters.  Default is the first 20 characters of the workflow name. Because the default might not result in a unique name, Informatica recommends that you enter a unique name.
RestartToken File Folder	Directory name on the PowerCenter Integration Service machine that contains the restart token override file.  Default is \$PMRootDir/Restart.
RestartToken File Name	The unique file name of the restart token file. This file is in the directory that is specified in the <b>RestartToken File Folder</b> attribute. PWXPC uses the contents of this file, if any, in conjunction with the state table or state file to determine the restart point for the CDC session.  Default is the <b>Application Name</b> value, or if you do not specify the application name, default is the workflow name.

**Attention:** The values for the **Application Name** and **RestartToken File Name** attributes must be unique for each CDC session. If either one of these values is not unique, unpredictable results might occur, including session failure and potential data loss.

## Event Table Processing

Use event table processing to stop the extraction of changes based on user-defined events, such as an end-of-day event.

For example, to stop an extraction process every night, after all changes for the day are processed, write a change to the event table at midnight. This change triggers PowerExchange to stop reading change data and shut down the extraction process after the current UOW completes.

Use the following rules and guidelines:

- You can use event table processing only with real-time or continuous extraction modes.
- You must create the event table and define the applications that can update the table.
- You must register the event table for change data capture from the PowerExchange Navigator.
- A CDC session monitors a single event table. Each user-defined event requires its own event table and a separate extraction process.
- The event table and all of the source tables in the CDC session must be of the same source type.

## Implementing Event Table Processing

Use this procedure to implement event table processing. With event table processing, you can stop change data extraction processing based on user-defined events.

1. Create an event table.

The event table must be of the same source type and on the same machine as the change data to be extracted. For example, if you extract DB2 change data on z/OS, the event table must be a DB2 table in the same DB2 subsystem as the DB2 source tables for the extraction.

2. In the PowerExchange Navigator, create a capture registration for the event table.

When you create the capture registration, the PowerExchange Navigator generates a corresponding extraction map.

3. In PowerCenter, create a CDC connection and session.

In the **Event Table** attribute on the PWX CDC Real Time application connection, enter the name of the extraction map associated with the capture registration that you created.

4. Define applications that write an update to the event table whenever the defined event occurs.

PowerExchange reads the update and places an end-of-file (EOF) in the change stream. PWXPC processes the EOF, passes it to the PowerCenter Integration Service, and then shuts down the PowerExchange reader. The PowerCenter Integration Service completes writing all of the data that is in the pipeline to the targets and then ends the CDC session.

## Flush Latency

PowerExchange reads change data into a buffer on the source system, or into a buffer on the PowerCenter Integration Service machine if you use offload processing. The PowerExchange Consumer API (CAPI) periodically flushes the buffer to transfer the change data to PWXPC on the PowerCenter Integration Service machine.

The CAPI flushes the buffer to PWXPC when the one of the following events occurs:

- The buffer becomes full.
- The CAPI timeout value that is specified by the **PowerExchange Latency in seconds** attribute on the PWX CDC Real Time connection expires.
- A commit point occurs.

To specify the flush latency for CDC sessions that run in real-time or continuous extraction mode, set the **PWX Latency in seconds** attribute on the PWX CDC Real Time application connection. This attribute specifies the maximum time that PowerExchange waits for more change data before flushing data to PWXPC. This attribute applies to PowerExchange on the source system, or to the PowerExchange client on the PowerCenter Integration Service machine if you use offload processing.

For CDC sessions that use batch extraction mode, PowerExchange always uses 2 seconds for the flush latency.

PowerExchange writes message PWX-09957 to the PowerExchange message log to identify the CAPI timeout value based on the **PWX Latency in seconds** attribute. If you select **Retrieve PWX Log Entries** on the application connection, PWXPC also writes this message in the session log.

After PowerExchange flushes the change data, PWXPC provides the data to the source qualifiers in the CDC session for further processing. Then the PowerCenter Integration Service commits the data to the targets.

**Note:** The **PWX Latency in seconds** value also affects how fast a CDC session responds to a stop command from Workflow Monitor or pmcmd program. Before PWXPC can process a stop request, it must wait for PowerExchange to return control to it. Use the default value of 2 seconds for the **PWX Latency in seconds** attribute to avoid unacceptable delays in stop command processing.

## Target Latency

Target latency is the total time for applying change data to the targets.

This total includes the time that PWXPC takes to extract change data from the change stream and the time that PowerCenter Integration Service takes to apply that change data to the targets. If extraction and apply processing occurs quickly, target latency is low.

The values for the commitment control attributes affect target latency. When you set the commitment control attributes, balance target latency requirements with resource consumption on the PowerCenter Integration Service machine and the target databases.

Lower target latency values result in higher resource use. The increased resource use occurs because the PowerCenter Integration Service must flush the change data more frequently. Also, the target databases must process more commit requests.

The following table describes the default values for the commitment control attributes, which provide the lowest latency:

Attribute	Default
Maximum Rows Per commit	0, which disables this attribute
Minimum Rows Per commit	0, which disables this attribute
Real-time Flush Latency in milli-seconds	0, which is equivalent to 2000 milliseconds or 2 seconds
UOW Count	1

These values decrease target latency because PWXPC commits changes after each UOW or on UOW boundaries. However, these values can have the following drawbacks:

- Highest resource consumption on the source system, PowerCenter Integration Service machine, and target databases
- Decreased throughput for the CDC sessions because PWXPC flushes change data too frequently for the PowerCenter Integration Service or target databases to handle this processing

To lower resource consumption and potentially increase throughput for CDC sessions, enter a value greater than the default value for one of the following attributes:

- **Minimum Rows Per commit**
- **UOW Count**
- **Real-time Flush Latency in milli-seconds**

Then disable the other attributes.

## Commitment Control Attributes

PWXPC, in conjunction with PowerExchange and the PowerCenter Integration Service, controls the timing of commit processing for CDC sessions based on commitment control attributes on PWX CDC connections.

Commit processing is not controlled by a single commitment control attribute. When setting these attributes, try to balance performance and resource consumption with latency requirements.

The **Maximum Rows Per commit**, **Real-Time Flush Latency in milli-seconds**, and **UOW Count** attributes control the timing of real-time flushes of change data to the targets. The **Minimum Rows Per commit** attribute controls if a commit can occur.

Set one or more of the following commitment control attributes on PWX CDC connections:

#### **Maximum Rows Per commit**

Maximum number of change records in a source UOW that PWXPC processes before flushing the data buffer to commit the change data to the targets.

Use this attribute to have PWXPC commit change data to the targets without waiting for the UOW boundary, or end-UOW, to be met. This type of commit is called a *subpacket commit*. By using subpacket commits for large UOWs, you can minimize use of storage on the PowerCenter Integration Service machine and locking contention on the target databases.

**Attention:** Because PWXPC can commit the change data to the targets between UOW boundaries, relational integrity (RI) might be compromised. Do not use this connection attribute if you have targets in the CDC session with RI constraints.

After the maximum rows limit is met, PWXPC flushes the change data from the buffer on the PowerCenter Integration Service machine and commits the data to the targets. PWXPC also writes message PWXPC\_12128 to the session log. After commit processing completes, the RDBMS releases locks on the target databases and PowerCenter Integration Service can reuse the buffer space for additional change records.

The maximum rows limit is cumulative across all sources in the CDC session. PWXPC issues a real-time flush when the limit is met, regardless of the number of sources with changes.

PWXPC resets the maximum rows limit when a real-time flush occurs. The flush can occur because of the maximum rows limit, UOW count limit, or real-time flush latency timer.

If PWXPC reaches a UOW boundary and the maximum row limit has not been met, PWXPC continues to process change records across UOW boundaries.

Use a maximum rows limit if you have extremely large UOWs in the change stream that might cause the following problems:

- Locking issues on the target database
- Resource issues on the PowerCenter Integration Service node

For example, you have a large UOW with 10,000 updates for a single source, and you set the **Maximum Rows per Commit** attribute to 1000. In this case, PWXPC issues a subpacket commit after each 1,000 change records.

Or, you might have a UOW that contains updates for more than one source. For example, the UOW contains 900 updates for source 1, 100 updates for source 2, and then 500 more updates for source 1. If you set the **Maximum Rows per Commit** attribute to 1000, PWXPC issues a subpacket commit after reading 1,000 change records, or after processing the updates for source 2.

Default is 0, which causes PWXPC to not use this maximum rows limit. If you specify 0 or do not enter a value for the maximum rows limit, commits occur only on UOW boundaries.

If you specify a low maximum rows limit, the CDC session uses more system resources on the PowerCenter Integration Service machine and target systems. This increased resource use occurs because PWXPC flushes data to the targets more frequently.

**Note:** The **Maximum Rows Per commit** attribute is a count of the records within a UOW. The **UOW Count** attribute is a count of complete UOWs.

#### **Minimum Rows Per commit**

Minimum number of change records that PowerExchange must pass to PWXPC before passing a commit record. Until the minimum rows limit is met, PowerExchange discards any commit records that it reads from the change stream and passes only change records to PWXPC. After the minimum rows limit

is met, PowerExchange passes the next commit record it encounters to PWXPC and then resets the minimum rows counter.

If the change stream has many small UOWs, you can set the **Minimum Rows Per commit** attribute to create larger UOWs of a more uniform size. Online transactions that run in transaction control systems such as CICS and IMS often commit after only a few changes, which results in many, small UOWs in the change stream. PowerExchange and PWXPC process a few large UOWs more efficiently than many small UOWs. By using the minimum rows limit to increase the size of UOWs, you can improve CDC processing efficiency.

The minimum rows limit does not impact the relational integrity of the change data because PowerExchange does not create additional commit points in the change data. PowerExchange skips some of the original commit records in the change stream.

Default is 0, which causes PowerExchange to not use this minimum rows limit.

If you enter a minimum rows limit, PowerExchange changes the number of change records in a UOW to match or exceed this limit.

**Note:** PWXPC does not commit change data to the targets based on the minimum rows limit. PWXPC commits change data to the targets based on the **Maximum Rows Per commit**, **Real-Time Flush Latency in milli-seconds**, and **UOW Count** attributes.

#### **Real-Time Flush Latency in milli-seconds**

For real-time or continuous extraction mode, the number of milliseconds that must elapse before PWXPC flushes the data buffer to commit change data to the targets. After the flush latency interval expires and PWXPC reaches a UOW boundary, PWXPC issues a real-time flush to commit change data and restart tokens to the targets. PWXPC also writes message PWXPC\_10082 in the session log.

PWXPC resets the flush latency interval when a real-time flush occurs. The flush can occur because of the maximum rows limit, UOW count limit, or real-time flush latency timer.

Valid values for the real-time flush latency are:

- -1. Disables data flushes based on time.
- 0 through 2000. Sets the interval to 2000 milliseconds, or 2 seconds.
- 2000 through 86400. Sets the interval to the specified number of milliseconds.

Default is 0.

If you set the flush latency interval value is 0 or greater, PWXPC flushes the change data for all complete UOWs after the interval expires and the next UOW boundary occurs. The lower you set the flush latency interval, the faster PWXPC commits change data to the targets. If you require a low latency for applying changes to the targets, enter a low value for the flush latency interval.

However, if you specify a low flush latency interval, the CDC session might consume more system resources on the PowerCenter Integration Service and target systems. This increased consumption occurs because PWXPC commits change data to the targets more frequently.

#### **UOW Count**

Number of complete UOWs that PWXPC reads from the change stream before flushing change data to the targets. When PWXPC reads change data from PowerExchange and provides it to the source qualifier in the CDC session, the count of the UOWs begins.

After the UOW count limit is met, PWXPC issues a real-time flush to commit the change data and restart tokens to the targets. PWXPC also writes message PWXPC\_10081 in the session log.

PWXPC resets the UOW count after a real-time flush occurs because of the UOW count limit or the real-time flush latency interval.



Valid values for UOW count are:

- -1 or 0. PWXPC does not use the **UOW Count** attribute to control commit processing.
- 1 through 999999999. PWXPC flushes change data after reading the specified number of UOWs.

Default is 1.

The lower you set the UOW count value, the faster the PowerCenter Integration Service commits change data to the target. If you require the lowest possible latency, enter a UOW count of 1. However, a low latency might result in the session using more system resources on the PowerCenter Integration Service and the target systems.

**Attention:** In the session properties, verify that the **Commit Type** attribute specifies **Source** and that the **Commit at End of File** attribute is disabled. The **Commit at End of File** attribute is enabled by default. If you accept the default, the PowerCenter Integration Service writes additional data to the targets after the CDC reader has committed the restart tokens and shut down. When you restart the CDC session, the session might write duplicate data to the targets.

## Examples of Controlling Commit Processing

Review the following examples to learn how to use the commitment control attributes to control commit processing with PWXPC.

### Example 1. Subpacket Commit and UOW Count

This example uses the **Maximum Rows Per commit** and **UOW Count** attributes to control commit processing.

The change data is composed of UOWs of the same size. Each UOW contains 1,000 change records.

The following table describes the commitment control attribute values that this example uses:

Attribute	Value
Maximum Rows Per commit	300
Minimum Rows Per commit	0, which disables this attribute
Real-time Flush Latency in milli-seconds	0, which is equivalent to 2 seconds
UOW Count	1

Based on the maximum rows value, PWXPC flushes the data buffer after reading the first 300 records in a UOW. This action commits the change data to the targets. PWXPC continues to commit change data to the targets every 300 records.

PWXPC commits on UOW boundaries only for the UOW count and real-time flush latency interval. If the real-time flush latency interval expires before PWXPC reads 300 change records, PWXPC still commits based on the maximum rows value because that threshold is met before a UOW boundary occurs.

When the end of the UOW is read, PWXPC commits the change data because the **UOW Count** value is 1. PWXPC resets the UOW and maximum row counters and the real-time flush latency timer each time it commits. Because all of the UOWs have the same number of change records, PWXPC continues to read change data and to commit the data to the targets at the same points in each UOW.

In this example, PWXPC commits change data at the following points:

- 300 change records based on the maximum rows value
- 600 change records based on the maximum rows value
- 900 change records based on the maximum rows value
- 1,000 change records based on the UOW count value

## Example 2. UOW Count and Time-Based Commits

This example uses the **UOW Count** and **Real-time Flush Latency in milli-seconds** attributes to control commit processing. The change data consists of UOWs of varying sizes.

The following table describes the commitment control attribute values that this example uses:

Attribute	Value
Maximum Rows Per commit	0, which disables this attribute
Minimum Rows Per commit	0, which disables this attribute
Real-time Flush Latency in milli-seconds	5000, which is equivalent to 5 seconds
UOW Count	1000

Initially, PWXPC reads 900 complete UOWs in 5 seconds. Because the real-time flush latency interval has expired, PWXPC flushes the data buffer to commit the change data to the targets. PWXPC then resets both the UOW counter and real-time flush latency timer. When PWXPC reaches UOW 1,000, PWXPC does not commit change data to the targets because the UOW counter was reset to 0 after the last commit.

PWXPC reads the next 1,000 UOWs in 4 seconds, which is less than the real-time flush latency timer. PWXPC commits this change data to the target because the UOW counter has been met. After this commit, PWXPC then resets the real-time flush latency timer and the UOW counter.

PWXPC continues to read change data and commit the data to the targets, based on the UOW count or the real-time flush latency flush time, whichever limit is met first.

In this example, PWXPC commits change data at the following points:

- After UOW 900 because the real-time latency flush latency timer matched first.
- After UOW 1,900 because the UOW count matched first during the second commit cycle.

## Example 3. Minimum Rows and UOW Count

This example uses the **Minimum Rows Per commit** and **UOW Count** attributes to control commit processing.

The change data consists of UOWs of the same size. Each UOW contains ten change records.

The following table describes the commitment control attribute values that this example uses:

Attribute	Value
Maximum Rows Per commit	0, which disables this attribute
Minimum Rows Per commit	100

Attribute	Value
Real-time Flush Latency in milli-seconds	-1, which disables this attribute
UOW Count	10

PWXPC passes the minimum rows value to PowerExchange and requests change data from the change stream. Because the minimum rows value is 100, PowerExchange skips the commit records of the first nine UOWs. When PowerExchange reads the last change record in the tenth UOW, the minimum rows limit is met. So, PowerExchange passes the commit record for the tenth UOW to PWXPC and resets the minimum rows counter. PWXPC increases the UOW counter to one.

PowerExchange and PWXPC continue to read the change data until the UOW counter is 10. At this point, PWXPC flushes the data buffer to commit the change data to the targets and resets the UOW counter.

PWXPC commits change data after 1,000 change records, or after every 10 UOWs, because each UOW contains 10 change records and the **UOW Count** is 10.

## Recovery and Restart Processing for CDC Sessions

If you select the **Resume from last checkpoint** option for the **Recovery Strategy** attribute in a CDC session that extracts change data, PWXPC and PowerCenter provide recovery and restart processing for that session.

If a session fails, the PowerCenter Integration Service recovers the session state of operation, and PWXPC recovers the restart information.

PWXPC saves restart information for all sources that are in a CDC session. The restart information for CDC sessions, including the restart tokens, originates from PowerExchange on the system from which the change data is extracted. You can include both relational and nonrelational targets in a single CDC session. PWXPC uses one of the following locations to store and retrieve restart information, based on the target type:

- For relational targets, PWXPC uses recovery state tables in the target databases. PWXPC, in conjunction with the PowerCenter Integration Service, commits both the change data and the restart tokens for that data in the same commit operation. This commit ensures that the applied data and the restart tokens are in-sync.
- For nonrelational targets, PWXPC uses the recovery state file that is in the shared location on the PowerCenter Integration Service machine. PWXPC, in conjunction with the PowerCenter Integration Service, writes the change data to the target files and then writes the restart tokens to the recovery state file. As a result, duplicate data might be applied to the targets when you restart the failed CDC sessions.

The PowerCenter Integration Service saves the session state of operation and maintains target recovery tables. The PowerCenter Integration Service stores the session state of operation in the shared location that is specified in \$PMStorageDir. The PowerCenter Integration Service saves relational target recovery information in the target database.

When you run a CDC session that uses a resume recovery strategy, PWXPC writes the following message to the session log to indicate that recovery is in effect:

```
PWXPC_12094 [INFO] [CDCRestart] Advanced GMD recovery in effect. Recovery is automatic.
```

When you recover or restart a CDC session, PWXPC uses the saved restart information to resume reading the change data from the point of interruption. The PowerCenter Integration Service restores the session state of operation, including the state of each source, target, and transformation. PWXPC, in conjunction with the

PowerCenter Integration Service, determines how much of the source data it needs to reprocess. PowerExchange and PWXPC use the restart information to determine the correct point in the change stream from which to restart extracting change data and then applying it to the targets.

If you run a session with resume recovery strategy and the session fails, do not change the mapping, the session, or the state information before you restart the session. PowerCenter and PWXPC cannot guarantee recovery if you make any of these changes.

**Restriction:** If any of the targets in the CDC session use the PowerCenter File Writer to write CDC data to flat files, do not use a resume recovery strategy. Restart tokens for all targets in the CDC session, including relational targets, will be compromised if a flat file target is in the same session. Data loss or duplication might occur.

## RELATED TOPICS:

- [“Creating Restart Tokens for Extractions” on page 241](#)
- [“Configuring the Restart Token File” on page 242](#)

## PowerCenter Recovery Tables for Relational Targets

When the PowerCenter Integration Service runs a CDC session that has a resume recovery strategy, it writes information to recovery tables on the target database system.

When the PowerCenter Integration Service recovers the session, it uses the information in the recovery tables to determine where to begin loading data to target tables. PWXPC also uses information in the recovery tables to determine where to begin reading the change stream.

If you want the PowerCenter Integration Service to create the recovery tables, grant table creation privileges to the database user name that is configured in the target database connection. Otherwise, you must create the recovery tables manually.

For relational targets, the PowerCenter Integration Service creates the following recovery tables in the target database:

### PM\_RECOVERY

Contains target load information for the session run. The PowerCenter Integration Service removes the information from this table after each successful session and initializes the information at the beginning of subsequent sessions.

### PM\_TGT\_RUN\_ID

Contains information the PowerCenter Integration Service uses to identify each target on the database. The information remains in the table between session runs. If you manually create this table, you must create a row and enter a value other than zero for LAST\_TGT\_RUN\_ID to ensure that the session recovers successfully.

### PM\_REC\_STATE

Contains state and restart information for CDC sessions. PWXPC stores the application name and restart information for all sources in the CDC session. The PowerCenter Integration Service stores any state information for the session. Unlike the session state information, restart information persists in this table across successful sessions. The PowerCenter Integration Service updates it with each commit to the target tables.

If you edit or drop the recovery tables before you recover a session, the PowerCenter Integration Service cannot recover the session. Also, PWXPC cannot restart the CDC session from the point of interruption.

If you disable recovery, the PowerCenter Integration Service does not remove the recovery information from the target database. Also, PWXPC no longer updates the restart information in the target database.

## Recovery State Table

The recovery state table, PM\_REC\_STATE, contains state and CDC restart information for a CDC session. This table resides in the same target database as the target tables.

The PowerCenter Integration Service creates an entry in the state table for each CDC session. These entries can comprise more than one row. CDC sessions with heterogeneous target tables have state table entries in each unique relational target database and an entry in a state file on the PowerCenter Integration Service machine for each nonrelational target. For example, a CDC session that targets Oracle and SQL Server tables and a MQ Series queue has an entry in the state table in the target Oracle database, in the state table in the target SQL Server database, and in the state file on the PowerCenter Integration Service machine.

Each session entry in a state table contains a number of repository identifiers and execution state data such as the checkpoint number and CDC restart information. The following columns can contain PWXPC-specific restart information:

### APPL\_ID

Contains the value the PWXPC creates by appending the task instance ID of the CDC session to the value that you specify in the **Application Name** attribute in the source PWX CDC application connection. When this value matches an APPL\_ID value for a row in the state table, the PowerCenter Integration Service, in conjunction with PWXPC, selects the row from the state table for the CDC session.

### STATE\_DATA

Contains the restart information for the session in a variable-length, 1,024-byte binary column. When the PowerCenter Integration Service commits change data to the targets tables, it also commits the restart information for that data in this column. PWXPC uses the restart information from this column to perform restart processing for the CDC session.

If the amount of restart information for a session exceeds 1,024 bytes, the PowerCenter Integration Service adds additional rows to accommodate the remainder of the restart information. For each row added, the PowerCenter Integration Service increases the value of the SEQ\_NUM column by one, starting from zero.

## PowerCenter Recovery Files for Nonrelational Targets

If you configure a resume recovery strategy for a CDC session, the PowerCenter Integration Service stores the session state of operation in the shared location, \$PMStorageDir, on the PowerCenter Integration Service machine. For nonrelational targets, the PowerCenter Integration Service also stores the target recovery status in a recovery state file in the shared location on the PowerCenter Integration Service machine. PWXPC stores the restart information for nonrelational target files in this state file.

## Recovery State File

For all nonrelational targets in a CDC session, the PowerCenter Integration Service uses a recovery state file on the PowerCenter Integration Service machine.

Nonrelational target files include MQ Series message queues, PowerExchange nonrelational targets, and other PowerCenter nonrelational targets.

CDC sessions with heterogeneous target tables have state table entries in each unique relational target database and an entry in a state file on the PowerCenter Integration Service machine for each nonrelational target.

The PowerCenter Integration Service creates the recovery state file in the shared location, \$PMStorageDir. The file name has the following prefix:

```
pm_rec_state_appl_id
```

PWXPC creates the value for the *appl\_id* variable in the file name by appending the task instance ID of the CDC session to the value that you specify in the **Application Name** attribute in the source PWX CDC application connection. The PowerCenter Integration Service uses various task and workflow repository attributes to complete the file name. The message CMN\_65003, which the PowerCenter Integration Service writes to the session log, contains the complete file name.

## Application Names

PWXPC, in conjunction with the PowerCenter Integration Service, uses the application name you specify as part of the key when it stores and retrieves the restart information for a CDC session.

When you configure the PWX CDC application connection for a CDC session, enter a unique value for the **Application Name** attribute. PWXPC appends the repository task instance ID for the CDC session to this value to create the APPL\_ID value in the recovery state table and the *appl\_id* portion of the recovery state file name.

Because the value of the APPL\_ID column and the state recovery file contains the task instance ID for the CDC session, changes to the session can affect restart processing. If you add or remove sources or targets for a CDC session, you must use the restart token file to provide restart tokens and then cold start the session.

## Restart Processing for CDC Sessions by Start Type

How you start a CDC session affects how PWXPC determines the restart points for sources in the session. Each source has its own restart point.

For each start type, PWXPC determines the restart point as follows:

- For a cold start, PWXPC uses the restart token file to acquire restart tokens for all data sources. PWXPC does not read the state tables or state file and does not attempt to recover the session. The CDC session continues to run until it is stopped or interrupted.
- For a warm start, PWXPC reconciles the restart tokens that are in the restart token file with the restart tokens in the state tables and state file. If necessary, PWXPC performs recovery processing. The session continues to run until it is stopped or interrupted.
- For a recover start, PWXPC reads the restart tokens from any applicable state tables and state file. If necessary, PWXPC performs recovery processing. PWXPC updates the restart token file with the restart tokens for each source in the CDC session, and then the session ends.

Before you run a CDC session for the first time, create and populate the restart token file with restart token pair for each source in the session. Each restart token pair should match a point in the change stream where the source and target are in a consistent state.

For example, materialize a target table and stop update activity on the source. To define a start or restart point, specify a special override statement that contains the CURRENT\_RESTART option in the restart token file. Use the restart token file that has the file name that matches the restart token file name in the PWX CDC application connection. When you cold start the CDC session, PWXPC requests that PowerExchange use the current end-of-log as the extraction start point. You can then resume update activity on the sources.

If you cold start a CDC session and a restart token file does not exist, the PowerCenter Integration Service runs the session. PWXPC passes Null restart tokens for all sources to PowerExchange. PowerExchange issues message PWXPC\_12060 to indicate that the restart tokens for each source are Null and then assigns the default restart point to each source.

**Attention:** If you use Null restart tokens, the CDC session might have incorrect results. Provide valid restart tokens when you cold start CDC sessions.

## Default Restart Points for Null Restart Tokens

If PowerExchange receives null restart tokens for all sources in a CDC session, it uses the default restart points.

The following table describes the default restart points for data sources on Linux, UNIX, and Windows, by source type and extraction mode:

Data Source	Batch and Continuous Extraction Modes	Real-time Extraction Mode
DB2 for Linux, UNIX, and Windows	Oldest PowerExchange Logger for Linux, UNIX, and Windows log file that is recorded in the CDCT file.	Current log position at the time the PowerExchange capture catalog was created.
Microsoft SQL Server	Oldest PowerExchange Logger log file that is recorded in the CDCT file.	Oldest data available in the Publication database.
Oracle	Oldest PowerExchange Logger log file that is recorded in the CDCT file, for both PowerExchange Oracle CDC with LogMiner and Express CDC for Oracle.	Earliest available point in the change stream: <ul style="list-style-type: none"><li>- For PowerExchange Oracle CDC with LogMiner, the most recent Oracle catalog dump in the archived logs.</li><li>- For PowerExchange Express CDC for Oracle, the beginning of the most recent archive log.</li></ul>

PowerExchange uses the default restart point only if all sources in a CDC session have null restart tokens. If some sources have non-null restart tokens, PWXPC assigns the oldest restart point from those tokens to any sources for which no restart tokens are specified.

For example, a new CDC session contains the sources A, B, and C. The restart token file contains restart tokens for sources A and B. The restart point for source A is older than that for source B. Source C does not have existing or supplied restart tokens. Because some sources in the CDC session have explicit restart points, PWXPC does not assign null restart tokens to source C. Instead, PWXPC assigns the restart point for source A to source C because this restart point is the oldest one supplied.

## Determining the Restart Tokens for Cold Start Processing

When you cold start a CDC session, PWXPC uses the restart token file to determine the restart tokens for all sources. PWXPC ignores any entries in the state tables or state file for the sources in the CDC session.

More specifically, PWXPC uses one of the following methods to determine the restart tokens:

- If the restart token file is empty or does not exist, PWXPC assigns null restart tokens to all sources in the CDC session.
- If the restart token file contains only explicit override statements, PWXPC performs the following processing:
  - Assigns the restart tokens in the explicit override statements to the specified sources.
  - Assigns the oldest supplied restart point to any sources for which an explicit override statement was not specified.
- If the restart token file contains only the special override statement, PWXPC assigns the restart tokens in the special override statement to all sources.
- If the restart token file contains a special override statement and explicit override statements, PWXPC performs the following processing:
  - Assigns the restart tokens in the explicit override statements to the specified sources.

- Assigns the restart tokens in the special override statement to all remaining sources.

## Determining the Restart Tokens for Warm Start Processing

When you warm start a CDC session, PWXPC uses the state tables and state file, in conjunction with restart token file, to determine the restart tokens for all sources.

More specifically, PWXPC uses one of the following methods to determine the restart tokens:

- If the restart token file is empty or does not exist and there is no matching entry in a state table or state file, PWXPC assigns null restart tokens to all sources in the session.
- If the restart token file is empty or does not exist and if some but not all sources have a matching entry in a state table or a state file, PWXPC performs the following processing:
  - Assigns any restart tokens found in a state table and state file to the appropriate sources.
  - Assigns the oldest available restart point to all sources that do not have restart tokens.
- If the restart token file is empty or does not exist and if all sources have an entry in a state table or state file, PWXPC uses the restart tokens from the state tables or state file.
- If the restart token file contains explicit override statements and no sources have a matching entry in a state table or no state file, PWXPC performs the following processing:
  - Assigns the restart tokens in the explicit override statements to the specified sources.
  - Assigns the oldest supplied restart point to all sources that do not have restart tokens.
- If the restart token file contains explicit override statements and if some but not all sources have a matching entry in a state table or a state file, PWXPC performs the following processing:
  - Assigns the restart tokens in the explicit override statements to the specified sources.
  - Assigns restart tokens from a state table or state file to the appropriate sources, provided that the tokens have not been supplied in the restart token file.
  - Assigns the oldest available restart point to all sources that do not have restart tokens supplied in the restart token file or from a state table or state file.
- If the restart token file contains explicit override statements and if all sources have an entry in a state table or a state file, PWXPC performs the following processing:
  - Assigns the restart tokens in the explicit override statements to the specified sources.
  - Assigns the restart tokens from state tables or the state file to all remaining sources that do not have restart tokens supplied in the restart token file.
- If the restart token file contains only the special override statement, PWXPC assigns the restart tokens in the special override statement to all sources.
- If the restart token file contains a special override statement and explicit override statements, PWXPC performs the following processing:
  - Assigns the restart tokens in the explicit override statements to the specified sources.
  - Assigns the restart tokens in the special override statement to all remaining sources.



# Creating Restart Tokens for Extractions

Before you begin extracting change data, you must create restart tokens to indicate the extraction start point.

When generating restart tokens, consider the following points:

- The optimal start point matches the point in the change stream at which you last synchronized the source and target. This point marks the end of the change stream, or current end-of-log (EOL), if you stop update activity on the source, as recommended, until after target materialization and restart token generation are complete.
- The length of restart tokens vary by source type.
- For Microsoft SQL Server, the sequence token represents the point from which PowerExchange starts reading change data from the SQL Server distribution database. The restart token includes the DBID of the distribution database and the name of the distribution server. PowerExchange verifies that the distribution database in the restart token matches the distribution database that is specified in the CAPI\_CONNECTION statement.
- PWXPC can generate restart tokens when it starts extraction processing for a CDC session. PowerExchange also provides methods of generating restart tokens.

To create current restart tokens for the current EOL, use one the following methods:

## **PWXPC restart token file**

To generate current restart tokens for a CDC session that uses real-time or continuous extraction mode, specify the CURRENT\_RESTART option on the RESTART1 and RESTART2 special override statements in the PWXPC restart token file. When the CDC session runs, PWXPC requests that PowerExchange provide restart tokens for the current EOL. PWXPC uses this restart information to locate the extraction start point.

## **Database row test**

In the PowerExchange Navigator, perform a database row test with a SELECT CURRENT\_RESTART SQL statement.

## **DTLUAPPL utility**

Run the DTLUAPPL utility with the GENERATE RSTKKN option.

If you use the DTLUAPPL utility or PowerExchange Navigator to generate restart tokens, enter the token values in the PWXPC restart token file before you start the CDC session.

# Displaying Restart Tokens

You can display restart token values in the output from a database row test, extraction session, or DTLUAPPL PRINT function.

If you run a database row test on an extraction map from the PowerExchange Navigator, the output includes a restart token pair for each row of change data. The following columns show the token values:

- DTL\_\_CAPXRESTART1 shows the sequence token value.
- DTL\_\_CAPXRESTART2 shows the restart token value.

If you include the DTL\_\_CAPXRESTART1 and DTL\_\_CAPXRESTART2 columns in the PowerCenter source definition, PowerExchange provides the restart tokens for each row when you extract change data in a CDC session.

When a CDC session runs, PowerExchange and PWXPC display restart token values in the following messages:

- In messages PWX-04565 and PWX-09959, the sequence token is in the Sequence field and restart token is in the PowerExchange Logger field.
- In messages PWXPC\_12060 and PWXPC\_12068, the sequence token is in the Restart Token 1 field and the restart token is in the Restart Token 2 field.
- In messages PWXPC\_10081, PWXPC\_10082, and PWXPC\_12128, the sequence token is the first token value and the restart token is the second token value.

If you use the DTLUAPPL utility to generate restart tokens, you can use the PRINT statement to display the generated values. In the PRINT output, DTLUAPPL displays the sequence token, without the usual trailing eight zeros, in the Sequence field and displays the restart token in the Restart field.

## Configuring the Restart Token File

When you configure the CDC session in PowerCenter, specify the name and location of the restart token file.

To specify the restart token file, enter the following attributes on the PWX CDC application connection for the source:

### RestartToken File Folder

Enter the name of the directory that contains the restart token file. If you use the default value of \$PMRootDir/Restart and the Restart directory does not exist, PWXPC creates the directory. PWXPC does not create any restart token directory under another name.

### RestartToken File Name

Enter a unique name for the restart token file. If you do not specify this value, PWXPC uses the value in the **Application Name** attribute, if present. Otherwise, PWXPC uses the workflow name. Because this name must be unique, Informatica recommends that you always specify a value for the **RestartToken File Name** attribute.

When you run a CDC session, PWXPC verifies that the restart token file exists. If one does not exist, PWXPC uses the name specified in this attribute to create an empty restart token file.

**Restriction:** The value of **RestartToken File Name** attribute must be unique for every CDC session. Non-unique file names can cause unpredictable results, such as change data loss and session failures.

To find the restart token file name for a CDC session, use the following methods:

- For CDC sessions that have run, look for message PWXPC\_12057 in the session log. This message indicates the restart token file directory and file name.
- In Workflow Manager, look for the restart token file folder and file name in the attributes on the PWX CDC application connection associated with the source in the CDC session. If the restart token file name is not present, PWXPC uses the application name, if specified. Otherwise, PWXPC uses the workflow name.

Before you run a CDC session the first time, configure the restart token file to indicate the point in the change stream from which to start extracting change data. Later, you might need to modify the restart token file to add sources to a CDC session or to indicate the point from which to restart change data extraction.

## Restart Token File Statements

You can specify explicit override, special override, and comment statements in the restart token file.

Use these statements as follows:

- Explicit override. Use this statement type to specify a restart token pair for a specific source. You must provide the PowerExchange extraction map name.
- Special override. Use this statement type to specify a restart token pair for one or more sources. You can provide a specific restart token pair or request that PowerExchange use the current restart point.
- Comment. Use this statement type to enter any comments that you want to add to the file.

### Restart Token File Syntax by Statement Type

In the restart token file, you can specify explicit override statements, special override statements, and comments.

For explicit override control statements, use the following syntax:

```
extraction_map_name={sequence_token|CURRENT_RESTART}  
extraction_map_name={restart_token|CURRENT_RESTART}
```

For special override control statements, use the following syntax:

```
RESTART1={sequence_token|CURRENT_RESTART}  
RESTART2={restart_token|CURRENT_RESTART}
```

For comments, use the following syntax:

```
<!-- comment_text
```

The following syntax rules and guidelines apply:

- Statements can begin in any column.
- All statements are optional.
- Do not include blank lines between statements.
- Comment lines must begin with:  

```
<!--
```
- In each file, you can specify one or more sets of explicit override statements and one set of special override statements.
- Explicit override statements for a source take precedence over any set of special override statements.

### Explicit Override Statements

Use explicit override statements to specify a restart token pair for a specific source.

You can specify explicit override statements for one or more sources in a CDC session. Also, you can use explicit override statements in conjunction with special override statements to provide restart tokens for all sources in a CDC session.

When you warm start a CDC session, the explicit override statements for a source override the restart tokens in the state table or state file for that source.

When defining explicit override statements for a source, specify a pair of statements, each containing an extraction map name and a restart or sequence token value. Because a source can have multiple extraction maps with distinct names, you might have multiple pairs of explicit override statements for a source.

Explicit override statements use the following parameters:

```
extraction_map_name={restart1_token|CURRENT_RESTART}  
extraction_map_name={restart2_token|CURRENT_RESTART}
```

Review the following parameter descriptions:

#### ***extraction\_map\_name***

The name of an extraction map for the data source. To determine the extraction map name, use one of the following methods:

- For CDC data map sources, see the **Schema Name Override** and **Map Name Override** attributes in the session properties. These attributes override the schema name and map name in the source extraction map. Or, in Designer, see the **Schema Name** and **Map Name** values in the source Metadata Extensions.
- For relational sources, see the **Extraction Map Name** attribute in the session properties.

#### ***restart1\_token***

The sequence token portion of a restart token pair. This value varies based on data source type.

#### ***restart2\_token***

The restart token portion of a restart token pair. This value based on data source type.

#### **CURRENT\_RESTART**

PowerExchange generates restart tokens for the current end of the change stream. The PWXPC CDC reader opens a separate connection to PowerExchange, requests the generation of current restart tokens, and then provides the token values to all applicable sources.

You can generate current restart tokens in the **Database Row Test** dialog box of the PowerExchange Navigator.

**Restriction:** Use CURRENT\_RESTART only for CDC sessions that use real-time extraction mode or continuous extraction mode.

## Special Override Statement

Use a special override statement to specify or generate restart tokens for one or more sources.

You can use a special override statement to provide restart tokens for all sources in a CDC session. Also, you can use a special override statement in conjunction with explicit override statements.

When you warm start a CDC session, the special override statement overrides the restart tokens in the state table or state file for all sources except those specified in explicit override statements.

A special override statement is composed of a pair of RESTART1 and RESTART2 statements, as shown in the following syntax:

```
RESTART1={restart1_token|CURRENT_RESTART}  
RESTART2={restart2_token|CURRENT_RESTART}
```

Specify only one set of these special override statements in the restart token file.

In RESTART1 and RESTART2 statements, use the following parameters to specify either a pair of sequence and restart token values or the current end of the change stream:

#### ***restart1\_token***

The sequence token portion of a restart token pair. This value varies based on data source type.

## restart2\_token

The restart token portion of a restart token pair. This value varies based on data source type.

## CURRENT\_RESTART

PowerExchange generates restart tokens for the current end of the change stream. The PWXPC CDC reader opens a separate connection to PowerExchange, requests the generation of current restart tokens, and then provides the token values to all applicable sources.

You can generate current restart tokens in the **Database Row Test** dialog box of the PowerExchange Navigator.

**Restriction:** Use CURRENT\_RESTART only for CDC sessions that use real-time extraction mode or continuous extraction mode.

## Comment Statements

You can insert a comment statement anywhere in the restart token file.

Comment statements must begin with:

```
<!--
```

## Example Restart Token File

This example restart token file is for a CDC session with seven source tables. The file includes explicit override statements that provide restart tokens for three source tables and a special override statement that provides restart tokens for the remaining tables.

The restart token file contains the following statements:

```
<!-- Restart Tokens for existing tables -->
Restart1=000000AD775600000000000000AD77560000000000000000
Restart2=C1E4E2D34040000000AD5F2C000000000
<!-- Restart Tokens for the Table: rrtb0001_RRTB_SRC_001 -->
dldsn9.rrtb0001_RRTB_SRC_001=0000060D1DB20000000000000000060D1DB200000000000000000
dldsn9.rrtb0001_RRTB_SRC_001=C1E4E2D3404000000013FF36200000000
<!-- Restart Tokens for the Table: rrtb0001_RRTB_SRC_002 -->
dldsn9.rrtb0002_RRTB_SRC_002=000000A371950000000000000000A3719500000000000000000
dldsn9.rrtb0002_RRTB_SRC_002=C1E4E2D340400000000968FC600000000
<!-- Restart Tokens for the Table: rrtb0001_RRTB_SRC_004 -->
dldsn9.rrtb0004_RRTB_SRC_004=000006D84E780000000000000006D84E7800000000000000000
dldsn9.rrtb0004_RRTB_SRC_004=C1E4E2D3404000000060D1E6100000000
```

When you warm start the CDC session, PWXPC reads the restart token file to process any override statements for restart tokens. In this case, the restart token file overrides all restart tokens for all sources in the CDC session. After resolving the restart tokens for all sources, PWXPC writes message PWXPC\_12060 to the session log with the following information:

```
=====
Session restart information:
=====
Extraction Map Name      Restart Token 1      Restart Token 2      Source
dldsn9.rrtb0001_RRTB_SRC_001 0000060D1DB20000000000000000060D1DB200000000000000000 C1E4E2D3404000000013FF36200000000 Restart file
dldsn9.rrtb0002_RRTB_SRC_002 000000A371950000000000000000A3719500000000000000000 C1E4E2D340400000000968FC600000000 Restart file
dldsn9.rrtb0003_RRTB_SRC_003 000000AD77560000000000000000AD77560000000000000000 C1E4E2D34040000000AD5F2C000000000 Restart file (special override)
dldsn9.rrtb0004_RRTB_SRC_004 000006D84E780000000000000006D84E78000000000000000 C1E4E2D3404000000060D1E6100000000 Restart file
dldsn9.rrtb0005_RRTB_SRC_005 000000AD77560000000000000000AD77560000000000000000 C1E4E2D34040000000AD5F2C000000000 Restart file (special override)
dldsn9.rrtb0006_RRTB_SRC_006 000000AD77560000000000000000AD77560000000000000000 C1E4E2D34040000000AD5F2C000000000 Restart file (special override)
dldsn9.rrtb0007_RRTB_SRC_007 000000AD77560000000000000000AD77560000000000000000 C1E4E2D34040000000AD5F2C000000000 Restart file (special override)
```

PWXPC indicates the source of the restart token values for each source. For the sources that had explicit override statements in the restart token file, PWXPC writes “Restart file” in the Source column.

For the sources to which PWXPC assigns the special override restart tokens, PWXPC writes “Restart file (special override)” in the Source column.

## CHAPTER 11

# Managing Change Data Extractions

This chapter includes the following topics:

- [Starting PowerCenter CDC Sessions, 246](#)
- [Stopping PowerCenter CDC Sessions, 248](#)
- [Changing PowerCenter CDC Sessions, 250](#)
- [Recovering PowerCenter CDC Sessions, 252](#)

## Starting PowerCenter CDC Sessions

Use Workflow Manager, Workflow Monitor, or pmcmd to start a workflow or task for a CDC session. You can do a cold start, warm start, or recovery start. The method you use determines how PWXPC gets the restart information.

Also, you can start the entire workflow, part of a workflow, or a task in the workflow.

Use one of the following methods to start a CDC session:

### **Cold start**

To cold start a CDC session, use the Cold Start command in Workflow Manager or Workflow Monitor. You can also use the pmcmd starttask or startworkflow commands with the norecovery option. A CDC session that uses real-time or continuous extraction mode runs continuously until it is stopped or interrupted. A CDC session that uses batch extraction mode runs until it reaches the end of log (EOL) or it is stopped or interrupted.

When you cold start a CDC session, PWXPC uses the restart token file to acquire restart tokens for all sources. PWXPC does not read the state tables or file or makes any attempt to recover the session.

### **Warm start**

To warm start a CDC session, use the Start or Restart commands in Workflow Manager or Workflow Monitor. You can also use the pmcmd starttask or startworkflow commands. A CDC session that uses real-time or extraction mode runs continuously until it is stopped or interrupted. A CDC session that uses batch extraction mode runs until it reaches EOL or it is stopped or interrupted.

When you warm start a CDC session, PWXPC reconciles any restart tokens provided in the restart token file with any restart tokens that exist in the state tables or file. If necessary, PWXPC performs recovery processing.

## Recovery start

To start recovery for a CDC session, use the Recover command from Workflow Manager or Workflow Monitor. You can also use the `pmcmd recoverworkflow` command or the `starttask` or `startworkflow` commands with the recovery option. When recovery completes, the CDC session ends.

When you recover a CDC session, PWXPC reads the restart tokens from any applicable state tables or state file. If necessary, PWXPC performs recovery processing. PWXPC updates the restart token file with the restart tokens for each source in the CDC session. Then the session ends. To begin extracting change data again, either cold start or warm start the session.

## Cold Start Processing

To cold start workflows and tasks, use the Cold Start command in Workflow Manager or Workflow Monitor. Alternatively, you can use the `pmcmd starttask` or `startworkflow` commands with the `norecovery` option.

After you request a cold start for a CDC session, the following processing occurs:

1. PWXPC writes the following message in the session log:  

```
PWXPC_12091 [INFO] [CDCRestart] Cold start requested
```
2. PWXPC reads the restart tokens from only the restart token file and associates a restart token with each source in the session.
3. PWXPC creates the initialization restart token file with the initial restart tokens.
4. PWXPC commits the restart tokens for each source to the appropriate state tables or state file and then writes message PWXPC\_12104 to the session log.
5. PWXPC passes the restart tokens to PowerExchange. PowerExchange begins extracting change data and passing the data to PWXPC for processing.
6. PWXPC continues processing change data from PowerExchange and committing the data and restart tokens to the targets. This processing continues until the session ends or you stop it.

## Warm Start Processing

To warm start workflows and tasks, use the Start or Restart command in Workflow Manager or Workflow Monitor. Alternatively, you can use the `pmcmd starttask` or `startworkflow` commands.

When you warm start a workflow or task, PWXPC automatically performs recovery. You do not need to recover failed workflows and tasks before you restart them.

After you request a warm start for a CDC session, the following processing occurs:

1. PWXPC writes the following message in the session log:  

```
PWXPC_12092 [INFO] [CDCRestart] Warm start requested. Targets will be resynchronized automatically if required
```
2. PWXPC queries the PowerCenter Integration Service about the commit levels of all targets. If all targets in the session have the same commit level, PWXPC skips recovery processing.
3. PWXPC reconciles the restart tokens from the restart token file and from the state tables or file.  
**Restriction:** If a CDC session requires recovery processing, PWXPC does not use the restart token file. Consequently, you cannot override restart tokens for sources.
4. PWXPC creates the initialization restart token file with the reconciled restart tokens.
5. If recovery is required, PWXPC re-reads the change data for the last unit-of-work (UOW) that was committed to the targets with the highest commit level and then flushes the data to the targets with lower commit levels. The PowerCenter Integration Service commits flushed change data and restart tokens to any relational targets and updates any nonrelational files.

6. If recovery is not required and the reconciled restart tokens differ from those in the state tables or state file, PWXPC commits the reconciled restart tokens and then writes message PWXPC\_12104 to the session log.
7. PWXPC passes the restart tokens to PowerExchange. PowerExchange begins extracting change data and passing the data to PWXPC for processing.
8. PWXPC continues processing change data from PowerExchange and commits the data and restart tokens to the targets. This processing continues until the session ends or you stop it.

## Recovery Processing

To recover workflows and tasks, use the Recover command in Workflow Manager or Workflow Monitor. Alternatively, you can use the `pmcmd recoverworkflow` command, or the `starttask` or `startworkflow` command with the `recovery` option.

Use the recovery start method to populate the restart token file with the restart tokens for all sources in a CDC session. You can then cold start the CDC session or verify that the targets and restart tokens are in a consistent state. However, you do not need to recover failed workflows and tasks before you restart them because PWXPC automatically performs recovery processing when you warm start a workflow or task.

After you request recovery for a CDC session, the following processing occurs:

1. PWXPC writes the following message in the session log:  

```
PWXPC_12093 [INFO] [CDCRestart] Recovery run requested. Targets will be  
resynchronized if required and processing will terminate
```
2. PWXPC queries the PowerCenter Integration Service about the commit levels of all targets. If all targets in the session have the same commit level, PWXPC skips recovery processing.
3. PWXPC reads the restart tokens from the recovery state tables or state file.  
**Restriction:** If a CDC session requires recovery processing, PWXPC does not use the restart token file. Consequently, you cannot override restart tokens for sources.
4. PWXPC creates the initialization restart token file with the reconciled restart tokens.
5. If recovery is required, PWXPC re-reads the change data for the last UOW that was committed to the targets with the highest commit level and then flushes the data to the targets with lower commit levels. The PowerCenter Integration Service commits any flushed change data and restart tokens to any relational targets, and updates any nonrelational files.
6. PWXPC updates the restart token file with the final restart tokens, creates the termination restart token file, and ends.

To process change data from the point of recovery, warm start or cold start the workflow or task.

## Stopping PowerCenter CDC Sessions

You can stop CDC sessions from PowerCenter or PowerExchange.

In PowerCenter, issue the Stop or Abort command in Workflow Monitor. Alternatively, use the `pmcmd stoptask`, `stopworkflow`, `aborttask`, or `abortworkflow` commands.

- If you issue the Stop command in Workflow Monitor or use the `pmcmd stoptask` or `stopworkflow` command, the PWXPC CDC reader and PowerCenter Integration Service complete processing all of the data in the pipeline and shut down. Then, the CDC session ends.



- If you issue the Abort command in Workflow Monitor or use the `pmcmd aborttask` or `abortworkflow` command, the PowerCenter Integration Service waits 60 seconds to allow the readers and writers to complete processing all of the data in the pipeline and shut down. If the PowerCenter Integration Service cannot finish processing and committing data within this period, it kills the DTM process and ends the CDC session.

For more information about these PowerCenter commands, see the *Informatica Command Reference* or *PowerCenter Workflow Basics Guide*.

In PowerExchange, issue the PowerExchange Listener `STOPTASK` command in one of the following ways:

- From the command line on the system where extraction processing occurs
- From the PowerExchange Navigator
- With the DTLUTSK utility
- With the `pwxcmd` program

When you issue the `STOPTASK` command, PowerExchange stops the extraction task in the PowerExchange Listener and passes an EOF to the PowerCenter Integration Service. Then the CDC session ends. For more information about the `STOPTASK` command, see the *PowerExchange Command Reference*.

## Stop Command Processing

After you issue a stop command in PowerCenter or PowerExchange, the following processing occurs:

**Note:** To stop CDC sessions and workflows, you can use the Stop command in Workflow Monitor or the `pmcmd stoptask` or `stopworkflow` command. Alternatively, you can use the PowerExchange `STOPTASK` command.

1. If you use a PowerCenter stop command, the PowerCenter Integration Service requests PWXPC to stop. If you use the PowerExchange `STOPTASK` command, PowerExchange sends an EOF to PWXPC.
2. When PWXPC receives an EOF, it flushes any complete and uncommitted UOWs and the associated restart tokens to the targets. PWXPC then writes messages `PWXPC_12101` and `PWXPC_12068` to the session log.
3. The PowerCenter Integration Service processes all of data in the pipeline and writes it to the targets.
4. The PowerCenter Integration Service sends an acknowledgment to PWXPC indicating that the targets have been updated.
5. PWXPC writes the termination restart token file, and then writes the message `PWXPC_12075` to the session log.
6. The PWXPC CDC reader shuts down.
7. The PowerCenter Integration Service performs any post-session tasks and ends the session.

## Terminating Conditions

You can have CDC sessions stop based on user-defined events or at EOL if you configure certain terminating conditions.

When PWXPC encounters a terminating condition, it stops reading change data from sources, flushes change data to the targets, and passes an EOF to the PowerCenter Integration Service. The PowerCenter Integration Service commits the data to the targets and ends the CDC session.

Use the following connection attributes and features as terminating conditions:

### Event table processing

Create an event table and a capture registration for the table. Then specify the extraction map for the table in the **Event Table** attribute of the PWX CDC Real Time application connection for the CDC session. After PowerExchange reads a change record for the event table, it passes an EOF to PWXPC to end the CDC session.

### Idle time

Enter 0 for the **Idle Time** attribute on a PWX CDC Real Time application connection. Then, whenever PowerExchange reaches EOL, it passes an EOF to PWXPC to end the CDC session.

### Batch extraction mode

If you use batch extraction mode, PowerExchange reads all closed PowerExchange Condense condense files or PowerExchange Logger for Linux, UNIX, and Windows log files. Then PowerExchange passes an EOF to PWXPC to end the CDC session.

## Changing PowerCenter CDC Sessions

Use this procedure to change CDC sessions. You might need to add or remove sources and targets.

After you change a CDC session, you must cold start it. Because a cold start is required, you must also get the latest restart tokens for the original sources before restarting the session. To do so, you can perform a recovery.

To change a CDC session:

1. Stop the workflow.
2. After the workflow ends, recover the CDC session.  
When you recover tasks, PWXPC writes the ending restart tokens for all sources in a CDC session to the restart token file that you specified on the PWX CDC application connection.
3. Make changes to the session or workflow, if necessary.
4. Verify that the restart token file in the source CDC connection points to the same restart token file updated in the recovery.
5. If you add sources to the CDC session, add statements to the restart token file that provide restart tokens for the additional sources.
6. If you remove sources from the CDC session, update the restart token file to remove their restart tokens.
7. Cold start the CDC session.

## Examples of Adding Sources and Creating Restart Tokens

The following examples show how to add sources to CDC sessions and create restart tokens for those sources.

The first example uses the `CURRENT_RESTART` option of the special override statement in the restart token file to generate current restart tokens. The second example uses `DTLUAPPL` to generate current restart tokens.

## Example 1. Creating Current Restart Tokens with Special Override Statements

This example adds a source table, RRTB\_SRC\_004, to a CDC session that has three other sources. You edit the restart token file to generate restart tokens that represent the current end of the change stream for the additional source.

In the restart token file, you define special override statements with CURRENT\_RESTART option for the RRTB\_SRC\_004 source.

For the other three sources, you retain the existing restart points.

To add a source with CURRENT\_RESTART restart tokens:

1. In Workflow Monitor, use the Stop command to stop the workflow.
2. After the workflow stops, select the Recover Task command to run a recovery session.

PWXPC writes the following messages in the session log:

```
PWXPC_12060 [INFO] [CDCRestart]
=====
Session restart information:
=====
Extraction Map Name      Restart Token 1      Restart Token 2      Source
dldsn9.rrtb0002_RRTB_SRC_002 000000AD220F00000000000000AD220F0000000000000000 C1E4E2D34040000000AD0D9C000000000 GMD storage
dldsn9.rrtb0001_RRTB_SRC_001 000000AD220F00000000000000AD220F0000000000000000 C1E4E2D34040000000AD0D9C000000000 GMD storage
dldsn9.rrtb0003_RRTB_SRC_003 000000AD220F00000000000000AD220F0000000000000000 C1E4E2D34040000000AD0D9C000000000 GMD storage
```

PWXPC also writes the restart tokens in the restart token file that is identified the CDC application connection attributes.

3. Edit the mapping, session, and workflow to add the source RRTB\_SRC\_004.
4. Edit the restart token file to add RESTART1 and RESTART2 special override statements that specify the CURRENT\_RESTART option for the RRTB\_SRC\_004 source.

The updated file appears as follows:

```
<!-- existing sources
dldsn9.rrtb0001_RRTB_SRC_001=000000AD220F00000000000000AD220F0000000000000000
dldsn9.rrtb0001_RRTB_SRC_001=C1E4E2D34040000000AD0D9C000000000
dldsn9.rrtb0002_RRTB_SRC_002=000000AD220F00000000000000AD220F0000000000000000
dldsn9.rrtb0002_RRTB_SRC_002=C1E4E2D34040000000AD0D9C000000000
dldsn9.rrtb0003_RRTB_SRC_003=000000AD220F00000000000000AD220F0000000000000000
dldsn9.rrtb0003_RRTB_SRC_003=C1E4E2D34040000000AD0D9C000000000
<!-- new source
RESTART1=CURRENT_RESTART
RESTART2=CURRENT_RESTART
```

5. Cold start the session.

PWXPC connects to PowerExchange and generates restart tokens that match the current end of the change stream for the RRTB\_SRC\_004 source. PWXPC passes the generated restart tokens to PowerExchange to begin change data extraction. Because the restart points for the other sources are earlier than the restart point for RRTB\_SRC\_004, PWXPC does not pass any change data for RRTB\_SRC\_004 until it reads the first change after the generated restart point.

## Example 2. Creating Current Restart Tokens with the DTLUAPPL Utility

This example adds the source table, RRTB\_SRC\_004, to a CDC session that has three other sources. You use the DTLUAPPL utility to generate restart tokens that represent the current end of the change stream.

For the other three sources, you retain the existing restart points.

1. In Workflow Monitor, use Stop command to stop the workflow.
2. After the workflow stops, select the Recover Task command to run a recovery session.

PWXPC writes the following messages in the session log:

```
PWXPC_12060 [INFO] [CDCRestart]
=====
Session restart information:
=====
Extraction Map Name      Restart Token 1      Restart Token 2      Source
dldsn9.rrtb0002_RRTB_SRC_002 000000AD220F00000000000000AD220F0000000000000000 C1E4E2D34040000000AD0D9C000000000 GMD storage
dldsn9.rrtb0001_RRTB_SRC_001 000000AD220F00000000000000AD220F0000000000000000 C1E4E2D34040000000AD0D9C000000000 GMD storage
dldsn9.rrtb0003_RRTB_SRC_003 000000AD220F00000000000000AD220F0000000000000000 C1E4E2D34040000000AD0D9C000000000 GMD storage
```

PWXPC also writes the restart tokens in the restart token file that is identified in the CDC application connection attributes.

3. Edit the mapping, session, and workflow to add the source RRTB\_SRC\_004.
4. Run the DTLUAPPL utility with RSTTKN GENERATE parameter to generate restart tokens that represent the current end of the change stream for the additional source.

Use the following DTLUAPPL control cards:

```
mod APPL dummy DSN7 rsttkn generate
  mod rsttkn rrtb004
end appl dummy
print appl dummy
```

The PRINT command produces the following output:

```
Registration name=<rrtb004.1> tag=<DB2DSN7rrtb0041>
Sequence=<00000DBF240A00000000000000DBF240A00000000>
Restart =<C1E4E2D340400000DBF238200000000>
```

You can add eight zeros to the end of the Sequence value to create the sequence value for the restart token file.

5. Edit the restart token file to add the source and its restart tokens.

The updated file contains the following lines:

```
<!-- existing sources
dldsn9.rrtb0001_RRTB_SRC_001=000000AD220F00000000000000AD220F0000000000000000
dldsn9.rrtb0001_RRTB_SRC_001=C1E4E2D34040000000AD0D9C000000000
dldsn9.rrtb0002_RRTB_SRC_002=000000AD220F00000000000000AD220F0000000000000000
dldsn9.rrtb0002_RRTB_SRC_002=C1E4E2D34040000000AD0D9C000000000
dldsn9.rrtb0003_RRTB_SRC_003=000000AD220F00000000000000AD220F0000000000000000
dldsn9.rrtb0003_RRTB_SRC_003=C1E4E2D34040000000AD0D9C000000000
<!-- new source
dldsn9.rrtb0004_RRTB_SRC_004=00000DBF240A0000000000000DBF240A000000000000000
dldsn9.rrtb0004_RRTB_SRC_004=C1E4E2D340400000DBF238200000000
```

6. Cold start the session.

PWXPC passes the restart tokens to PowerExchange to begin change data extraction. Because the restart points for the other sources are earlier than the restart point for RRTB\_SRC\_004, PWXPC does not pass any change data for RRTB\_SRC\_004 until it reads the first change after the generated restart point.

## Recovering PowerCenter CDC Sessions

You can use Workflow Manager, Workflow Monitor, or pmcmd to recover an entire workflow or a task in a workflow for a CDC session that fails.

A CDC session can fail for the following reasons:

- Permanent errors, such as source or target data errors
- Transitory or environmental errors, such as infrastructure problems, server failures, and network availability issues

If you run a session with a resume recovery strategy and the session fails, do not edit the state information or the mapping for the session before you restart the session.

If a session fails because of transitory or environmental errors, restart the session after you have corrected the errors. When you warm start a CDC session, PWXPC automatically performs recovery, if required. Alternatively, you can recover a CDC session, and then restart the session.

If a CDC session fails because of permanent errors, such as SQL or other database errors, you must correct the errors before restarting the CDC session. With some failures, you can correct the error and then restart the CDC session. In other cases, you might need to rematerialize the target table from the source table before you start extracting and applying change data again. If you rematerialize the target table, provide restart tokens that match the materialization point in the change stream, and then cold start the CDC session.

**Restriction:** If a CDC session requires recovery processing, you cannot override the restart tokens because PWXPC does not read the restart token file.

## Example of Session Recovery

This example describes recovery processing for a CDC session with relational targets.

Assume that you aborted the CDC session from the Workflow Monitor and then issued the Restart Task command to restart the session.

PWXPC automatically performs a recovery processing when the session warm starts and writes the following message in the session log:

```
PWXPC_12092 [INFO] [CDCRestart] Warm start requested. Targets will be resynchronized automatically if required
```

PWXPC then reads the restart tokens from the state tables and writes message PWXPC\_12060 in the session log. This message records the restart tokens for the session and its sources, for example:

```
PWXPC_12060 [INFO] [CDCRestart]
```

```
=====
Session restart information:
=====
```

Extraction Map Name	Restart Token 1	Restart Token 2	Source
dldsn8.rrtb0004_RRTB_SRC_004	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0009_RRTB_SRC_009	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0005_RRTB_SRC_005	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0006_RRTB_SRC_006	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0008_RRTB_SRC_008	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0003_RRTB_SRC_003	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0002_RRTB_SRC_002	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0001_RRTB_SRC_001	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage
dldsn8.rrtb0007_RRTB_SRC_007	00000FCA6584000000000000D2E004A00000000FFFFFFFF	C1E4E2D3404000000D21B1A500000000	GMD storage

If PWXPC detects that recovery is required, PWXPC writes message PWXPC\_12069 in the session log. This message usually includes the restart tokens for both the begin-UOW and end-UOW for the oldest uncommitted UOW that PWXPC re-reads during recovery. PWXPC usually stores end-UOW restart tokens in the state table or state file. However, if you specify a **Maximum Rows Per commit** threshold on the connection, PWXPC can commit change data and restart tokens between UOW boundaries. As a result, the restart tokens might not represent an end-UOW.

The following example PWXPC\_12069 message includes “from” restart tokens that are the same as those in the example PWXPC\_12060 message:

```
PWXPC_12069 [INFO] [CDCRestart] Running in recovery mode. Reader will resend the oldest uncommitted UOW to resync targets:
from: Restart 1 [00000FCA6584000000000000D2E004A00000000FFFFFFFF] : Restart 2 [C1E4E2D3404000000D21B1A500000000]
to: Restart 1 [00000FCA6584000000000000D300D80000000000FFFFFFFF] : Restart 2 [C1E4E2D3404000000D21B1A500000000].
```

Because this session specifies a maximum rows threshold, the restart token values in the Restart 2 fields in both the “from” and “to” restart tokens is the begin-UOW value. The sequence token values in the Restart 1 fields represent the start and end change records in the UOW that is displayed in the Restart 2 field.

During recovery processing, PWXPC reads the change data records between the restart points that are defined by the two restart token values in the PWXPC\_12069 message. Then PWXPC issues a commit for the

change data and restart tokens. The PowerCenter Integration Service writes the change data to the target tables and writes the restart tokens to the state table. Then the session ends.

# Part V: Monitoring and Tuning

This part contains the following chapters:

- [Monitoring CDC Sessions, 256](#)
- [Tuning CDC Sessions, 268](#)

## CHAPTER 12

# Monitoring CDC Sessions

This chapter includes the following topics:

- [Monitoring Overview, 256](#)
- [Monitoring CDC Sessions in PowerExchange, 256](#)
- [Monitoring CDC Sessions in PowerCenter, 263](#)

## Monitoring Overview

PowerExchange, PWXPC, and PowerCenter issue messages that you can use to monitor the progress of CDC sessions.

PWXPC can also display progress and statistical information about CDC sessions in PowerCenter Workflow Monitor.

### RELATED TOPICS:

- [“Monitoring CDC Sessions in PowerExchange” on page 256](#)
- [“Monitoring CDC Sessions in PowerCenter” on page 263](#)

## Monitoring CDC Sessions in PowerExchange

You can use certain PowerExchange messages and commands to monitor the extraction of change data by CDC sessions.

Use the following types of PowerExchange messages and output for monitoring extractions:

- Read progress messages. You can request that PowerExchange write messages that indicate the number of change records read by a CDC session.
- Extraction statistics messages. When extraction sessions end, PowerExchange writes messages that include statistical information about the change records processed.
- Multithreaded processing statistics messages. You can request that PowerExchange write statistical information about CDC sessions that use multithreaded processing.
- DISPLAY ACTIVE or LISTTASK command. Use one of these PowerExchange Listener commands, based on your operating system and mode of command execution, to list active CDC sessions. For more information about these commands, see the *PowerExchange Command Reference*.



## Read Progress Messages

You can request that PowerExchange write read progress messages to the PowerExchange message log file. These messages indicate the number of change records read for a CDC session.

If you select the **Retrieve PWX log entries** option on the PWX CDC application connection, PWXPC also writes these messages in the session log.

To have PowerExchange write read progress messages, include the following statements in the DBMOVER configuration file:

### **PRGIND=Y**

Enter Y to have PowerExchange write PWX-04587 messages to the PowerExchange message log file. These messages indicate the number of records read for a CDC session. Default is N.

### **PRGINT=records**

Enter the number of records that PowerExchange must read before writing PWX-04587 messages to the PowerExchange message log file. Default is 250 records.

For example, to have PowerExchange write read progress messages after reading 100 records, specify the following statements:

```
PRGIND=Y
PRGINT=100
```

PWX-04587 messages have the following format:

```
PWX-04587 int_server/workflow_name/session_name: Records read=records
```

Where:

- *int\_server* is the name of the PowerCenter Integration Service.
- *workflow\_name* is the name of the workflow that contains the CDC session.
- *session\_name* is the name of the CDC session.
- *records* is the cumulative number of records read since the CDC session started.

For example, for a CDC session named *s\_cdc\_DB2\_SQL\_stats* runs, PowerExchange writes the following messages:

```
PWX-04587 intserv/wf_cdc_mon_stats/s_cdc_DB2_SQL_stats: Records read=100
PWX-04587 intserv/wf_cdc_mon_stats/s_cdc_DB2_SQL_stats: Records read=200
PWX-04587 intserv/wf_cdc_mon_stats/s_cdc_DB2_SQL_stats: Records read=300
```

PowerExchange continues to write PWX-04587 messages for this CDC session until the session ends. In the PowerExchange message log file, each of these messages has a date and time stamp. Use this information to determine the speed with which PowerExchange processes change data from the change stream.

## Extraction Statistics Messages

When a CDC session ends, PowerExchange writes messages that contain statistical information about extraction processing for the session.

These messages are:

- PWX-04578. PowerExchange writes this message for each source in the CDC session. The message includes the number of Insert, Update, Delete, Commit, and total records read for the source.
- PWX-04588. PowerExchange writes this message for the entire CDC session. This message includes the total number of records read for the session.

**Important:** The statistical information in the PowerExchange messages represents the change data that PowerExchange read for a CDC session. This information might not reflect the data that was applied to the targets. For statistical information about the change data applied to a target, review the session log.

## Multithreaded Processing Statistics

If you use multithreaded processing, you can configure PowerExchange to issue messages that contain statistics on multithreaded extraction processing.

To issue these messages, you must specify the `SHOW_THREAD_PERF` statement in the `DBMOVER` configuration file on the PowerCenter Integration Service machine:

```
SHOW_THREAD_PERF=number_of_records
```

This statement specifies the number of records PowerExchange must process before writing statistics messages about multithreaded extraction processing to the PowerExchange message log file. For more information about this statement, see the *PowerExchange Reference Manual*.

If you select the **Retrieve PWX log entries** attribute on the application connection for the CDC session, PWXPC writes these messages in the session log. Also, you must specify 1 or greater for the **Worker Threads** attribute on the application connection to implement multithreaded processing so that statistics can be generated.

PowerExchange writes the following messages during each statistics interval:

- **PWX-31255.** Cycle time, which is the total time that PowerExchange on the PowerCenter Integration Service machine spent processing the change data before passing it to PWXPC. This message includes the total percentage of time and average, minimum, and maximum times in microseconds.
- **PWX-31256.** I/O time, which is the time that PowerExchange on the PowerCenter Integration Service machine spent reading change data from the PowerExchange Listener on the source system. This message includes the I/O percentage of the total time and average, minimum, and maximum times in microseconds.
- **PWX-31257.** Parsing time, which is the time that PowerExchange on the PowerCenter Integration Service machine spent in column-level processing for the change records on all threads. This message includes the parsing percentage of the total time and average, minimum, and maximum times in microseconds.
- **PWX-31258.** External time, which is the time that PowerExchange on the PowerCenter Integration Service machine spent combining the change records from all threads back into a single UOW to pass to PWXPC and for PWXPC to flush the data to PowerCenter. This message includes the external percentage of the total time and average, minimum, and maximum times in microseconds.
- **PWX-31259.** Delay time, which is the time that the PowerExchange on the PowerCenter Integration Service machine waited to receive new change records to process from the PowerExchange Listener on the source system. This message includes the delay percentage of the total time and average, minimum, and maximum times in microseconds.

For example, if you specify `SHOW_THREAD_PERF=10000`, PowerExchange writes the following messages after reading 10,000 change records and reaching the next UOW boundary:

```
PWX-31254 PowerExchange threading stats for last 10000 rows. Cycle (array) size is 25
rows. 0 out of array occurred.
PWX-31255 Cycle   time: 100% (avg:      5709 min:      4741 max:      7996 usecs)
PWX-31256 IO      time:  4% (avg:       235 min:        51 max:      1021 usecs)
PWX-31257 Parse   time:  79% (avg:     4551 min:     4102 max:     5495 usecs)
PWX-31258 Extern  time: 20% (avg:     1145 min:      618 max:     3287 usecs)
PWX-31259 Delay   time:  0% (avg:         7 min:         4 max:       165 usecs)
PWX-31254 PowerExchange threading stats for last 100000 rows. Cycle (array) size is 25
rows. 0 out of array occurred.
PWX-31255 Cycle   time:  99% (avg:     5706 min:     4735 max:     7790 usecs)
PWX-31256 IO      time:  4% (avg:      234 min:        51 max:       950 usecs)
```

PWX-31257	Parse time:	79% (avg:	4549 min:	4108 max:	5425 usecs)
PWX-31258	Extern time:	20% (avg:	1144 min:	616 max:	3242 usecs)
PWX-31259	Delay time:	0% (avg:	7 min:	4 max:	115 usecs)

If the parsing and external processing times are greater than the I/O time, you can increase the number of threads for the CDC session to try to improve throughput.

## PowerExchange Listener DISPLAY ACTIVE or LISTTASK Command

Run the PowerExchange Listener DISPLAY ACTIVE or LISTTASK command to display CDC sessions that are active in the PowerExchange Listener.

The specific command name and syntax depends on how you issue the command, as follows:

- Issue the DISPLAY ACTIVE command from the command line on the system where the PowerExchange Listener runs. For more information, see the *PowerExchange Command Reference*.
- Use the pwxcmd program to issue the listtask command to a PowerExchange Listener that runs on the local system or a remote system. For more information, see the *PowerExchange Command Reference*.
- In the PowerExchange Navigator, issue the LISTTASK command from the **Database Row Test** dialog box. For more information, see the *PowerExchange Navigator User Guide*.
- If you run the PowerExchange Listener as an application service in the Informatica domain, run the infacmd pwx program to issue the ListTaskListener command. For more information, see the *Informatica Command Reference*.

In the command output, the **PwrCntrSess** field displays the PowerCenter session name in the following format:

```
integration_server_name/workflow_name/session_name
```

For example, when two CDC sessions are active, the DISPLAY ACTIVE or LISTTASK command produces the following output:

```
PWX-00711 Active tasks:
PWX-00712 TaskId=1, Partner=10.10.10.01, Port=2480, PwrCntrSess=intserv1/workflow1/
cdc_sess1,
Application=appl_name1, Status=Active, AM=CAPXRT, Mode=Read, Process=, SessId=
PWX-00712 TaskId=2, Partner=10.10.10.02, Port=2480, PwrCntrSess=intserv2/workflow2/
cdc_sess2,
Application=appl_name2, Status=Active, AM=CAPXRT, Mode=Read, Process=, SessId=
PWX-00713 2 active tasks
PWX-00709 0 Dormant TCBs
```

## PowerExchange Listener DISPLAYSTATS Command

You can use the PowerExchange Listener DISPLAYSTATS or pwxcmd displalstats command to publish monitoring statistics for a PowerExchange Listener that runs on i5/OS, Linux, zLinux, UNIX, Windows, or z/OS.

Before you run the command, configure the following statements in the DBMOVER configuration file:

- Specify the MONITOR parameter in the STATS statement in the DBMOVER configuration file to enable PowerExchange to collect these statistics. You can include the *interval* subparameter to publish the statistics at a regular interval as well as on demand.
- For the proper display of monitoring output on z/OS, set the LOG\_LINE\_LIMIT statement to 132. Otherwise, the lines might wrap awkwardly, making the output hard to read.

You can issue the command in any of the followings ways:

- From the command line on the Linux, UNIX, Windows, or zLinux system where the PowerExchange Listener runs.
- With the MVS MODIFY (F) command on the z/OS system where the PowerExchange Listener runs.

- With the pwxcmd program from a remote Linux, UNIX, and Windows system to a Listener on any supported operating system.

**Note:** You must use this method to publish monitoring statistics for a PowerExchange Listener on i5/OS on demand.

The command syntax depends on the operating system type and whether you use pwxcmd. For more information, see the *PowerExchange Command Reference*.

Depending on which command parameter you use, you can publish one of the following types of reports:

- **Listener.** Reports PowerExchange Listener summary statistics on memory usage, CPU processing time, and activity on behalf of client requests. These statistics include counts of client tasks, connections, number of messages sent and received, bytes of data sent and received, and netport jobs (z/OS only). These statistics include both bulk data movement and CDC tasks.  
**Note:** If you run a PowerExchange Listener Service in the Informatica Domain, you can use the infacmd pwx displayStatsListener command to publish these statistics. For more information, see the *Informatica Command Reference*.
- **Accessmethods.** Reports statistics on PowerExchange Listener message and data transfer activity by client task and access method. For each active task and access method combination, these statistics include the number of rows read and written, bytes of data read and written, the source or target file name or data map file name, and the CPU processing time. For CDC requests that use the CAPX or CAPXRT access method, the report also includes counts of the SQL inserts, updates, and deletes that the task processed.
- **Clients.** Reports information about the active client tasks that are running under the PowerExchange Listener. For each task, the statistics show some or all of the following information: the status, access method, read or write mode, process name and session ID if available, CPU processing time, and start date and time. The statistics also include the client port number and IP address. If the client is PowerCenter, the statistics include the PowerCenter session ID and the application name for CDC.

By default, the Listener report is published.

The reports for a PowerExchange Listener on z/OS are similar to those for a PowerExchange Listener on i5/OS, Linux, zLinux, UNIX, or Windows.

The following example Listener report is for a PowerExchange Listener on z/OS:

```
PWX-00723 Command <displaystats Listener> succeeded
PWX-37101 Listener <PWXLST > ASID=375 (x'177') UserID=AUSRID
PWX-37102 Memory
PWX-37103 Region below 16-MB line: In Use      108 KB Limit Value      9192 KB Free      9084 KB
PWX-37104 Region above 16-MB line: In Use    53912 KB Limit Value    1675264 KB Free    1621352 KB
PWX-37117 CPU Time
PWX-37118 TCB Time = 0 SRB Time = 0 zIIP-NTime = 0
PWX-37119 Listener = 0 hrs, 0 mins, 1 secs, 275762 mcrs
PWX-37106 Cumulative Requests
PWX-37107 Total Tasks=      11 Active Tasks =      3 HWM Tasks =      3 Maxtasks =      50
PWX-37108 Connections=     11 Accepted =      11 Active =      0
PWX-37109 Msgs Sent =      0 Msgs Received=     22
PWX-37110 Data Sent =      0 Data Received=    7304
PWX-37111 NetportJobs=      0
```

The **Memory**, **TCB Time**, **SRB Time**, and **NetportJobs** values are specific to the PowerExchange Listener on z/OS. For a PowerExchange Listener on i5/OS, Linux, UNIX, or Windows, the report displays the total memory usage.

You can use this report determine if the number of client tasks is reaching the limit that is set in the MAXTASKS statement of the DBMOVER configuration file. Compare the **HWM Tasks** value to the **Maxtasks** value. If the HWM Tasks value reaches the MAXTASKS limit, PowerExchange Listener processing might be delayed, which can cause reduced throughput and connection timeouts.

The following example accessmethods report is for a PowerExchange Listener on z/OS, but the same fields are displayed for a PowerExchange Listener on i5/OS, Linux, UNIX, Windows, or zLinux:

```
PWX-00723 Command <displaystats AccessMethods> succeeded
PWX-37201 Active Access Methods
PWX-37202 Task ID = 42412 AM = CAPXRT
PWX-37203 Rows read = 1029 Rows written = 0
PWX-37204 Bytes read = 116277 Bytes written = 0
PWX-37205 File = d2ivd0.d002root_ROOT
PWX-37206 Table = <Capture Extract Realtime>
PWX-37208 Inserts = 564 Updates = 0 Deletes = 465
PWX-37121 CPU time = 0 hrs, 0 mins, 0 secs, 299809 mcrs
PWX-37202 Task ID = 42414 AM = NRDB
PWX-37203 Rows read = 10 Rows written = 0
PWX-37204 Bytes read = 570 Bytes written = 0
PWX-37205 File = ABC.VSAM.MASTER_REC
PWX-37206 Table = <Non-relational source>
PWX-37202 Task ID = 42414 AM = KSDS
PWX-37203 Rows read = 10 Rows written = 0
PWX-37204 Bytes read = 800 Bytes written = 0
PWX-37205 File = XYQ.TEST.V1.KSDS
PWX-37206 Table = XYQ.TEST.V1.KSDS
PWX-37121 CPU time = 0 hrs, 0 mins, 0 secs, 76151 mcrs
```

For the CAPXRT and CAPX access methods, the report includes the number of SQL inserts, updates, and deletes that the task processed for a CDC request.

A client task can have multiple access methods, for example, one for reading source data and one for mapping nonrelational source data to a relational format. In the example output, task 42414 uses the NRDB access method with the data map file specified in the **File** field to map nonrelational data to a relational format. The same task uses the KSDS access method to retrieve data from the KSDS data set specified in the **File** field.

The following example clients report is for a PowerExchange Listener on Windows, but the same fields are displayed for a PowerExchange Listener on i5/OS, Linux, zLinux, UNIX, or z/OS:

```
PWX-00723 Command <displaystats Clients> succeeded
PWX-37112 Active Tasks
PWX-37113 Task ID = 41942 Status = Active
PWX-37114 Port = 2480 Partner = 127.0.0.1
PWX-37115 PwrCntrSess = N/A
PWX-37207 Application = N/A
PWX-37116 AM = NRDB Mode = Read Process = DTLST3 SessionId =
PWX-37121 CPU time = 0 hrs, 0 mins, 0 secs, 62400 mcrs
PWX-37122 Start time = 2014-05-01 14:21:37
PWX-37113 Task ID = 41943 Status = Active
PWX-37114 Port = 2480 Partner = 127.0.0.1
PWX-37115 PwrCntrSess = N/A
PWX-37207 Application = N/A
PWX-37116 AM = NRDB Mode = Read Process = DTLST3 SessionId =
PWX-37121 CPU time = 0 hrs, 0 mins, 0 secs, 124800 mcrs
PWX-37122 Start time = 2014-05-01 14:22:01
```

The **Partner** field displays the IP address of the client that issued the request that caused the PowerExchange Listener to create the task. This value begins with ::ffff for an IPv6 address.

For more information about the fields in each of these reports, see the *PowerExchange Command Reference*.

## PowerExchange Logger for Linux, UNIX, and Windows Monitoring Statistics

You can use the PowerExchange Logger DL and DG commands or the `pwxcmd displaystats -tp {logger|groups}` command to publish monitoring statistics for a PowerExchange Logger process and its tasks or PowerExchange Logger group definitions.

Before you run one of these commands, you must configure the `STATS=(MONITOR)` parameter in the PowerExchange Logger `pwxccl.cfg` configuration file to enable collection of the statistics. In this parameter, you can include the optional *interval* subparameter to also publish the statistics at a regular interval.

Issue the commands in the following ways:

- Issue the DL and DG commands from the command line window on the Linux, UNIX, or Windows system where the PowerExchange Logger runs. The PowerExchange Logger must be running in the foreground.
- Issue the `pwxcmd displaystats -tp logger` or `pwxcmd displaystats -tp groups` command from a Linux, UNIX, or Windows system to the PowerExchange Logger on a remote system or the same system. You must use this method to issue the command to a PowerExchange Logger process that runs in background mode.

The command output is displayed on screen and printed to the PowerExchange message log.

### Logger Report

The DL and `pwxcmd displaystats -tp logger` commands produce statistics for the PowerExchange Logger process and its tasks. The following example report shows these statistics:

```
PWX-26011 Command handler received command "DS"
PWX-00723 Command <display L stats> succeeded
PWX-37130   PWXCCL pid = 7144           Writer status = Reading or waiting for source data
PWX-37134   CPU Time =      0:00:02.589616
PWX-37131   Memory (Current/Total/Maximum)
PWX-37132   Controller: (981/983/1849) KB   Command Handler: (0/0/34) KB   Writer: (5127/5147/5181)
KB
PWX-37135   Status 7144                  Totals  I=0000000024344 U=0000000000000 D=0000000024336
C=000000004004 Total=0000000052684
PWX-37136   CurrFileOpened : 2015-08-11 13:20:39 I=0000000024344 U=0000000000000 D=0000000024336
C=000000004004 Total=0000000052684
PWX-37137   Active Cycle : 2015-08-11 13:21:01 I=0000000024344 U=0000000000000 D=0000000024336
C=000000004004 Total=0000000052684
```

This report contains the following fields:

- PWXCCL pid. The process ID of the PowerExchange Logger process.
- Writer status. The status of the PowerExchange Logger Writer subtask at the time the command was issued.
- CPU Time. The amount of CPU time used by the PowerExchange Logger since it started.
- Controller. The amount of memory in kilobytes that the PowerExchange Logger Controller has used.
- Command Handler. The amount of memory in kilobytes that the PowerExchange Logger Command Handler has used.
- Writer. The amount of memory in kilobytes that the PowerExchange Logger Writer subtask has used.
- Status. The process ID of the PowerExchange Logger process.
- I, U, D, C, and Total. Counts of inserts, updates, deletes, and commits that the PowerExchange processed, with the total for all of these operation types. These values are reported for the PowerExchange Logger process, the currently open log file, and the active logging cycle.
- CurrentFileOpened. The timestamp that indicates when the log file was opened.
- Active Cycle. The timestamp that indicates when the active logging cycle started.

### Logger Group Definition Reports

The DG and `pwxcmd displaystats -tp groups` commands produce statistics for each PowerExchange Logger group definition that is definition. A *group definition* defines a set of PowerExchange Logger log files for a group of registered source tables. The following example report shows these statistics:

```
PWX-26011 Command handler received command "DG"
PWX-37138   Grp: dtld004 Regs=1 IUD=0000000000000 C=0000000000000 Unflushed=0000000000000
PWX-37138   Grp: dtld003 Regs=2 IUD=0000000000470 C=0000000000028 Unflushed=0000000000000
PWX-37138   Grp: dtld002 Regs=2 IUD=0000000003276 C=0000000000196 Unflushed=0000000000000
```

This report contains the following fields:

- Grp. The name of the group definition.

- Regs. The number of capture registrations in the group.
- IUD. the total number of inserts, updates, and deletes processed for the group.
- C. The number of commits processed for the group.
- Unflushed. The number of change records for the group that have not yet been flushed to PowerExchange Logger log files on disk.

If no PowerExchange Logger groups are defined, the command reports the following monitoring statistics for the PowerExchange Logger, as if all of the registrations were in one group named "condenseO":

```
PWX-26011 Command handler received command "DG"
PWX-37138 Grp: c:\pwx\capture\condenseO Regs=5 IUD=000000032292 C=000000001931 Unflushed=000000034223
PWX-37139 FirstRec=2015-05-22 13:59:10.603648 Open file=c:\pwx\capture/
condenseO.CND.CP150707.T1816001
PWX-37140 BeginSeq =000000009DE600000000000000000088D800000000 BeginRstrt
=D4C9C7D340400000000037DA00000000
PWX-37141 LastSeq =000001587438000000000000000158728600000000
PWX-37142 CommitSeq=000001589B240000000000000001589B2400000000
CommitRstrt=D4C9C7D340400000000037DA00000000
```

- FirstRec. The timestamp of the first record in the open Logger log file.
- BeginSeq. The sequence token of the earliest record in the open Logger log file.
- BeginRstrt. The restart token of the earliest record in the open Logger log file.
- LastSeq. The sequence token of the last change record in the Logger log file that is not followed by a commit record. This value should be greater than the CommitSeq value.
- CommitSeq. The sequence token of the last commit record in the Logger log file.
- CommitRstrt. The restart token of the last commit record in the Logger log file.

For more information about the command syntax, see the *PowerExchange Command Reference*.

## Monitoring CDC Sessions in PowerCenter

In PowerCenter, you can monitor the progress of CDC sessions.

Use the following information to monitor CDC sessions:

- Messages in the session log. PWXPC writes messages to the session log.
- Performance details in Workflow Monitor. If you configure a CDC session to report performance details, you can monitor the progress of the session in Workflow Monitor.

For more information about PowerCenter monitoring options, see the *PowerCenter Performance Tuning Guide*.

## Session Log Messages

You can use messages that PWXPC and PowerCenter write to the session log to monitor the progress of CDC sessions.

When PWXPC flushes change data, PWXPC writes one of the following messages in the session log to indicate the reason for the flush:

```
PWXPC_10081 [INFO] [CDCDispatcher] raising real-time flush with restart tokens
[restart1], [restart2] because the UOW Count [count] is reached

PWXPC_10082 [INFO] [CDCDispatcher] raising real-time flush with restart tokens
[restart1], [restart2] because Real-time Flush Latency [latency] is reached

PWXPC_12128 [INFO] [CDCDispatcher] raising real-time flush with restart tokens
[restart1], [restart2] because the Maximum Rows Per commit [count] is reached
```

You can use the restart tokens in these PWXPC flush messages to monitor the processing of the change data.

For each PWXPC flush message, PowerCenter writes a WRT\_8160 message after committing change data to the targets. This message displays the source-based commit statistics.

## Performance Details in Workflow Monitor

In Workflow Monitor, you can view performance details in the run properties for a CDC session to assess the efficiency of the CDC session and extraction processing.

If session performance is degraded, you can use data in the Performance Counter column to determine the bottleneck.

PWXPC does not store performance details in the repository so you cannot view performance details for previous executions of a CDC session.

To enable the collection of performance details, select **Collect performance data** on the **Properties** tab of the CDC session properties.

When the CDC session runs, PWXPC refreshes performance statistics every 10 seconds.

If you enable a resume recovery strategy for the CDC session, PWXPC displays data for all Performance Counter fields.

The following table describes the Performance Counter fields:

Performance Counter Field	Description
1 PowerExchange CDC Reader Status:	Current status of the PWXPC reader, as indicated by one of the following values: <ul style="list-style-type: none"><li>- <b>No Data To Process.</b> In the last read, PowerExchange did not pass data to PWXPC.</li><li>- <b>Restart Advance.</b> PowerExchange passed restart tokens to PWXPC but did not pass change data.</li><li>- <b>Processing Data.</b> PowerExchange passed change data and restart tokens to PWXPC for processing.</li></ul>
1.1 Time Last Data Row Read	Time, in milliseconds, that PWXPC took to read the data last received from PowerExchange.
1.2 Data Rows In Current Interval	Number of change records received from PowerExchange during the current statistics interval.



Performance Counter Field	Description
1.3 End Packets In Current Interval	Number of UOWs received from PowerExchange during the current statistics interval.
1.4 Data Read Rate In Current Interval (rows/sec)	<p>Number of change records read per second by PowerExchange during the current statistics interval.</p> <p>The value varies, based on the quantity of change data:</p> <ul style="list-style-type: none"> <li>- If PowerExchange is reading large amounts of change data from the change stream, this value is usually large and reflects the maximum PowerExchange throughput.</li> <li>- If PowerExchange is waiting for change data at the end of the change stream, this value is small.</li> </ul> <p>The following factors can increase this value:</p> <ul style="list-style-type: none"> <li>- Large network bandwidth</li> <li>- CDC offload processing</li> <li>- Multithreaded processing</li> </ul>
1.5 Mean Data Read Rate (rows/sec)	Mean number of change records that PowerExchange read per second, from the start of the CDC session.
1.6 Max Data Read Rate (rows/sec)	Maximum number of change records that PowerExchange read per second during a statistics interval, from the start of the CDC session.
2 PowerCenter Processing Status:	<p>Overall status of the CDC session, as indicated by one of the following values:</p> <ul style="list-style-type: none"> <li>- <b>Idle.</b> Waiting for change data.</li> <li>- <b>Processing Data.</b> Data is being processed.</li> <li>- <b>Recovery Disabled.</b> If a resume recovery strategy is not enabled, the PWXPC CDC reader cannot get PowerCenter status information.</li> </ul>
2.1 Time Of Last Commit	Time stamp of the last commit to a target.
2.2 Rows Processed To Commit In Current Interval	Number of change records that the PWXPC reader flushed during the current statistics interval. This count includes the change records in all committed UOWs. Some of these UOWs might have started before the current statistics interval began.
2.3 Commit Rate In Current Interval (rows/sec)	<p>Processing rate, in number of change records per second, for the change records for the UOW that was last committed during the current statistics interval. This processing includes reading the UOW from PowerExchange and committing the change data to the targets.</p> <p>The following factors can affect this rate:</p> <ul style="list-style-type: none"> <li>- Number of available DTM buffers</li> <li>- Responsiveness of the target</li> <li>- Number of transformations in the pipeline</li> </ul>
2.4 Mean Commit Rate (rows/sec)	<p>Mean number of change records per second for the rate displayed in <b>2.3 Commit Rate In The Current Interval</b>.</p> <p>This value differs from the <b>2.6 Mean Throughput Rate</b> value in that it takes into account only the time when the session is actively processing data. This value does not reflect processing overlap in PowerCenter.</p>
2.5 Max Commit Rate (rows/sec)	Maximum number of change records per second for the commit rate displayed in <b>2.3 Commit Rate In The Current Interval</b> , from the start of the CDC session.

Performance Counter Field	Description
2.6 Mean Throughput (rows/sec)	Mean rate of processing for the CDC session.
2.7 Max Throughput (rows/sec)	Maximum throughput for the CDC session.
2.8 Commits In Current Interval	Number of commits processed to completion by the target during the current statistics interval.
2.9 Commits Pending	Number of commits that the PWXPC reader issued but that have not yet reached the targets. A large value might indicate problems with target responsiveness.
3 Capture Timestamps	-
3.1 Timestamp On Last End Packet Read	The capture timestamp, DTL__CAPXTIMESTAMP, from the last UOW read for a source in the CDC session.
3.2 Timestamp On Last Target Commit	The capture timestamp, DTL__CAPXTIMESTAMP, from the last UOW committed to the target.
4 Totals	-
4.1 Elapsed Time	Total elapsed time for the CDC session.
4.2 Rows Read	Total number of change records read from PowerExchange.
4.3 End Packets Read	Total number of UOWs read.
4.4 Time in PowerExchange Processing	Total time of PowerExchange processing for the CDC session.
4.5 Rows Processed	Total number of change records processed through PowerCenter and committed to the targets.
4.6 Commits to Target	Total number of flushes that the PWXPC reader issued and that were committed to the targets.
4.7 TS on Last Commit minus TS at Commit (2.1-3.2)	Result from subtracting <b>3.2 Timestamp On Last Target Commit</b> value from the <b>2.1 Time Of Last Commit</b> value. If this result is negative, the value is enclosed in parentheses.

## Viewing Performance Details in Workflow Monitor

In Workflow Monitor, view performance details for a CDC session to assess the efficiency of the CDC sessions.

1. In Workflow Monitor, right-click a session and select **Get Run Properties**.
2. In the **Properties** window, click the **Performance** area.

The **Performance Counter** column displays a source qualifier from the CDC session. The **Counter Value** column displays the PowerCenter node name.

3. To view performance details, select the source qualifier.

**Note:** For PWXPC to display performance details for a CDC session that ended, you must select performance details while the session is still running.

PowerCenter displays data in the **Performance Counter** fields in the **Performance** area.

## CHAPTER 13

# Tuning CDC Sessions

This chapter includes the following topics:

- [Tuning Overview, 268](#)
- [PowerExchange DBMOVER Statements for Tuning CDC Sessions, 269](#)
- [PowerCenter Connection Attributes and Session Properties, 272](#)
- [CDC Offload Processing, 275](#)
- [Multithreaded Processing, 277](#)

## Tuning Overview

PowerExchange and PowerCenter provide options that you can use to tune CDC sessions. These tuning options can help you increase throughput, reduce overhead on the source system, and improve CDC efficiency.

Use any of the following options to tune CDC sessions:

- PowerExchange DBMOVER statements. Customize certain statements in the DBMOVER configuration file to make tuning adjustments such as changing buffer sizes or disabling compression or traces.
- PowerCenter connection attributes. Customize PWX CDC application connection attributes to make tuning adjustments such as disabling encryption or compression, reducing commit processing, or enabling offload processing and multithreaded processing.
- Buffer memory. Set the PowerCenter **DTM Buffer Size** and **Default Buffer Block Size** session properties to generate a large number of small blocks. For CDC, this strategy improves session performance and prevents wasted buffer space.
- Offload processing. Use offload processing to transfer column-level extraction processing from the PowerExchange Listener on the source system to the PowerExchange client on the PowerCenter Integration Service machine. Also, if the data source type requires use of the UOW Cleanser (UOWC), offloading transfers UOWC processing to the PowerCenter Integration Service machine. Offloading helps increase throughput when resources available for the PowerExchange Listener are constrained on the source system.
- Multithreaded processing. Enable the use of multiple worker threads for resource-intensive, column-level extraction processing. You can use multithreading on the source system to process data from Linux, UNIX, or Windows data sources if the PWX connection for the CDC session has a location of local. You can also use multithreading for extracting change data from the systems other than the source system when offload processing is in effect. Enable multithreading only when extractions appear to be CPU bound.

**Note:** You can also log data to a PowerExchange Logger for Linux, UNIX, and Windows instance on a system that is remote from the source system. In certain situations, this configuration can reduce resource consumption on the source system, move column-level and UOW Cleanser processing to the remote system, and reduce the network overhead of data transfer. For more information, see [Chapter 8, “Remote Logging of Data” on page 188](#).

#### RELATED TOPICS:

- [“PowerCenter Connection Attributes for Tuning CDC Sessions ” on page 272](#)
- [“PowerExchange DBMOVER Statements for Tuning CDC Sessions” on page 269](#)
- [“Tuning Commit Processing” on page 274](#)

## PowerExchange DBMOVER Statements for Tuning CDC Sessions

You can customize certain statements in the dbmover.cfg configuration file to tune CDC sessions.

Customize any of the following parameters to try to increase throughput or reduce CPU use:

#### **APPBUFSIZE=bytes**

The maximum application data buffer size, in bytes, that PowerExchange uses to read or write data. This buffer type can exist on a source or target system.

If you use a remote target system, PowerExchange usually writes change data to its application data buffer on the source system until the buffer is full. PowerExchange then sends the change data to a sending TCP/IP buffer on the source system. TCP/IP transports the change data to a receiving TCP/IP buffer on the target system. PowerExchange on the target system reads the change data from the TCP/IP buffer into its application data buffer. PWXPC then reads the change data and passes it to PowerCenter. PowerCenter processes the data and applies it to the targets.

Enter an APPBUFSIZE value that is greater than the maximum size of any single data row to be sent.

Valid values are from 34816 through 8388608. Default is 256000.

If the target is remote, enter the same APPBUFSIZE value in the DBMOVER configuration files on the source and target systems.

When the APPBUFSIZE value is not optimal, PowerExchange writes message PWX-01295 in the PowerExchange message log file on the source system. This message recommends a minimum application buffer size.

If dynamic application buffer sizing is enabled, the APPBUFSIZE statement defines the initial size of the application data buffer for all connections made during a PowerExchange Listener run. PowerExchange resizes the application data buffer dynamically for individual connections as needed. Dynamic application buffer sizing is enabled by default. You can explicitly enable it by specifying Y for the APPBUFSIZEDYN statement in the DBMOVER configuration file.

#### **APPBUFSIZEDYN={N|Y}**

Specifies whether to enable dynamic application buffer sizing.

The DBMOVER APPBUFSIZE statement defines the initial size of the application buffer for all connections made during a PowerExchange Listener run. If APPBUFSIZEDYN=Y, PowerExchange resizes the application buffers for individual connection as needed.

The APPBUFSIZEDYN statement applies to PowerExchange connections to data sources with either fixed-length or variable-length records. A variable-length record is a record with at least one variable-length field. A variable-length field has a datatype of VARCHAR or VARBIN.

For each connection to a data source with variable-length records, PowerExchange resizes the application buffer when it encounters a record that is too large to fit into the buffer. PowerExchange increases the size of the application buffer to a value of ten times the size of the record that has overflowed, up to the maximum application buffer size of 8 MB. The new size remains in effect for the duration of the Listener run or until the application buffer is resized again. PowerExchange never decreases the application buffer size for a connection after the Listener run has started.

For each connection to a data source with fixed-length records, PowerExchange determines the record length when the connection is opened and resizes the application buffer once, up to the maximum application buffer size of 8 MB, as needed.

**CAPI\_CONNECTION=( ..., (TYPE={UDB|UOWC}, MEMCACHE=cache\_size, ...))**

The maximum memory cache size, in kilobytes, that PowerExchange can allocate to reconstruct complete UOWs. This MEMCACHE parameter is specified only in the UDB or UOWC CAPI\_CONNECTION statements.

Enter a number from 0 through 2147483647. Default is 1024. If you enter 0, the memory cache size is unlimited.

PowerExchange keeps all changes in each UOW in cache until processing the end-UOW record. PowerExchange incrementally allocates memory cache up to the limit that this parameter specifies. If the MEMCACHE value is too small to hold all of the changes in a UOW in cache, the changes spill to a disk file.

Each UOW spill file contains one UOW. A UOW might require multiple UOW spill files to hold all of the changes for that UOW. If the change stream contains multiple large UOWs and the memory cache is insufficient, PowerExchange might create numerous UOW spill files.

PowerExchange processes the change stream more efficiently if it does not need to use UOW spill files. In addition to degrading extraction performance, large numbers of UOW spill files can cause a disk space shortage.

The default value of 1024 is appropriate if the change stream contains many small UOWs. If you have UOWs larger than 1024 KB, increase this value or enter 0. PowerExchange processes a UOW more efficiently if all of the changes are cached in memory. For most environments, a good starting value is 10240.

**Attention:** PowerExchange allocates memory cache for each connection for change data extraction processing. To prevent excessive memory use, use a MEMCACHE value that is reasonable for the extraction processing load and number of CDC sessions that run concurrently. If the value is too large and you run many concurrent sessions, memory constraints might occur.

**CAPI\_CONNECTION=( ..., (TYPE={MSQL|UDB|UOWC}, RSTRADV=rstr\_seconds, ...))**

Time interval, in seconds, that PowerExchange waits before advancing restart and sequence tokens for a registered data source during periods when UOWs do not include any changes of interest for the data source. When the wait interval expires, PowerExchange returns the next committed "empty UOW," which includes only updated restart information.

This RSTRADV parameter is specified only in CAPI\_CONNECTION statements of the following types:

- MSQL
- UDB
- UOWC

Enter a number from 0 through 86400.

If you do not specify RSTRADV, PowerExchange does not advance restart and sequence tokens for a registered source during periods when PowerExchange receives no changes of interest. In this case, when PowerExchange warm starts, it reads all changes, including those not of interest for CDC, from the restart point.

PowerExchange resets the wait interval to 0 when one of the following events occur:

- PowerExchange completes processing a UOW that includes changes of interest.
- PowerExchange returns an empty UOW because the wait interval expired without PowerExchange receiving any changes of interest.

For sources with low change activity, you can use the RSTRADV parameter to periodically advance to the restart tokens for those sources. Advancing the restart tokens speeds up restart processing for CDC sessions by minimizing the amount of change data that must be reprocessed.

For example, if you specify 5, PowerExchange waits 5 seconds after it completes processing the last UOW or after the previous wait interval expires. Then PowerExchange returns the next committed empty UOW that includes the updated restart information and resets the wait interval to 0.

A low value can cause the **UOW Count** option on the PWX CDC connection to match more quickly than expected. When the UOW counter matches, PWXPC flushes the data buffer and commits restart tokens to the targets. Excessive flush activity can adversely affect performance on the PowerCenter Integration Service machine and on the target databases.

**Attention:** A value of 0 can degrade performance. In addition to the UOWs that contain changes for registered sources of interest, PowerExchange returns an empty UOW for every UOW that does not contain changes for the registered sources of interest.

**LISTENER=(node\_name,TCPIP,port,send\_bufsize,receive\_bufsize,send\_size,receive\_size, ...)**

A TCP/IP port on which a named PowerExchange Listener process listens for work requests.

The *send\_bufsize* and *receive\_bufsize* positional parameters define the data portion of the TCP/IP send and receive buffer sizes that PowerExchange uses. If you do not specify these values, PowerExchange uses the operating system defaults.

To increase throughput, try increasing the *send\_bufsize* and *receive\_bufsize* values in the LISTENER statement in the DBMOVER configuration file on the source system. For help in determining the best values to use, contact your network administrator.

**NODE=(node\_name,TCPIP,host\_name,port,send\_bufsize,receive\_bufsize,send\_size,receive\_size, ...)**

A TCP/IP host name and port that PowerExchange uses to contact a PowerExchange Listener process.

The *send\_bufsize* and *receive\_bufsize* positional parameters define the data portion of the send and receive buffer sizes that PowerExchange uses. If you do not specify these values, PowerExchange uses the operating system defaults.

To increase throughput, try increasing the *send\_bufsize* and *receive\_bufsize* values in the NODE statement in the DBMOVER configuration file on the target system. For help in determining the best values to use, contact your network administrator.

**TRACE=(trace\_id,trace\_level,99)**

Activates PowerExchange diagnostic traces that Informatica Global Customer Support uses to solve problems with PowerExchange code.

TRACE statements can severely impact PowerExchange performance. Use these statements only at the direction of Informatica Global Customer Support.

After Informatica Global Customer Support diagnoses the problem, remove or comment out all TRACE statements in the DBMOVER configuration files on all systems.

For more information about these DBMOVER statements, see the *PowerExchange Reference Manual*.

## PowerCenter Connection Attributes and Session Properties

You can use certain PowerCenter connection attributes and session properties for tuning CDC sessions.

### PowerCenter Connection Attributes for Tuning CDC Sessions

In PowerCenter, you can customize some attributes on the PWX CDC connections to tune CDC sessions.

The following table describes the connection attributes that you can optionally use for tuning:

Connection Option	Description	Tuning Suggestion
Compression	Controls whether to compress source data during the PowerCenter session. Default disables compression.	Do not use compression.
Encryption Type	Type of data encryption that PowerExchange uses. Default is None for no encryption.	Do not use encryption.



Connection Option	Description	Tuning Suggestion
Image Type	<p>Indicates how PWXPC passes captured Updates to CDC sessions that extract and apply the updates to the target.</p> <p>Options are:</p> <ul style="list-style-type: none"> <li>- <b>AI.</b> Process Updates as Update operations. PWXPC passes each Update as a single Update record. An Update record includes after images of the data only, unless you add before image (BI) and change indicator (CI) fields to the extraction map that you import for the source definition for the CDC session.</li> <li>- <b>BA.</b> Process Updates as Deletes followed by Inserts. PWXPC passes each Update as a Delete record followed by an Insert record. The Delete record contains the before image of the data, and the Insert record contains the after image.</li> </ul> <p>Default is BA.</p> <p>If you specify AI, you can still use before images of the data, if available, in extraction processing. PWXPC can embed before-image data and after-image data in the same Update row. To embed before-image data, you must complete the following configuration tasks:</p> <ul style="list-style-type: none"> <li>- In the PowerExchange Navigator, add BI and CI fields to the extraction map that you plan to import for the source definition in PowerCenter.</li> <li>- If you use batch or continuous extraction mode, enter BA for the CAPT_IMAGE parameter in the PowerExchange Condense or PowerExchange Logger for Linux, UNIX, and Windows configuration file. This setting stores both before and after images in the PowerExchange Logger log files or PowerExchange Condense condense files. When CDC sessions run, they extract data from these files.</li> </ul>	Set to AI.
UOW Count	<p>The number of UOWs that PWXPC reads from the source before it flushes the data buffer to commit the change data to the targets.</p> <p>Default is 1.</p>	To improve efficiency on the PowerCenter Integration Service machine and the target databases, increase this value to reduce commit processing.
Real-time Flush Latency in milliseconds	<p>The frequency, in milliseconds, with which PWXPC flushes the data buffer to commit the change data to the targets.</p> <p>Default is 0, which is equivalent to 2 seconds.</p>	To improve efficiency on the PowerCenter Integration Service machine and the target databases, increase this value to reduce commit processing.
PWX Latency in seconds	<p>Maximum time, in seconds, that the PowerExchange instance on the source waits for more change data before flushing data to PWXPC on the PowerCenter Integration Service machine.</p> <p>Default is 2.</p>	Use the default value.

Connection Option	Description	Tuning Suggestion
Minimum Rows Per commit	Minimum number of change records that PowerExchange reads from the change stream before it passes any commit records to PWXPC. Default is 0, which means that PWXPC ignores this option.	If UOWs usually contain few changes, increase this value to increase the size of the UOWs. This practice can improve efficiency on the PowerCenter Integration Service machine and on the target databases by reducing commit processing.
Offload Processing	Controls whether PowerExchange uses CDC offload processing. Offload processing transfers resource-intensive column-level and UOW Cleanser processing from the source system to another system. Default is No.	If resource constraints exist on the source system and you need to increase CDC throughput, consider enabling offload processing.
Worker Threads	Controls whether PowerExchange uses multiple threads for resource-intensive, column-level extraction processing. You can use multithreading on the source system to process data from Linux, UNIX, or Windows data sources, or on another system for extraction processing when offload processing is in effect. Enable multithreading only when extractions appear to be CPU bound. Enter the number of threads that you want PowerExchange to use. Valid values are 1 through 64. Default is 0, which causes PowerExchange to not use multithreaded processing.	Enter a number greater than 0.
Array Size	If the <b>Worker Threads</b> value is greater than zero, indicates the size of the storage array, in number of records, for the threads. Valid values are from 25 through 100000. Default is 25.	Use 25. <b>Attention:</b> If you enter a large array size value, have large records, or run many sessions that use multithreading concurrently, memory shortages might occur on the PowerCenter Integration Service machine.

For more information about PWX CDC connection attributes, see *PowerExchange Interfaces for PowerCenter*.

## Tuning Commit Processing

To tune commit processing and performance of CDC sessions, you can adjust commitment control attributes on the PWX CDC application connection.

If the session log for a CDC session contains PWXPC flush messages followed by PowerCenter source-based commit messages, the session might be reading change data faster than the data is applied to the targets. To try to resolve this issue, adjust the following commitment control attributes on the PWX CDC connection, based on the most prevalent type of flush message in the session log:

- If PWXPC\_10081 flush messages are the most prevalent messages, try increasing the **UOW Count**.
- If PWXPC\_10082 flush messages are the most prevalent messages, try increasing the **Real-time Flush Latency in milli-seconds**.

If PWXPC flushes change data too frequently, too many commitment control attributes might be specified on the PWX CDC connection. In this case, specify a single commitment control attribute and disable the other ones.

If the change stream contains many small UOWs, you can use the **Minimum Rows Per commit** option to create larger UOWs of more uniform size. PowerExchange and PWXPC can process a few large UOWs more efficiently than many small UOWs. By using the **Minimum Rows Per commit** option to increase the size of UOWs, you can improve CDC processing efficiency.

Also, performance of the target database can impact the performance of the CDC session. Contact your database administrator to verify that database access is optimal.

## PowerCenter Session Properties for Tuning Buffer Memory

When you run a CDC session, the PowerCenter Data Transformation Manager (DTM) allocates buffer memory to the session based on the **DTM Buffer Size** value on the **Properties** tab of the session properties. The DTM divides the memory into buffer blocks based on the **Default Buffer Block Size** setting on the **Config Object** tab of the session properties.

If you suspect that buffer memory is insufficient, enable the collection of performance details in the CDC session. Then review the difference between the performance counters

**4.1 Time in PowerExchange Processing** and **4.4 Elapsed Time**. If the elapsed time is much larger than the PowerExchange processing time, buffer memory constraints might exist. To improve performance of the CDC session, try adjusting the **DTM Buffer Size** and **Default Buffer Block Size** properties.

For optimal CDC performance, set these session properties to create a large number of small blocks. Informatica recommends the following settings:

- For the **DTM Buffer Size**, specify 128 MB, 256 MB, 512 MB, 1 GB, or 2 GB.
- For the **Default Buffer Block Size**, specify 32 KB.

Do *not* set these session properties to **auto**. The **auto** option creates a small number of large blocks, which can degrade CDC session performance. The **auto** option is intended for bulk data load processing.

## CDC Offload Processing

CDC offload processing transfers column-level processing of change data from the PowerExchange Listener on the source system to the PowerExchange client on the PowerCenter Integration Service machine.

For data sources for which PowerExchange uses the UOW Cleanser (UOWC), offload processing also transfers UOWC processing to the PowerCenter Integration Service machine. These data sources include z/OS data sources, DB2 for i5/OS, and Oracle CDC with LogMiner.

Use offload processing when resources on the source system are constrained. In this situation, offload processing can help increase throughput for CDC sessions.

### RELATED TOPICS:

- [“Rules and Guidelines for CDC Offload Processing” on page 276](#)
- [“Example of CDC Offload Processing with an Oracle Source” on page 277](#)
- [“Enabling Offload Processing for CDC Sessions” on page 276](#)

## Rules and Guidelines for CDC Offload Processing

Before implementing CDC offload processing, review the following rules and guidelines:

- You must copy the appropriate source-specific CAPI\_CONNECTION statements from the DBMOVER configuration file on the source system to the PowerCenter Integration Service machine.
- PowerExchange does not support CDC offload processing for capture registrations that you create from data maps that use any of the following options:
  - User access methods
  - User-defined fields that invoke programs by using the CALLPROG function
  - Record-level exits

## Enabling Offload Processing for CDC Sessions

To use CDC offload processing, you must configure some PWX CDC connection attributes. Also, you must add the source-specific CAPI\_CONNECTION statements to the DBMOVER configuration file on the PowerCenter Integration Service machine.

1. Configure attributes for offload processing on the PWX CDC Real Time application connection for the CDC session.

The following table describes the attributes that are required for offload processing:

Connection Attribute	Description
Location	Specifies the node name of the system on which the change data resides. This node name must match the name of a NODE statement in the dbmover.cfg configuration file on the PowerCenter Integration Service machine.
Offload Processing	Controls whether PowerExchange uses CDC offload processing. When offload processing is enabled, PowerExchange transfers column-level processing of the change data and any UOW Cleanser (UOWC) processing from the source system to the PowerCenter Integration Service machine.  Options are: <ul style="list-style-type: none"><li>- <b>No.</b> Disables offload processing.</li><li>- <b>Yes.</b> Enables offload processing.</li><li>- <b>Auto.</b> PowerExchange determines whether to enable or disable offload processing.</li></ul> Default is No.
CAPI Connection Name	Specifies the name of the source CAPI_CONNECTION statement in the dbmover configuration file on the PowerCenter Integration Service machine.

2. Copy the source-specific CAPI\_CONNECTION statements from the dbmover configuration file on the source system to the dbmover configuration file on the PowerCenter Integration Service machine.

The following table identifies the CAPI\_CONNECTION statement types to copy for each Linux, UNIX, and Windows source type:

Source Type	CAPI_CONNECTION Statement Types
DB2 for Linux, UNIX, and Windows	UDB
Microsoft SQL Server	MSQL
Oracle	ORCL and UOWC, for PowerExchange Oracle CDC with LogMiner sources ORAD, for PowerExchange Express CDC for Oracle sources

## Example of CDC Offload Processing with an Oracle Source

In this example, you enable a CDC session with a PWX CDC Real Time connection to use offload processing for change data extraction from an Oracle source. You are using the PowerExchange Oracle CDC with LogMiner method of Oracle change capture.

The source data remains on the Oracle system but all column-level and UOW Cleanser processing is offloaded to the PowerCenter Integration Service machine.

1. Copy the UOWC and ORCL CAPI\_CONNECTION statements from the dbmover configuration file on the Oracle source system to the dbmover configuration file on the PowerCenter Integration Service machine.

This example uses the following CAPI\_CONNECTION statements:

```
CAPI_CONNECTION=(NAME=UOWCORA,TYPE=(UOWC,CAPINAME=CAPIORA,RSTRADV=600))
CAPI_CONNECTION=(NAME=CAPIORA,TYPE=(ORCL,catint=120,ORACOLL=PRODORA))
```

The PowerExchange Listener uses these statements to access change data from the specified Oracle instance.

2. Stop the CDC session.
3. Update the following attributes on the PWX CDC Real Time application connection for the CDC session:
  - For the **Offload Processing** attribute, select **Yes**.
  - For the **CAPI Connection Name** attribute, enter the name of the UOWC CAPI\_CONNECTION statement. In this example, the name is UOWCORA.
4. Restart the CDC session.

## Multithreaded Processing

Multithreaded processing uses multiple worker threads to distribute resource-intensive, column-level processing across multiple CPUs. Use multithreading if a single CPU cannot optimally handle extraction processing.

By default, PWXPC uses a single thread to process change data on the PowerCenter Integration Service machine. When you enable multithreading, PWXPC uses multiple threads to process change records.

## Rules and Guidelines for Multithreaded Processing

Multithreaded processing can help improve performance for CDC sessions in specific situations.

Use the following rule and guidelines to determine when multithreaded processing is useful and how to set the **Worker Threads** attribute:

- Use multithreaded processing when the PWX reader thread of a CDC session uses 100% of a single CPU on a multiple-CPU server on the PowerCenter Integration Service machine. In this situation, multithreading improves throughput by spreading PowerExchange processing across multiple threads. Otherwise, multithreading does not improve throughput.
- For optimal performance, verify that the value of the **Worker Threads** attribute does not exceed the number of installed or available processors on the PowerCenter Integration Service machine.
- When defining the PWX CDC application connection, you must either set the **Location** attribute to "local" to enable the extraction to access the source locally, or set the **Offload Processing** attribute to **Yes** to offload extraction processing.
- If processing slows or hangs for CDC sessions that use multiple worker threads, increase the MAXTASKS value in the DBMOVER configuration file to help improve performance.

## Enabling Multithreaded Processing for CDC Sessions

To use multithreaded processing, you must configure some PWX CDC connection attributes.

The following table describes the PWX CDC Real Time application connection attributes that are required to enable multithreaded processing for a CDC session:

Connection Attribute	Description
Worker Threads	Specifies the number of threads that PowerExchange uses on the PowerCenter Integration Service machine to process change data. Default is 0.
Array Size	If the <b>Worker Threads</b> value is greater than zero, specifies the size of the storage array, in number of records, for each thread. Default is 25.

# APPENDIX A

## DTL\_\_CAPXTIMESTAMP Time Stamps

This appendix includes the following topic:

- [Time Stamps That Are Reported in the DTL\\_\\_CAPXTIMESTAMP Field by Data Source, 279](#)

### Time Stamps That Are Reported in the DTL\_\_CAPXTIMESTAMP Field by Data Source

The time stamp that PowerExchange reports in the generated DTL\_\_CAPXTIMESTAMP field in change records depends on the data source type and on certain parameter settings.

For PowerExchange data sources on z/OS and for PowerExchange Oracle CDC with LogMiner sources, the TIMESTAMP parameter in the UOWC CAPI\_CONNECTION controls the type of time stamp that PowerExchange reports in the DTL\_\_CAPXTIMESTAMP field. If you set the TIMESTAMP parameter to COMMIT, PowerExchange reports the time stamp of the transaction commit on the source for all changes in the transaction. If you use the default parameter value of LOG, PowerExchange retrieves the time stamp from the source database logs. In this case, the time stamp type depends on the source type.

The following table describes the time stamps that PowerExchange reports when you use the default value of LOG for the TIMESTAMP parameter:

Data Source Type	Time Stamp Type
Adabas	The HDDATE time stamp from the PLOG block header, which indicates when the block was written. <b>Note:</b> In Adabas environments with a low level of update activity, the same time stamp might be reported for multiple updates that occurred at different times.
Datacom table-based CDC	The Coordinated Universal Time (UTC) time or local time when the change record was written to the Datacom LXX log. The LOCAL_TIME parameter in the ECCR configuration member, ECCRDCMP, controls whether the UTC or local time is used.
DB2 for i5/OS	An i5/OS journal time stamp that reflects when the change was recorded in the journal.
DB2 for z/OS	The time at which the DB2 ECCR captured the change data record. Each record in a UOW has a different time stamp. Usually, this time stamp is a UTC value that reflects the time zone of the DB2 for z/OS system.

Data Source Type	Time Stamp Type
IDMS	The time at which the change data record was written to the IDMS log file. This time stamp is equivalent to the storeclock (STCK) time stamp. It does not reflect the local time zone.
IMS log-based CDC	The time at which the change was recorded in the IMS logs.
IMS synchronous CDC	The time at which the change occurred.
Oracle CDC with LogMiner	The time stamp of the change on the source database, as recorded in the redo logs. This time reflects the local time zone.
Batch VSAM and CICS/VSAM	The time at which the change record was captured. Each record in a UOW has a different time stamp. Usually, this time stamp is a UTC value.

For other data sources that do not use the UOWC CAPI\_CONNECTION statement, PowerExchange determines the appropriate time stamp to report in the DTL\_\_CAPXTIMESTAMP field. For PowerExchange Express CDC for Oracle sources, the TIME\_STAMP\_MODE parameter in the OPTIONS statement of the Express CDC configuration file controls the time stamp type.

The following table describes the time stamp types that PowerExchange reports for these data sources:

Data Source Type	Time Stamp Type
DB2 for Linux, UNIX, and Windows	The time stamp of the transaction commit. This time stamp is an ascending virtual time stamp (VTS) of the DB2 system, which usually corresponds to the UTC value.
Microsoft SQL Server	The time at which the change was written to the distribution database.
PowerExchange Express CDC for Oracle	<p>The time stamp type is controlled by the TIME_STAMP_MODE parameter setting in the OPTIONS statement of the Express CDC configuration file.</p> <ul style="list-style-type: none"> <li>- If you use the default value of LOGTIME, PowerExchange reports the time stamp of the change on source database, as recorded in the redo logs. This time stamp reflects the local time zone.</li> <li>- If you specify COMMITTIME, PowerExchange reports the time stamp of the transaction commit on the source database.</li> <li>- If you specify BEGINTIME, PowerExchange reports the time stamp of the begin UOW log record.</li> </ul>



# INDEX

## A

- alternative logging [40, 56](#)
- Application Name attribute [228](#)
- application names [238](#)
- architectural diagrams
  - batch or continuous extraction processing [20](#)
  - real-time extraction processing [20](#)
- architecture
  - PowerExchange Express CDC for Oracle [141](#)
- architecture, PowerExchange CDC [20](#)
- archive log destination [159](#)
- ARCHIVELOG mode
  - enabling for Oracle CDC with LogMiner [115](#)
  - enabling for PowerExchange Express CDC for Oracle [159](#)
- ASM
  - creating an ASM user [161](#)
- ASM environments
  - Express CDC configuration considerations [154](#)

## B

- batch extraction mode
  - use for terminating CDC sessions [249](#)
- before indicator (BI) fields
  - use cases [214](#)
- benefits
  - PowerExchange Express CDC for Oracle [140](#)

## C

- CAPL Connection Name Override attribute [226](#)
- CAPL connection statements
  - CAPL\_CONNECTION statement [26](#)
  - CAPL\_SRC\_DFLT statement [26](#)
  - CAPX parameters [57](#)
  - introduction [27](#)
  - MEMCACHE parameter [269](#)
  - MSQL parameters [102](#)
  - ORAD parameters [168](#)
  - ORCL parameters [125](#)
  - RSTRADV parameter [269](#)
  - UDB parameters [85](#)
  - UOWC parameters [132](#)
- CAPL\_CONNECTION - CAPX statement
  - DBMOVER configuration file [57](#)
- CAPL\_CONNECTION - MSQL statement
  - DBMOVER configuration file [102](#)
- CAPL\_CONNECTION - ORAD statement
  - DBMOVER configuration file [168](#)
- CAPL\_CONNECTION - ORCL statement
  - DBMOVER configuration file [125](#)
- CAPL\_CONNECTION - UDB statement
  - DBMOVER configuration file [85](#)

- CAPL\_CONNECTION - UOWC statement
  - DBMOVER configuration file [132](#)
- CAPL\_SRC\_DFLT statement
  - DBMOVER configuration file [28](#)
- CAPT\_PATH statement
  - DBMOVER configuration file [30](#)
- CAPT\_XTRA statement
  - DBMOVER configuration file [30](#)
- capture catalog table
  - creating [83](#)
  - DTLUCUDB SNAPSHOT command [83](#)
  - initializing the table [83](#)
- capture registrations
  - adding another registration for Express CDC for Oracle [185](#)
  - grouping in PowerExchange Logger group definition file [60](#)
  - settings for the PowerExchange Logger [43](#)
- CAPX CAPL\_CONNECTION parameters
  - parameters and syntax [57](#)
- catalog, Oracle
  - copying for Oracle CDC with LogMiner [118](#)
  - parameters in ORCL CAPL\_CONNECTION [131](#)
- CDC sessions
  - adding sources with DTLUAPPL-generated CURRENT\_RESTART tokens [251](#)
  - adding sources with special override CURRENT\_RESTART tokens [251](#)
  - changing and restarting [250](#)
  - cold starting [247](#)
  - commit processing [218](#)
  - default restart points [239](#)
  - defining terminating conditions [249](#)
  - methods of starting [238](#)
  - monitoring in PowerCenter [263](#)
  - monitoring in PowerExchange [256](#)
  - multithreaded processing [277](#)
  - offload processing [275](#)
  - performance details in Workflow Monitor [264](#)
  - PowerExchange Logger for LUW logging of data from remote source [188](#)
  - processing of multiple source definitions [216](#)
  - recovering [252](#)
  - recovery example [253](#)
  - recovery start [248](#)
  - restart points for warm starts [240](#)
  - restart token file [243](#)
  - session and connection attributes for CDC [224](#)
  - start methods [246](#)
  - stop command processing [249](#)
  - stopping [248](#)
  - tuning [268](#)
  - tuning buffer memory [274](#)
  - tuning overview [219](#)
  - warm starting [247](#)
- CDCT file [37, 72, 75](#)
- change data capture (CDC)
  - architecture [20](#)

- change data capture (CDC) (*continued*)
  - data source types [16](#)
  - DB2 for Linux, UNIX, and Windows CDC [77](#)
  - Oracle CDC with LogMiner [109](#)
  - overview [14](#)
  - PowerExchange components [19](#)
  - SQL Server CDC [94](#)
  - task summary [24](#)
- change data extraction
  - connection attributes for Logger for LUW log files remote from the source [199](#)
  - creating restart tokens [241](#)
  - extraction modes [209](#)
  - monitoring in PowerCenter [263](#)
  - monitoring in PowerExchange [256](#)
  - multithreaded processing [277](#)
  - offload processing [275](#)
  - overview [15](#), [208](#)
  - overview of extracting change data [220](#)
  - task flow [222](#)
  - testing extraction maps [222](#)
  - tuning CDC sessions [268](#)
- change indicator (CI) fields [214](#)
- close command (pwxcmd) [33](#)
- closeforce command (pwxcmd) [33](#)
- cold starts
  - CDC sessions [247](#)
  - determining restart tokens [239](#)
- commit processing
  - commitment control attributes [230](#)
  - examples [233](#)
  - in CDC sessions [218](#)
  - target latency [230](#)
  - tuning [274](#)
- components
  - PowerExchange Express CDC for Oracle [141](#)
- components, PowerExchange
  - for CDC [19](#)
  - PowerExchange Listener [19](#), [26](#)
  - PowerExchange Logger [19](#)
  - PowerExchange Navigator [20](#)
- configuration in Oracle
  - creating an Oracle user and granting privileges [160](#)
  - enabling ARCHIVELOG mode for PowerExchange Express CDC for Oracle [159](#)
  - enabling minimal global supplemental logging for Express CDC [162](#)
  - specifying an archive log destination for PowerExchange Express CDC for Oracle [159](#)
- configuration in PowerExchange
  - configuration tasks [162](#)
  - customizing the Express CDC configuration file [173](#)
  - example dbmover.cfg configuration file [172](#)
- configuration script files, Oracle CDC with LogMiner
  - Copy\_Oracle\_Dictionary.sql script file [114](#)
  - Create\_ORACAPT\_user.sql script file [114](#)
  - DataGuard\_Logon\_Trigger.sql script file [114](#)
  - Optional\_GRANT\_for\_registrations.sql script file [114](#)
  - readme\_oracapt.txt file [114](#)
  - Setup\_Archive\_Logging.sql script file [114](#)
  - Setup\_Logminer\_Tablespace.sql script file [114](#)
- configuration tasks
  - DB2 for Linux, UNIX, and Windows CDC [81](#), [82](#)
  - Oracle CDC with LogMiner [114](#)
  - PowerExchange Listener [26](#)
  - PowerExchange Logger [43](#)
  - SQL Server CDC [99](#)
- connection attributes
  - Application Name [228](#)

- connection attributes (*continued*)
  - attributes to set for CDC [224](#)
  - CAPI Connection Name Override [226](#)
  - commitment control attributes [230](#)
  - Event Table [229](#)
  - Idle Time [226](#)
  - Image Type attribute [225](#)
  - PWX Latency in seconds [229](#)
  - restart control attributes [228](#)
  - RestartToken File Folder [228](#)
  - RestartToken File Name [228](#)
- continuous extraction mode [209](#)

## D

- data maps
  - use in DB2 for Linux, UNIX, and Windows CDC [89](#)
- data sources, types [16](#)
- database row tests
  - testing data access with an extraction map [222](#)
- DATABASE statement
  - PowerExchange Express CDC for Oracle configuration file [174](#)
- datatypes
  - DB2 for Linux, UNIX, and Windows [79](#)
  - SQL Server [96](#)
- datatypes, Oracle
  - supported for PowerExchange Express CDC for Oracle [148](#)
- DB2 for Linux, UNIX, and Windows CDC
  - changing a source table definition [90](#)
  - configuring in DB2 [81](#)
  - configuring in PowerExchange with the Logger [82](#)
  - configuring in PowerExchange without the Logger [82](#)
  - creating the capture catalog table [83](#)
  - datatypes supported [79](#)
  - dbmover.cfg parameters [84](#)
  - example dbmover.cfg statements [84](#)
  - IBM APARs [92](#)
  - initializing the capture catalog table [83](#)
  - overview [77](#)
  - planning [78](#)
  - planning considerations [80](#)
  - prerequisites [78](#)
  - stopping [90](#)
  - troubleshooting [92](#)
  - user authority requirement [78](#)
  - using a data map [89](#)
- DB2 partitioned databases
  - adding or dropping partitions [91](#)
  - reconfiguration tasks [91](#)
  - reconfiguring a database partition group [92](#)
- DB2 SQL1224 error [92](#)
- DB2CODEPAGE environment variable [81](#)
- DB2NOEXITLIST environment variable [81](#)
- DBMOVER configuration file
  - APBUBFSIZE statement [269](#)
  - CAPI\_CONNECTION statements [26](#)
  - CAPI\_SRC\_DFLT statement [26](#)
  - CAPT\_PATH parameter [56](#)
  - CAPT\_PATH statement [26](#)
  - CAPT\_XTRA statement [26](#)
  - DB2 for Linux, UNIX, and Windows CDC parameters [84](#)
  - DB2 for Linux, UNIX, and Windows example statements [84](#)
  - example for capture system [172](#)
  - general CDC parameters [26](#)
  - LOGPATH parameter [56](#)
  - NODE and LISTENER buffer size parameters [269](#)
  - Oracle CDC with LogMiner example statements [123](#)

DBMOVER configuration file (*continued*)

PowerExchange Logger for LUW logging of remote source data [198](#), [199](#)

PowerExchange Logger parameters [56](#)

SQL Server CDC example statements [101](#)

SQL Server CDC parameters [101](#)

SVCNODE parameter [56](#)

TRACE statement [269](#)

TRACING parameter [56](#)

types of CAPI connection statements for CDC [27](#)

DBMOVER statements

CAPI\_CONNECTION - CAPX [57](#)

CAPI\_CONNECTION - MSQL [102](#)

CAPI\_CONNECTION - ORAD [168](#)

CAPI\_CONNECTION - ORCL [125](#)

CAPI\_CONNECTION - UDB [85](#)

CAPI\_CONNECTION - UOWC [132](#)

CAPI\_SRC\_DFLT [28](#)

CAPT\_PATH [30](#)

CAPT\_XTRA [30](#)

key statements for PowerExchange Express CDC for Oracle [164](#)

Oracle CDC with LogMiner statements [122](#)

ORACLE\_CAPTURE\_TYPE statement [168](#)

ORACLE\_UNHANDLED\_NUMASCHAR statement [170](#)

ORACLEID [123](#), [165](#)

detail.log [40](#)

diagrams

batch or continuous extraction processing [20](#)

real-time extraction processing [20](#)

DICTIONARY statement

PowerExchange Express CDC for Oracle configuration file [176](#)

DIRSUB statement

PowerExchange Express CDC for Oracle configuration file [177](#)

DISPLAY ACTIVE command [259](#)

DTL\_\_BI\_columnname column

described [210](#)

DTL\_\_CAPXACTION

described [210](#)

DTL\_\_CAPXCASDELIND

described [210](#)

DTL\_\_CAPXRESTART1 column

described [210](#)

displaying sequence token [241](#)

DTL\_\_CAPXRESTART2 column

described [210](#)

displaying restart token [241](#)

DTL\_\_CAPXROWID column

described [210](#)

DTL\_\_CAPXRRN column

described [210](#)

DTL\_\_CAPXTIMESTAMP column

described [210](#)

DTL\_\_CAPXTIMESTAMP field

types of reported time stamps by data source [279](#)

DTL\_\_CAPXUOW column

described [210](#)

DTL\_\_CAPXUSER column

described [210](#)

DTL\_\_CI\_columnname column

described [210](#)

DTLUAPPL utility

displaying restart tokens in generated columns [241](#)

DTLUCUDB SNAPSHOT command [83](#)

DTLUTSK utility

stopping CDC sessions [248](#)

## E

environment information checklist [145](#)

Event Table attribute [229](#)

event table processing

guidelines for using [228](#)

implementing [229](#)

use for terminating CDC sessions [249](#)

extraction map columns, PowerExchange-generated

DTL\_\_BI\_columnname [210](#)

DTL\_\_CAPXACTION [210](#)

DTL\_\_CAPXCASDELIND [210](#)

DTL\_\_CAPXRESTART1 [210](#)

DTL\_\_CAPXRESTART2 [210](#)

DTL\_\_CAPXROWID [210](#)

DTL\_\_CAPXRRN [210](#)

DTL\_\_CAPXTIMESTAMP [210](#)

DTL\_\_CAPXUOW [210](#)

DTL\_\_CAPXUSER [210](#)

DTL\_\_CI\_columnname [210](#)

extraction maps

BI and CI fields [214](#)

PowerExchange-generated columns [210](#)

extraction modes [209](#)

extraction of change data

connection attributes for Logger for LUW log files remote from the source [199](#)

creating restart tokens [241](#)

extraction modes [209](#)

monitoring in PowerCenter [263](#)

monitoring in PowerExchange [256](#)

multithreaded processing [277](#)

offload processing [275](#)

overview [208](#)

overview of extracting change data [220](#)

task flow [222](#)

testing extraction maps [222](#)

tuning CDC sessions [268](#)

## F

file switches

description [40](#)

FILESWITCH command [68](#)

flush latency [229](#), [230](#)

## G

group definition file

configuring for PowerExchange Logger [60](#)

example file [62](#)

GROUP statement [61](#)

REG statement [61](#)

SCHEMA statement [61](#)

statements and parameters [61](#)

## I

Idle Time attribute

use for terminating CDC sessions [249](#)

Image Type attribute [225](#)

integration with PowerCenter [22](#)

## L

- listtask (pwxcmd) [33](#), [259](#)
- LISTTASK command [259](#)
- lock files [39](#)
- log files of PowerExchange Logger
  - file switches [40](#)
- log files, PowerExchange Logger
  - maintaining with PWXUCDCT commands [72](#)
  - name format [38](#)
- LogMiner, Oracle
  - configuring for Oracle CDC [116](#)

## M

- management and maintenance tasks
  - adding a capture registration [185](#)
  - changing the structure of an Oracle table [186](#)
  - monitoring Express CDC for Oracle read progress [184](#)
  - overview for PowerExchange Express CDC for Oracle [184](#)
  - stopping CDC processing for a table [186](#)
  - stopping CDC processing temporarily [186](#)
- Maximum Rows Per commit attribute [230](#)
- message log files [40](#)
- Microsoft SQL Server CDC
  - changing a source table definition [106](#)
  - configuration tasks [99](#)
  - configuring in PowerExchange with the Logger [100](#)
  - configuring in PowerExchange without the Logger [100](#)
  - datatypes supported [96](#)
  - dbmover.cfg parameters [101](#)
  - example dbmover.cfg statements [101](#)
  - operational considerations [97](#)
  - overview [94](#)
  - planning [95](#)
  - prerequisites [95](#)
  - stopping [106](#)
  - user authority requirements [95](#)
- minimal global supplemental logging
  - enabling for Express CDC for Oracle [162](#)
- Minimum Row Per commit attribute [230](#)
- monitoring CDC sessions
  - methods [256](#)
  - performance details in Workflow Monitor [264](#)
  - PowerCenter [263](#)
  - PowerCenter session log messages [264](#)
  - PowerExchange extraction statistics messages [257](#)
  - PowerExchange multithreaded processing statistics [258](#)
  - PowerExchange read progress messages [257](#)
  - viewing performance details in PowerCenter [266](#)
- monitoring Express CDC for Oracle read progress [184](#)
- MSQL CAPI\_CONNECTION parameters
  - parameters and syntax [102](#)
- multiple-source processing
  - in CDC sessions [216](#)
- multithreaded processing
  - guidelines for usage [278](#)
  - overview [219](#), [277](#)
  - statistics messages [258](#)

## O

- offload processing
  - enabling for CDC sessions [276](#)
  - example of offload processing [277](#)
  - overview [219](#), [275](#)

- offload processing (*continued*)
  - rules and guidelines [276](#)
- operational considerations
  - PowerExchange Express CDC for Oracle [150](#)
- OPTIONS statement
  - PowerExchange Express CDC for Oracle configuration file [177](#)
- Oracle CDC with LogMiner
  - changing a source table definition [138](#)
  - configuration in a RAC environment [119](#)
  - configuration script files [114](#)
  - configuring Oracle for CDC [114](#)
  - configuring Oracle LogMiner [116](#)
  - configuring PowerExchange with the Logger [121](#)
  - configuring PowerExchange without the Logger [120](#)
  - copying the Oracle catalog [118](#)
  - Data Guard logical standby database as source [119](#)
  - DBMOVER statements [122](#)
  - enabling ARCHIVELOG mode [115](#)
  - example dbmover statements [123](#)
  - implementation considerations [110](#)
  - overview [109](#)
  - performance considerations [114](#)
  - planning [110](#)
  - restrictions [112](#)
  - stopping CDC processing for selected tables [137](#)
  - stopping CDC processing temporarily [138](#)
  - supplemental logging requirement [115](#)
  - supported datatypes [113](#)
  - user privileges required [117](#)
- Oracle configuration
  - creating an Oracle user and granting privileges [160](#)
  - enabling ARCHIVELOG mode for PowerExchange Express CDC for Oracle [159](#)
  - enabling minimal global supplemental logging for Express CDC [162](#)
  - specifying an archive log destination for PowerExchange Express CDC for Oracle [159](#)
- Oracle Data Guard
  - capturing data from physical standby databases [155](#)
  - logical standby database as a CDC source [119](#)
- Oracle datatypes
  - supported for PowerExchange Express CDC for Oracle [148](#)
- Oracle Express CDC
  - capturing data from physical standby databases [155](#)
- Oracle LogMiner
  - configuring for Oracle CDC [116](#)
- Oracle multitenant databases
  - configuration tasks for PowerExchange Express CDC [157](#)
  - using as PowerExchange Express CDC sources [157](#)
- Oracle user privileges
  - required for CDC [160](#)
- ORACLE\_CAPTURE\_TYPE statement
  - DBMOVER configuration file [168](#)
- ORACLE\_UNHANDLED\_NUMASCHAR statement
  - DBMOVER configuration file [170](#)
- ORACLEID statement
  - DBMOVER configuration file [123](#), [165](#)
- ORAD CAPI\_CONNECTION parameters
  - parameters and syntax [168](#)
- ORCL CAPI\_CONNECTION parameters
  - parameters and syntax [125](#)
- ORCL CAPI\_CONNECTION statement
  - CATBEGIN parameter [131](#)
  - CATEND parameter [131](#)
  - CATINT parameter [131](#)
  - Oracle catalog parameters [131](#)
- output files, PowerExchange Logger
  - CDCT file [37](#)

## P

- partitioned DB2 databases
  - adding or dropping partitions [91](#)
  - reconfiguring [91](#)
  - reconfiguring a database partition group [92](#)
- performance
  - CDC session performance details [266](#)
  - multithreaded processing [277](#)
  - Oracle CDC with LogMiner considerations [114](#)
- performance considerations
  - PowerExchange Express CDC for Oracle [153](#)
- planning considerations
  - PowerExchange Express CDC for Oracle restrictions [147](#)
- PowerCenter integration with PowerExchange [22](#)
- PowerExchange Client for PowerCenter (PWXPC) [22](#)
- PowerExchange components
  - for CDC [19](#)
  - PowerExchange Listener [19](#), [26](#)
  - PowerExchange Logger [19](#)
  - PowerExchange Navigator [20](#)
- PowerExchange Express CDC for Oracle
  - ASM considerations [154](#)
  - benefits [140](#)
  - capturing data from physical standby databases [155](#)
  - changing a source table definition [186](#)
  - configuration overview [145](#)
  - dbmover configuration file statements [164](#)
  - gathering environment information [145](#)
  - operational considerations [150](#)
  - Oracle datatypes supported and not supported [148](#)
  - Oracle multitenant databases as sources [157](#)
  - performance considerations [153](#)
  - restrictions [147](#)
  - specifying an archive log destination [159](#)
  - stopping CDC processing for a table [186](#)
  - stopping CDC processing temporarily [186](#)
  - tasks for capturing data from Oracle multitenant databases [157](#)
- PowerExchange Express CDC for Oracle configuration file
  - customization of statements and parameters [173](#)
  - DATABASE statement [174](#)
  - DICTIONARY statement [176](#)
  - DIRECTORY statement [177](#)
  - example file [184](#)
  - OPTIONS statement [177](#)
  - RAC statement [180](#)
  - READER statement [181](#)
  - STANDBY statement [183](#)
- PowerExchange Listener
  - CLOSE command [33](#)
  - DISPLAY ACTIVE command [33](#), [259](#)
  - displaying active listener tasks [33](#)
  - LISTTASK command [259](#)
  - overview [26](#)
  - starting [31](#)
  - stopping [33](#)
  - STOPTASK command [33](#)
- PowerExchange Logger for Linux, UNIX, and Windows
  - backing up CDCT, checkpoint, and log files [74](#)
  - batch mode [42](#)
  - capture registrations for logging data from remote sources [199](#)
  - CDCT file [37](#)
  - cold starting [67](#)
  - CONDENSE command [68](#)
  - configuration parameters for logging from a remote source [194](#)
  - configuration tasks for remote logging [194](#)
  - configuring [43](#)
  - connection attributes for log files remote from the source [199](#)
- PowerExchange Logger for Linux, UNIX, and Windows (*continued*)
  - continuous mode [41](#)
  - controlling [68](#)
  - DBMOVER statements for logging data from remote sources [198](#), [199](#)
  - dbmover.cfg parameters [56](#)
  - DG command [69](#)
  - DISPLAY CPU command [69](#)
  - DISPLAY EVENTS command [69](#)
  - DISPLAY MEMORY command [69](#)
  - DISPLAY RECORDS command [69](#)
  - DISPLAY STATUS command [69](#)
  - DL command [69](#)
  - example of logging data from a remote source [200](#), [202](#), [204](#)
  - FILESWITCH command [68](#)
  - group definition file [60](#)
  - lock files [39](#)
  - log file switches [40](#)
  - log files [38](#)
  - logging data from remote sources [188](#)
  - maintaining CDCT file and log files [72](#)
  - memory requirement on Linux and UNIX [42](#)
  - message log files [40](#)
  - monitoring [69](#)
  - monitoring statistics [261](#)
  - NFS access to log files [42](#)
  - operational modes [41](#)
  - output files [37](#)
  - overview [35](#)
  - pwxccl.cfg parameters [44](#)
  - regenerating the CDCT file after a failure [75](#)
  - required capture registration settings [43](#)
  - rules and guidelines for logging data from a remote source [193](#)
  - running in background mode on Linux or UNIX [42](#)
  - SHUTCOND command [68](#)
  - SHUTDOWN command [68](#)
  - start point in change stream [67](#)
  - starting [63](#)
  - stopping [68](#)
  - subtasks [37](#)
  - verbose messages [72](#)
- PowerExchange-generated extraction map columns
  - DTL\_\_BI\_columnname [210](#)
  - DTL\_\_CAPXACTION [210](#)
  - DTL\_\_CAPXCASDELIND [210](#)
  - DTL\_\_CAPXRESTART1 [210](#), [241](#)
  - DTL\_\_CAPXRESTART2 [210](#), [241](#)
  - DTL\_\_CAPXROWID [210](#)
  - DTL\_\_CAPXRRN [210](#)
  - DTL\_\_CAPXTIMESTAMP [210](#)
  - DTL\_\_CAPXUOW [210](#)
  - DTL\_\_CAPXUSER [210](#)
  - DTL\_\_CI\_columnname [210](#)
- preparing for implementation
  - environment information checklist [145](#)
- product overview
  - PowerExchange Express CDC for Oracle [139](#)
- PWX Latency in seconds attribute [229](#)
- pwxccl command
  - syntax [64](#)
- pwxccl configuration file
  - COLL\_END\_LOG parameter [44](#)
- pwxccl statement
  - parameters [65](#)
- pwxccl.cfg
  - configuring [44](#)
  - example file [56](#)
  - parameters [44](#)

- pwxccl.cfg configuration file
  - CAPT\_IMAGE parameter [44](#)
  - CAPTURE\_NODE parameter [44](#)
  - CAPTURE\_NODE\_EPWD parameter [44](#)
  - CAPTURE\_NODE\_PWD parameter [44](#)
  - CAPTURE\_NODE\_UID parameter [44](#)
  - COND\_CDCT\_RET\_P parameter [44](#)
  - CONDENSE\_SHUTDOWN\_TIMEOUT parameter [44](#)
  - CONDENSENAME parameter [44](#)
  - CONN\_OVR parameter [44](#)
  - DB\_TYPE parameter [44](#)
  - DBID parameter [44](#)
  - DISTDB parameter [44](#)
  - DISTSVR parameter [44](#)
  - ENCRYPTPWD parameter [44](#)
  - ENCRYPTOPT parameter [44](#)
  - ENCRYPTPWD parameter [44](#)
  - EXT\_CAPT\_MASK parameter [44](#)
  - FILE\_FLUSH\_VAL parameter [44](#)
  - FILE\_SWITCH\_CRIT parameter [44](#)
  - FILE\_SWITCH\_VAL parameter [44](#)
  - GROUPDEFS parameter [44](#)
  - LOGGER\_DELETES\_EXPIRED\_CDCT\_RECORDS parameter [44](#)
  - NO\_DATA\_WAIT parameter [44](#)
  - NO\_DATA\_WAIT2 parameter [44](#)
  - PROMPT parameter [44](#)
  - RESTART\_TOKEN parameter [44](#)
  - SEQUENCE\_TOKEN parameter [44](#)
  - SIGNALLING parameter [44](#)
  - STATS parameter [44](#)
  - VERBOSE parameter [44](#)
- pwxcmd
  - close [33](#)
  - closeforce [33](#)
  - listtask [33](#)
  - listtask command [259](#)
- pwxorad configuration file
  - DATABASE statement [174](#)
  - DICTIONARY statement [176](#)
  - DIRSUB statement [177](#)
  - OPTIONS statement [177](#)
  - RAC statement [180](#)
  - READER statement [181](#)
  - STANDBY statement [183](#)
- PWXPC [22](#)

## R

- RAC environments
  - Express CDC considerations [153](#)
  - PowerExchange Express CDC for Oracle
  - RAC considerations [153](#)
- RAC statement
  - PowerExchange Express CDC for Oracle configuration file [180](#)
- READER statement
  - PowerExchange Express CDC for Oracle configuration file [181](#)
- real application clusters (RACs)
  - configuring for Oracle CDC with LogMiner [119](#)
- real-time extraction mode [209](#)
- Real-Time Flush Latency in milli-seconds attribute [230](#)
- recovery
  - CDC sessions [252](#)
  - example of session recovery processing [253](#)
  - PM\_REC\_STATE table [236](#), [237](#)
  - PM\_RECOVERY table [236](#)
  - PM\_TGT\_RUN\_ID table [236](#)
  - recovery information for nonrelational targets [237](#)

- recovery (*continued*)
  - recovery state file for nonrelational targets [237](#)
  - recovery tables for relational targets [236](#)
- recovery and restart processing [235](#)
- recovery start
  - CDC sessions [248](#)
- restart
  - default restart points [239](#)
  - methods of starting CDC sessions [246](#)
  - warm starting CDC sessions [247](#)
  - warm starts
    - CDC sessions [247](#)
- restart control options
  - Application Name connection attribute [228](#)
  - RestartToken File Folder attribute [228](#)
  - RestartToken File Name attribute [228](#)
- restart of CDC sessions
  - processing by start type [238](#)
- restart token file
  - \$PMRootDir/Restart [242](#)
  - comment statements [245](#)
  - example file [245](#)
  - explicit override statements [243](#)
  - overview [215](#)
  - special override statement [244](#)
  - statement types [243](#)
  - syntax [243](#)
- restart tokens
  - creating for extraction sessions [241](#)
  - determining for cold starts [239](#)
  - displaying in DTL\_\_CAPXRESTART2 column [241](#)
  - overview [215](#)
  - recovery state file [237](#)
  - recovery state table [237](#)
- RestartToken File Folder attribute [228](#)
- RestartToken File Name attribute [228](#)
- restrictions
  - PowerExchange Express CDC for Oracle [147](#)
- row tests
  - testing data access with an extraction map [222](#)

## S

- security
  - extracting z/OS data during PowerCenter CDC sessions [221](#)
  - extracting z/OS data from remote PowerExchange Logger for LUW log files [193](#)
- sequence tokens
  - displaying in DTL\_\_CAPXRESTART1 column [241](#)
- session attributes
  - attributes to set for CDC [224](#)
- SHOW\_THREAD\_PERF parameter [258](#)
- source RDBMSs [16](#)
- source table definitions
  - changing a DB2 table definition [90](#)
  - changing a SQL Server table definition [106](#)
  - changing an Oracle table definition [138](#), [186](#)
- SQL Server CDC
  - changing a source table definition [106](#)
  - configuration tasks [99](#)
  - configuring in PowerExchange with the Logger [100](#)
  - configuring in PowerExchange without the Logger [100](#)
  - datatypes supported [96](#)
  - dbmover.cfg parameters [101](#)
  - example dbmover.cfg statements [101](#)
  - extracting data for multiple publication databases [98](#)
  - overview [94](#)

SQL Server CDC (*continued*)

- planning [95](#)
- prerequisites [95](#)
- restrictions [97](#)
- stopping [106](#)
- user authority requirements [95](#)

STANDBY statement

- PowerExchange Express CDC for Oracle configuration file [183](#)

stopping CDC processing

- for a table [186](#)
- temporarily [186](#)

STOPTASK command

- stopping CDC sessions [248](#)

supplemental logging

- enabling for Express CDC for Oracle [162](#)

supplemental logging, Oracle [115](#)

SYSASM authority [161](#)

## T

target latency [230](#)

task flow

- CDC implementation [24](#)
- extracting change data [222](#)

terminating conditions

- Idle Time attribute for CDC sessions [226](#)

troubleshooting

- DB2 for Linux, UNIX, and Windows CDC [92](#)

tuning

- overview of tuning options [219](#)

tuning CDC sessions

- APPBUFSIZE statement [269](#)

tuning CDC sessions (*continued*)

- buffer memory [274](#)
- CAPL\_CONNECTION MEMCACHE parameter [269](#)
- CAPL\_CONNECTION RSTRADV parameter [269](#)
- commit processing attributes [274](#)
- DBMOVER tuning parameters [269](#)
- methods [268](#)
- NODE and LISTENER buffer size parameters [269](#)
- PowerCenter CDC connection attributes [272](#)
- TRACE statement [269](#)

## U

UDB CAPL\_CONNECTION parameters

- parameters and syntax [85](#)

UOW Count attribute [230](#)

UOWC CAPL\_CONNECTION parameters

- parameters and syntax [132](#)

use cases [214](#)

user authority

- DB2 for Linux, UNIX, and Windows CDC requirement [78](#)
- Oracle CDC with LogMiner requirements [117](#)
- SQL Server CDC requirement [95](#)

## W

warm starts

- restart points used [240](#)