

Novell Identity Manager Driver for JDBC*

3.5.1

February 12, 2008

IMPLEMENTATION GUIDE

www.novell.com

N

Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. For more information on exporting Novell software, see the [Novell International Trade Services Web page \(http://www.novell.com/info/exports/\)](http://www.novell.com/info/exports/). Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2007 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at [Novell Legal Patents \(http://www.novell.com/company/legal/patents/\)](http://www.novell.com/company/legal/patents/) and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see [Novell Documentation \(http://www.novell.com/documentation/\)](http://www.novell.com/documentation/).

Novell Trademarks

For a list of Novell trademarks, see [Trademarks \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide	13
1 Introducing the Identity Manager Driver for JDBC	15
1.1 What's New in the Driver for JDBC	15
1.2 Driver Concepts	15
1.2.1 JDBC	16
1.2.2 Identity Manager Driver for JDBC	16
1.2.3 Third-Party JDBC Driver	16
1.2.4 Identity Vault	17
1.2.5 Directory Schema	17
1.2.6 Application Schema	17
1.2.7 Database Schema	17
1.2.8 Synchronization Schema	17
1.2.9 Logical Database Class	18
1.2.10 XDS	18
1.3 Database Concepts	18
1.3.1 Structured Query Language	18
1.3.2 Data Manipulation Language	18
1.3.3 Data Definition Language	18
1.3.4 View	19
1.3.5 Identity Columns/Sequences	19
1.3.6 Transaction	20
1.3.7 Stored Procedures or Functions	20
1.3.8 Trigger	21
1.3.9 Instead-Of-Trigger	21
1.4 Driver Features	22
1.4.1 Local and Remote Platforms	23
1.4.2 Role-Based Entitlements	23
1.4.3 Password Synchronization	23
1.4.4 Data Synchronization Models	23
1.4.5 Triggerless vs. Triggered Publication	26
2 Installing the Driver for JDBC	29
2.1 Before You Install	29
2.1.1 Driver Prerequisites	29
2.1.2 Known Issues	29
2.1.3 Limitations	30
2.1.4 Placing Jar Files	31
2.2 Installing the Driver Separately	31
2.3 Setting Up a Remote Loader	32
2.4 Installing and Configuring Database Objects	32
2.4.1 SQL Script Conventions	32
2.4.2 Installing IBM DB2 Universal Database (UDB)	34
2.4.3 Installing Informix Dynamic Server (IDS)	35
2.4.4 Installing Microsoft SQL Server	36
2.4.5 Installing MySQL	36
2.4.6 Installing Oracle	36
2.4.7 Installing PostgreSQL 7	37
2.4.8 Installing PostgreSQL 8	37

2.4.9	Installing Sybase Adaptive Server Enterprise (ASE)	38
2.5	Test Scripts	38
2.6	Troubleshooting	39
3	Uninstalling the IDM Driver for JDBC	41
3.1	Deleting Identity Manager Driver Objects	41
3.2	Running the Product Uninstaller	41
3.3	Executing Database Uninstallation Scripts	41
3.3.1	IBM DB2 Universal Database (UDB) Uninstallation	42
3.3.2	Informix Dynamic Server (IDS) Uninstallation	42
3.3.3	Microsoft SQL Server Uninstallation	42
3.3.4	MySQL Uninstallation	43
3.3.5	Oracle Uninstallation	43
3.3.6	PostgreSQL Uninstallation	43
3.3.7	Sybase Adaptive Server Enterprise (ASE) Uninstallation	44
4	Upgrading the JDBC Driver	45
4.1	Upgrading While Installing Identity Manager 3.5.1	45
4.1.1	Backward Incompatibilities	46
4.2	Upgrading after Identity Manager Is Installed	46
4.2.1	Upgrading the Driver by Using Designer	46
4.2.2	Upgrading the Driver by Using iManager	49
5	Importing an Example JDBC Configuration File	51
5.1	Using Designer to Import	51
5.2	Using iManager to Import	51
5.3	JDBC Driver Settings	52
6	Configuring the JDBC Driver	55
6.1	Smart Configuration	55
6.2	Configuration Parameters	57
6.2.1	Viewing Driver Parameters	57
6.2.2	Deprecated Parameters	57
6.2.3	Authentication Parameters	58
6.3	Driver Parameters	59
6.3.1	Uncategorized Parameters	60
6.3.2	Database Scoping Parameters	64
6.3.3	Connectivity Parameters	69
6.3.4	Compatibility Parameters	71
6.4	Subscription Parameters	81
6.4.1	Uncategorized Parameters	82
6.4.2	Primary Key Parameters	84
6.5	Publication Parameters	89
6.5.1	Uncategorized Parameters	90
6.5.2	Triggered Publication Parameters	93
6.5.3	Triggerless Publication Parameters	95
6.5.4	Polling Parameters	95
6.6	Trace Levels	98
6.7	Configuring Third-Party JDBC Drivers	99

7	Activating the JDBC Driver	101
8	Managing the JDBC Driver	103
8.1	Starting, Stopping, or Restarting the JDBC Driver	103
8.2	Migrating and Resynchronizing Data	103
8.3	Using the DirXML Command Line Utility	104
8.4	Viewing Driver Versioning Information	104
8.4.1	Viewing a Hierarchical Display of Versioning Information	104
8.4.2	Viewing the Versioning Information As a Text File	106
8.4.3	Saving Versioning Information	108
8.5	Reassociating a Driver Set Object with a Server Object	109
8.6	Changing the Driver Configuration	110
8.7	Storing Driver Passwords Securely with Named Passwords	110
8.7.1	Using Designer to Configure Named Passwords	111
8.7.2	Using iManager to Configure Named Passwords	111
8.7.3	Using Named Passwords in Driver Policies	113
8.7.4	Configuring Named Passwords by Using the DirXML Command Line Utility	114
8.8	Adding a Driver Heartbeat	117
9	Synchronizing Objects	119
9.1	What Is Synchronization?	119
9.2	When Does Synchronization Occur?	119
9.3	How Does the Metadirectory Engine Decide Which Object to Synchronize?	120
9.4	How Synchronization Works	121
9.4.1	Scenario One	121
9.4.2	Scenario Two	123
9.4.3	Scenario Three	124
10	Schema Mapping	127
10.1	High-Level View	127
10.2	Logical Database Classes	127
10.3	Indirect Synchronization	127
10.3.1	Mapping eDirectory Classes to Logical Database Classes	128
10.3.2	Parent Tables	130
10.3.3	Parent Table Columns	130
10.3.4	Child Tables	131
10.3.5	Referential Attributes	132
10.3.6	Single-Value Referential Attributes	132
10.3.7	Multivalued Referential Attributes	133
10.4	Direct Synchronization	135
10.4.1	View Column Meta-Identifiers	136
10.4.2	Primary Key Columns	138
10.4.3	Schema Mapping	138
10.5	Synchronizing Primary Key Columns	138
10.6	Synchronizing Multiple Classes	138
10.7	Mapping Multivalued Attributes to Single-Value Database Fields	139
11	Mapping XDS Events to SQL Statements	141
11.1	Mapping XDS Events for Indirect Synchronization	141

12 The Event Log Table	143
12.1 Event Log Columns	143
12.2 Event Types	145
13 Embedded SQL Statements in XDS Events	153
13.1 Common Uses of Embedded SQL	154
13.2 Embedded SQL Basics	154
13.2.1 Elements	154
13.2.2 Namespaces	154
13.2.3 Embedded SQL Example	155
13.3 Token Substitution	155
13.4 Virtual Triggers	158
13.5 Manual vs. Automatic Transactions	159
13.6 Transaction Isolation Level	160
13.7 Statement Type	161
13.8 SQL Queries	162
13.9 Data Definition Language (DDL) Statements	163
13.10 Logical Operations	164
13.11 Implementing Password Set with Embedded SQL	164
13.12 Implementing Modify Password with Embedded SQL	165
13.13 Implementing Check Object Password	165
13.14 Calling Stored Procedures and Functions	166
13.14.1 Using Embedded SQL to Call Stored Procedures or Functions	166
13.14.2 Using the jdbc:call-procedure Element	167
13.14.3 Using the jdbc:call-function Element	170
13.15 Best Practices	174
14 Supported Databases	177
14.1 Database Interoperability	177
14.2 Supported Databases	177
14.3 Database Characteristics	178
14.3.1 Database Features	178
14.3.2 Current Time Stamp Statements	179
14.3.3 Syntaxes for Calling Stored Procedures and Functions	180
14.3.4 Left Outer Join Operators	180
14.3.5 Undelimited Identifier Case Sensitivity	181
14.3.6 Supported Transaction Isolation Levels	181
14.3.7 Commit Keywords	182
14.3.8 IBM DB2 Universal Database (UDB)	182
14.3.9 Informix Dynamic Server (IDS)	183
14.3.10 Microsoft SQL Server	184
14.3.11 MySQL	184
14.3.12 Oracle	185
14.3.13 PostgreSQL	186
14.3.14 Sybase Adaptive Server Enterprise (ASE)	186
15 Third-Party JDBC Drivers	189
15.1 Third-Party JDBC Driver Interoperability	189
15.2 JDBC Driver Types	189
15.2.1 Which Type To Use?	190
15.3 Third-Party Jar File Placement	190

15.3.1	Identity Manager File Paths	190
15.3.2	Remote Loader File Paths	190
15.4	Supported Third-Party JDBC Drivers	191
15.4.1	Third-Party JDBC Driver Features	191
15.4.2	JDBC URL Syntaxes	192
15.4.3	JDBC Driver Class Names	192
15.4.4	IBM DB2 Universal Database Type 4 JDBC Driver	193
15.4.5	Informix JDBC Driver	194
15.4.6	jTDS JDBC Driver	195
15.4.7	MySQL Connector/J JDBC Driver	196
15.4.8	Oracle Thin Client JDBC Driver	197
15.4.9	Oracle OCI JDBC Driver	199
15.4.10	PostgreSQL JDBC Driver	201
15.4.11	Sybase Adaptive Server Enterprise JConnect JDBC Driver	201
15.5	Supported-But-Discouraged Third-Party JDBC Drivers	202
15.5.1	Third-Party JDBC Driver Features	202
15.5.2	JDBC URL Syntaxes	202
15.5.3	JDBC Driver Class Names	203
15.5.4	IBM DB2 Universal Database Type 3 JDBC Driver	203
15.5.5	Microsoft SQL Server 2000 Driver for JDBC	204
15.5.6	Microsoft SQL Server 2005 JDBC Driver	206
15.6	Deprecated Third-Party JDBC Drivers	207
15.6.1	BEA Weblogic jDriver for Microsoft SQL Server	207
15.7	Other Third-Party JDBC Drivers	208
15.7.1	IBM Toolbox for Java/JTOpen	209
15.7.2	Minimum Third-Party JDBC Driver Requirements	209
15.7.3	Considerations When Using Other Third-Party JDBC Drivers	209
15.8	Security Issues	210
16	The Association Utility	211
16.1	Independent Operations	211
16.2	Before You Begin	212
16.3	Using the Association Utility	213
16.4	Editing Associations	213
17	Troubleshooting the JDBC Driver	215
17.1	Recognizing Publication Events	215
17.2	Executing Test Scripts	215
17.3	Troubleshooting Driver Processes	215
17.3.1	Viewing Driver Processes	215
18	Backing Up the JDBC Driver	223
18.1	Exporting the Driver in Designer	223
18.2	Exporting the Driver in iManager	223
A	Best Practices	225
B	FAQ	227
B.1	Can't See Tables or Views	227
B.2	Synchronizing with Tables	227
B.3	Processing Rows in the Event Log Table	228

B.4	Managing Database User Accounts	228
B.5	Synchronizing Large Data Types	228
B.6	Slow Publication	228
B.7	Synchronizing Multiple Classes	229
B.8	Encrypted Transport	229
B.9	Mapping Multivalued Attributes	229
B.10	Synchronizing Garbage Strings	229
B.11	Running Multiple Driver for JDBC Instances	229
C	Supported Data Types	231
D	java.sql.DatabaseMetaData Methods	233
E	JDBC Interface Methods	235
F	Third-Party JDBC Driver Descriptor DTD	241
G	Third-Party JDBC Driver Descriptor Import DTD	243
H	Database Descriptor DTD	245
I	Database Descriptor Import DTD	247
J	Policy Example: Triggerless Future Event Processing	249
K	Setting Up an OCI Client on Linux	251
K.1	Downloading the Instant Client	251
K.2	Setting Up the OCI Client	251
K.3	Configuring the OCI Driver	252
L	Sybase Chain Modes and the Identity Manager Driver for JDBC	253
L.1	Error Codes	253
L.2	Procedures and Modes	254
L.2.1	Using Stored Procedure sp_proxmode	254
L.2.2	Chained and Unchained Modes	254
L.2.3	Managing Transactions in a Policy	255
L.2.4	Useful Links	255
M	The DirXML Command Line Utility	257
M.1	Interactive Mode	257
M.2	Command Line Mode	266
N	Properties of the JDBC Driver	271
N.1	Driver Configuration	271
N.1.1	Driver Module	272
N.1.2	Driver Object Password	273

N.1.3	Authentication	273
N.1.4	Startup Option	274
N.1.5	Driver Parameters	275
N.2	Global Configuration Values	276
N.3	Named Passwords	278
N.4	Engine Control Values	278
N.5	Log Level	280
N.6	Driver Image	281
N.7	Security Equals	281
N.8	Filter	282
N.9	Edit Filter XML	282
N.10	Misc	283
N.11	Excluded Users	283
N.12	Driver Manifest	284
N.13	Driver Cache Inspector	284
N.14	Driver Inspector	285
N.15	Server Variables	285
O	Documentation Updates	289
O.1	October 10, 2007	289
O.1.1	Upgrading the JDBC Driver	289

About This Guide

The Identity Manager Driver for Java* Database Connectivity (JDBC*) provides a generic solution for synchronizing data between an Identity Vault and relational databases.

This guide provides an overview of the driver's technology as well as configuration instructions.

This guide is version 3.5.1. The preceding version was 3.5.

Audience

This guide is for Novell® eDirectory and Identity Manager administrators who are using the Identity Manager Driver for JDBC.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with Novell Identity Manager. Please use the User Comment feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html and enter your comments there.

Documentation Updates

For the most recent version of this document, see the [Identity Manager 3.5.1 drivers \(http://www.novell.com/documentation/idm35drivers/index.html\)](http://www.novell.com/documentation/idm35drivers/index.html) documentation Web site.

Additional Documentation

For documentation on using Identity Manager and the other drivers, see the [Identity Manager 3.5.1 \(http://www.novell.com/documentation/idm35/index.html\)](http://www.novell.com/documentation/idm35/index.html) documentation Web site.

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items within a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell® trademark. An asterisk (*) denotes a third-party trademark.

What's New

See [Section 1.1, "What's New in the Driver for JDBC," on page 15.](#)

Introducing the Identity Manager Driver for JDBC

1

The Identity Manager Driver for Java DataBase Connectivity (JDBC) provides a generic solution for synchronizing data between Identity Manager and JDBC-accessible relational databases.

The principal value of this driver resides in its generic nature. Unlike most drivers that interface with a single application, this driver can interface with most relational databases and database-hosted applications.

- ♦ [Section 1.1, “What’s New in the Driver for JDBC,” on page 15](#)
- ♦ [Section 1.2, “Driver Concepts,” on page 15](#)
- ♦ [Section 1.3, “Database Concepts,” on page 18](#)
- ♦ [Section 1.4, “Driver Features,” on page 22](#)

1.1 What’s New in the Driver for JDBC

The JDBC Driver 3.5 included the new feature Stored Procedures. See [Section 13.14, “Calling Stored Procedures and Functions,” on page 166](#).

The following have been updated since the JDBC 3.5 driver guide:

- ♦ See [Chapter 14, “Supported Databases,” on page 177](#) and [Chapter 15, “Third-Party JDBC Drivers,” on page 189](#).
- ♦ See [Chapter 14, “Supported Databases,” on page 177](#) and [Chapter 15, “Third-Party JDBC Drivers,” on page 189](#).
- ♦ [Appendix F, “Third-Party JDBC Driver Descriptor DTD,” on page 241](#)
- ♦ [Appendix G, “Third-Party JDBC Driver Descriptor Import DTD,” on page 243](#)
- ♦ [Appendix H, “Database Descriptor DTD,” on page 245](#)
- ♦ [Appendix I, “Database Descriptor Import DTD,” on page 247](#)

You can now call stored procedures and functions from a policy. See [Section 13.14, “Calling Stored Procedures and Functions,” on page 166](#).

For information on what’s new in Identity Manager, see “[What's New in Identity Manager 3.5.1?](#)” in the *Identity Manager 3.5.1 Installation Guide*.

1.2 Driver Concepts

- ♦ [Section 1.2.1, “JDBC,” on page 16](#)
- ♦ [Section 1.2.2, “Identity Manager Driver for JDBC,” on page 16](#)
- ♦ [Section 1.2.3, “Third-Party JDBC Driver,” on page 16](#)
- ♦ [Section 1.2.4, “Identity Vault,” on page 17](#)
- ♦ [Section 1.2.5, “Directory Schema,” on page 17](#)

- ◆ [Section 1.2.6, “Application Schema,” on page 17](#)
- ◆ [Section 1.2.7, “Database Schema,” on page 17](#)
- ◆ [Section 1.2.8, “Synchronization Schema,” on page 17](#)
- ◆ [Section 1.2.9, “Logical Database Class,” on page 18](#)
- ◆ [Section 1.2.10, “XDS,” on page 18](#)

1.2.1 JDBC

Java DataBase Connectivity (JDBC) is a cross-platform database interface standard that Sun* Microsystems* developed.

Most enterprise database vendors provide a unique implementation of the JDBC interface. Three versions of the JDBC interface are available:

- ◆ JDBC 1 (Java 1.0)
- ◆ JDBC 2 (Java 1.2 or 1.3)
- ◆ JDBC 3 (Java 1.4 or 1.5)

The Identity Manager Driver for JDBC primarily uses the JDBC 1 interface. It uses a small subset of JDBC 2 or JDBC 3 methods when supported by third-party JDBC drivers.

1.2.2 Identity Manager Driver for JDBC

The Identity Manager Driver for JDBC uses the JDBC interface to synchronize data and identities between an Identity Vault and relational databases.

The driver consists of four jar files:

- ◆ `JDBCShim.jar`
- ◆ `JDBCUtil.jar`
- ◆ `JDBConfig.jar`
- ◆ `CommonDriverShim.jar`

In addition to these files, you need a third-party JDBC driver to communicate with each individual database.

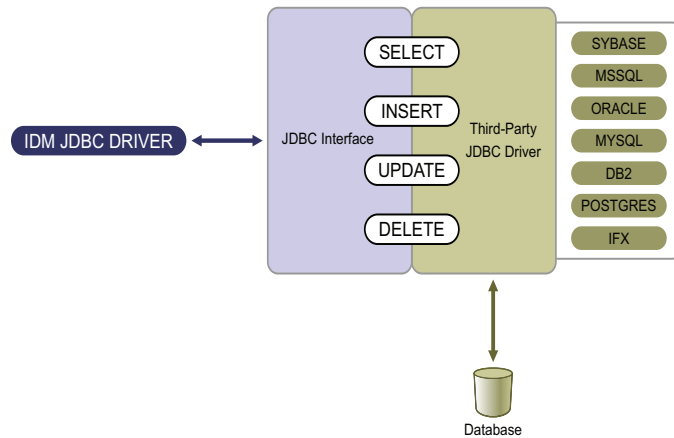
1.2.3 Third-Party JDBC Driver

A third-party JDBC driver is one of the numerous JDBC interface implementations that the Identity Manager Driver for JDBC uses to communicate with a particular database.

For example, `classes12.zip` is one of the Oracle* JDBC drivers. Different third-party JDBC drivers implement different portions of the JDBC interface specification and implement the interface in a relatively consistent manner.

The following illustration indicates the relationship between the Driver for JDBC and third-party JDBC drivers.

Figure 1-1 IDM JDBC Driver vs. Third-Party JDBC Drivers



1.2.4 Identity Vault

An Identity Vault is the data store that Identity Manager uses.

1.2.5 Directory Schema

The directory schema is the set of object classes and attributes in the directory.

For example, the eDirectory™ User class and Given Name attribute are part of the eDirectory schema.

1.2.6 Application Schema

The application schema is the set of classes and attributes in an application.

Because databases have no concept of classes or attributes, the Driver for JDBC maps eDirectory classes to tables or views, and maps eDirectory attributes to columns.

1.2.7 Database Schema

Database schema is essentially synonymous with ownership. A database schema consists of database objects (for example, tables, views, triggers, stored procedures, and functions) that a database user owns.

With the Driver for JDBC, schema is useful to scope the database (reduce the number of database objects visible to the driver at runtime).

Ownership is often expressed by using a qualified dot notation (for example, `indirect.usr`, where `indirect` is the name of the database user that owns the table `usr`). All of the database objects owned by `indirect` constitute the `indirect` database schema.

1.2.8 Synchronization Schema

The synchronization schema is the database schema visible to the driver at runtime.

1.2.9 Logical Database Class

The logical database class is the set of tables or view used to represent an eDirectory class in a database.

1.2.10 XDS

XDS format is the defined Novell® subset of possible XML formats that Identity Manager can use.

XDS is the initial format for data coming from the Identity Vault. By modifying default rules and changing the style sheets, you can configure the Driver for JDBC to work with any XML format.

1.3 Database Concepts

- ◆ [Section 1.3.1, “Structured Query Language,” on page 18](#)
- ◆ [Section 1.3.2, “Data Manipulation Language,” on page 18](#)
- ◆ [Section 1.3.3, “Data Definition Language,” on page 18](#)
- ◆ [Section 1.3.4, “View,” on page 19](#)
- ◆ [Section 1.3.5, “Identity Columns/Sequences,” on page 19](#)
- ◆ [Section 1.3.6, “Transaction,” on page 20](#)
- ◆ [Section 1.3.7, “Stored Procedures or Functions,” on page 20](#)
- ◆ [Section 1.3.8, “Trigger,” on page 21](#)
- ◆ [Section 1.3.9, “Instead-Of-Trigger,” on page 21](#)

1.3.1 Structured Query Language

Structured Query Language (SQL) is the language used to query and manipulate data in relational databases.

1.3.2 Data Manipulation Language

Data Manipulation Language (DML) statements are highly standardized SQL statements that manipulate database data.

DML statements are essentially the same, regardless of the database that you use. The Driver for JDBC is DML-based. It maps Identity Manager events expressed as XDS XML to standardized DML statements.

The following example shows several DML statements:

```
SELECT * FROM usr;  
INSERT INTO usr(lname) VALUES ('Doe');  
UPDATE usr SET fname = 'John' WHERE idu = 1;
```

1.3.3 Data Definition Language

Data Definition Language (DDL) statements manipulate database objects such as tables, indexes, and user accounts.

DDL statements are proprietary and differ substantially between databases. Even though the Driver for JDBC is DML-based, you can embed DDL statements in XDS events. For additional information, refer to [Chapter 13, “Embedded SQL Statements in XDS Events,” on page 153](#),

The following examples show several DDL statements:

```
CREATE TABLE usr
(
    idu    INTEGER,
    fname  VARCHAR2(64),
    lname  VARCHAR2(64)
);

CREATE USER idm IDENTIFIED BY novell;
```

NOTE: Examples used throughout this guide are for the Oracle database.

1.3.4 View

A view is a logical table.

When queried by using a `SELECT` statement, the view is constituted by executing the SQL query supplied when the view was defined. Views are a useful abstraction mechanism for representing multiple tables of arbitrary structure as a single table or logical database class.

```
CREATE VIEW view_usr
(
    pk_idu,
    fname,
    lname
)
AS
SELECT idu, fname, lname from usr;
```

1.3.5 Identity Columns/Sequences

Identity columns and sequences are used to generate unique primary key values. Identity Manager can associate with these values, among other things.

An identity column is a self-incrementing column used to uniquely identify a row in a table. Identity column values are automatically filled in when a row is inserted into a table.

A sequence object is a counter that can be used to uniquely identify a row in a table. Unlike an identity column, a sequence object is not bound to a single table. However, if it is used by a single table, a sequence object can be used to achieve an equivalent result.

The following is an example of a sequence object:

```
CREATE SEQUENCE seq_idu
    START WITH 1
    INCREMENT BY 1
    NOMINVALUE
    NOMAXVALUE
    ORDER;
```

1.3.6 Transaction

A transaction is an atomic database operation that consists of one or more statements.

When a transaction is complete, all statements in the transaction are committed. When a transaction is interrupted or one of the statements in the transaction has an error, the transaction is said to roll back. When a transaction is rolled back, the database is left in the same state it was before the transaction began.

Transactions are either manual (user-defined) or automatic. Manual transactions can consist of one or more statements and must be explicitly committed. Automatic transactions consist of a single statement and are implicitly committed after each statement is executed.

Manual (User-Defined) Transactions

Manual transactions usually contain more than one statement. DDL statements typically cannot be grouped with DML statements in a manual transaction.

The following example illustrates a manual transaction:

```
SET AUTOCOMMIT OFF
INSERT INTO usr(lname) VALUES('Doe');
UPDATE usr SET fname = 'John' WHERE idu = 1;
COMMIT; -- explicit commit
```

Automatic Transactions

Automatic transactions consist of only one statement. They are often referred to as auto-committed statements because changes are implicitly committed after each statement. An auto-committed statements is autonomous of any other statement.

The following example illustrates an automatic transaction:

```
SET AUTOCOMMIT ON
INSERT INTO emp(lname) VALUES('Doe');
-- implicit commit
```

1.3.7 Stored Procedures or Functions

A stored procedure or function is programmatic logic stored in a database. Stored procedures or functions can be invoked from almost any context.

The Subscriber channel can use stored procedures or functions to retrieve primary key values from rows inserted into tables, to create associations. Stored procedures or functions can also be invoked from within embedded SQL statements or triggers.

The distinction between stored procedures and functions varies by database. Typically, both can return output, but they differ in how they do it. Stored procedures usually return values through parameters. Functions usually return values through a scalar return value or result set.

The following example illustrates a stored procedure definition that returns the next value of a sequence object:

```
CREATE SEQUENCE seq_idu
  START WITH 1
  INCREMENT BY 1
  NOMINVALUE
```

```

        NOMAXVALUE
        ORDER;
CREATE
PROCEDURE sp_idu(io_idu IN OUT INTEGER)
IS
BEGIN
    IF (io_idu IS NULL) THEN
        SELECT seq_idu.nextval INTO io_idu FROM DUAL;
    END IF;
END sp_idu;

```

1.3.8 Trigger

A database trigger is programmatic logic associated with a table, which executes under certain conditions. A trigger is said to fire when its execution criteria are met.

Triggers are often useful for creating side effects in a database. In the context of the Driver for JDBC, triggers are useful to capture event publications. The following is an example of a database trigger on the `usr` table.

```

CREATE TABLE usr
(
    idu    INTEGER,
    fname VARCHAR2(64),
    lname VARCHAR2(64)
);
-- t = trigger; i = insert
CREATE TRIGGER t_usr_i
    AFTER INSERT ON usr
    FOR EACH ROW
BEGIN
    UPDATE usr SET fname = 'John';
END;

```

When a statement is executed against a table with triggers, a trigger fires if the statement satisfies the conditions specified in the trigger. For example, using the above table, suppose the following insert statement is executed:

```
INSERT INTO usr(lname) VALUES('Doe')
```

Trigger `t_emp_i` fires after the insert statement is executed, and the following update statement is also executed:

```
UPDATE usr SET fname = 'John'
```

A trigger can typically be fired before or after the statement that triggered it. Statements that are executed as part of a database trigger are typically included in the same transaction as the triggering statement. In the above example, both the `INSERT` and `UPDATE` statements are committed or rolled back together.

1.3.9 Instead-Of-Trigger

An instead-of-trigger is programmatic logic associated with a view, which executes under certain conditions.

Instead-of-triggers are useful for making views writable or subscribeable. They are often used to define what it means to INSERT, UPDATE, and DELETE from a view. The following is an example of an instead-of-trigger on the `usr` table.

```
CREATE TABLE usr
(
    idu    INTEGER,
    fname  VARCHAR2(64),
    lname  VARCHAR2(64)
);

CREATE VIEW view_usr
(
    pk_idu,
    fname,
    lname
)
AS
SELECT idu, fname, lname from usr;
-- t = trigger; i = insert
CREATE TRIGGER t_view_usr_i
    INSTEAD OF INSERT ON usr
BEGIN
    INSERT INTO usr(idu, fname, lname)
        VALUES (:NEW.pk_idu, :NEW.fname, :NEW.lname);
END;
```

When a statement is executed against a view with instead-of-triggers, an instead-of-trigger executes if the statement satisfies the conditions specified in the trigger. Unlike triggers, instead-of-triggers always execute before the triggering statement. Also, unlike regular triggers, instead-of-triggers are executed instead of, not in addition to, the triggering statement.

For example, using the above view, suppose the following insert statement is executed instead of the original insert statement:

```
INSERT INTO view_usr(pk_idu, fname, lname)
    VALUES(1, 'John', 'Doe')
```

Rather than executing the original statement, instead-of-trigger `t_view_usr_i` fires and executes the following statement:

```
INSERT INTO usr(idu, fname, lname)
    VALUES (:NEW.pk_idu, :NEW.fname, :NEW.lname);
```

In this example, the statements happen to be equivalent.

1.4 Driver Features

- ◆ [Section 1.4.1, “Local and Remote Platforms,” on page 23](#)
- ◆ [Section 1.4.2, “Role-Based Entitlements,” on page 23](#)
- ◆ [Section 1.4.3, “Password Synchronization,” on page 23](#)
- ◆ [Section 1.4.4, “Data Synchronization Models,” on page 23](#)
- ◆ [Section 1.4.5, “Triggerless vs. Triggered Publication,” on page 26](#)

1.4.1 Local and Remote Platforms

The driver runs on all Identity Manager-enabled platforms, including Windows* NT*/2000, NetWare[®], Solaris*, Linux*, and AIX*. The JDBC driver runs in any Identity Manager 3.5.1 installation or Remote Loader installation. See “[System Requirements for Identity Manager](#)” in the *Identity Manager 3.5.1 Installation Guide*.

For information on supported databases, see “[Database Interoperability](#)” on page 177.

For information on supported third-party JDBC drivers, see “[Third-Party JDBC Driver Interoperability](#)” on page 189.

1.4.2 Role-Based Entitlements

The JDBC driver does not support entitlements.

1.4.3 Password Synchronization

The JDBC driver supports password set and check on the Subscriber channel. The driver does not support bi-directional password synchronization.

1.4.4 Data Synchronization Models

The JDBC driver supports two data synchronization models: direct and indirect. Both terms are best understood with respect to the final destination of the data being synchronized.

Model	Association	Description
Direct	Usually associated with views	Views provide the abstraction mechanism that best facilitates integration with existing customer tables.
Indirect	Usually associated with tables	Customer tables probably don't match the structure required by the driver. Therefore, it's usually necessary to create intermediate staging tables that do match the structure that the driver requires. Although the structures might match, it is highly unlikely.

The following sections describe how direct and indirect synchronization work on both the Subscriber and Publisher channels.


Indirect Synchronization

Indirect synchronization uses intermediate staging tables to synchronize data between the Identity Vault and a database.

The following diagrams illustrate how indirect synchronization works on the Subscriber and Publisher channels. In the following scenarios, you can have one or more customer tables and intermediate staging tables.

Subscriber Channel

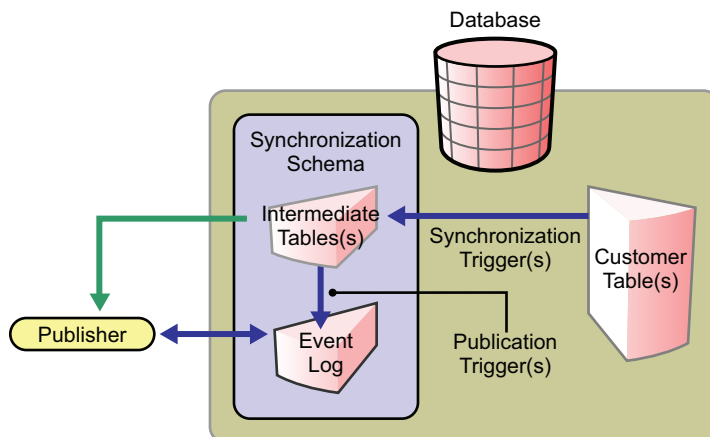
Figure 1-2 *Indirect Synchronization on the Subscriber Channel*



The Subscriber channel updates the intermediate staging tables in the synchronization schema. The synchronization triggers then update customer tables elsewhere in the database.

Publisher Channel

Figure 1-3 *Indirect Synchronization on the Publisher Channel*



When customer tables are updated, synchronization triggers update the intermediate staging tables. Publication triggers then insert one or more rows into the event log table. The Publisher channel then reads the inserted rows and updates the Identity Vault.

Depending on the contents of the rows read from the event log table, the Publisher channel might need to retrieve additional information from the intermediate tables before updating the Identity Vault. After updating the Identity Vault, the Publisher channel then deletes or marks the rows as processed.


Direct Synchronization

Direct synchronization typically uses views to synchronize data between Identity Manager and a database. You can use tables if they conform to the structure that the Driver for JDBC requires.

The following diagrams illustrate how direct synchronization works on the Subscriber and Publisher channels. In the following scenarios, you can have one or more customer views or tables.

Subscriber Channel

Figure 1-4 Direct Synchronization on the Subscriber Channel

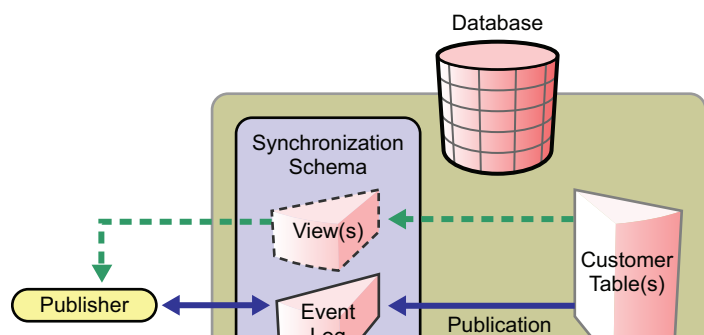


The Subscriber channel updates existing customer tables through a view in the synchronization schema.

NOTE: Direct synchronization without a view is possible only if customer tables match the structure that the Driver for JDBC requires. For additional information, see [Section 10.3, “Indirect Synchronization,”](#) on page 127.

Publisher Channel

Figure 1-5 Direct Synchronization on the Publisher Channel



When a customer table is updated, publication triggers insert rows into the event log table. The Publisher channel then reads the inserted rows and updates the Identity Vault.

Depending on the contents of the rows read from the event log table, the Publisher channel might need to retrieve additional information from the view before updating the Identity Vault. After updating the Identity Vault, the Publisher channel then deletes or marks the rows as processed.

1.4.5 Triggerless vs. Triggered Publication

Triggers are no longer required to log publication events. In situations where triggers cannot be used to capture granular events, the Publisher channel can derive database changes by inspecting database data.

Triggerless publication is particularly useful when support contracts forbid the use of triggers on database application tables or for rapid prototyping.

Triggerless publication is less efficient than triggered publication. With triggered publication, what changed is already known. With triggerless publication, change calculation must occur before events can be processed.


Triggerless publication, unlike triggered publication, does not preserve event order. It only guarantees that by the end of a polling cycle, objects in the database and the Identity Vault are in sync.

Triggerless publication, unlike triggered publication, does not provide historical data such as old values. It provides information on the current state of an object, not the previous state.

Triggerless publication does have the advantage of being much simpler because it reduces database-side dependencies. Writing database triggers can be complicated and requires extensive knowledge of database-specific SQL syntaxes.


The following figure illustrates direct triggerless publication:

Figure 1-6 Direct Triggerless Synchronization



The following figure illustrates indirect triggerless publication:

Figure 1-7 *Indirect Triggerless Synchronization*



If you move the driver without moving the state files, the driver must build-up new state files by resynchronizing. For information on this situation, see [“State Directory” on page 63](#).

Installing the Driver for JDBC

2

This section assumes that you have already installed the Metadirectory engine (and, most likely, other drivers) on the server and need to install only the JDBC driver. See “[Installing Identity Manager](#)” in the *Identity Manager 3.5.1 Installation Guide*.

Typically, an Identity Manager installation installs all drivers, including the JDBC driver, at the same time that the Metadirectory engine is installed. If the JDBC driver wasn’t installed at that time, you can install the driver separately. The schema won’t be extend during this driver install because the Identity Manager installation already extended it when the Metadirectory engine was installed.

IMPORTANT: Novell recommends installing or uninstalling driver configurations and database scripts as a unit. To prevent unintentional mismatching, database scripts and driver configurations contain headers with a version number, the target database name, and the database version.

- ♦ [Section 2.1, “Before You Install,” on page 29](#)
- ♦ [Section 2.2, “Installing the Driver Separately,” on page 31](#)
- ♦ [Section 2.3, “Setting Up a Remote Loader,” on page 32](#)
- ♦ [Section 2.4, “Installing and Configuring Database Objects,” on page 32](#)
- ♦ [Section 2.5, “Test Scripts,” on page 38](#)
- ♦ [Section 2.6, “Troubleshooting,” on page 39](#)

For information on uninstalling the driver, see [Chapter 3, “Uninstalling the IDM Driver for JDBC,” on page 41](#)

2.1 Before You Install

- ♦ [Section 2.1.1, “Driver Prerequisites,” on page 29](#)
- ♦ [Section 2.1.2, “Known Issues,” on page 29](#)
- ♦ [Section 2.1.3, “Limitations,” on page 30](#)
- ♦ [Section 2.1.4, “Placing Jar Files,” on page 31](#)

2.1.1 Driver Prerequisites

The Identity Manager Driver for JDBC 3.5.1 requires the following:

- Novell® Identity Manager 3.5.1 installed on the server or an Identity Manager 3.5.1 Remote Loader
- Java Virtual Machine (JVM*) 1.4 or later
- A supported third-party JDBC driver

2.1.2 Known Issues

- ♦ Identity Vault Time and Timestamp syntaxes are inadequate for expressing the range and granularity of their database counterparts.

This is a publication problem because database time-related types typically have a wider range and greater degree of granularity (typically nanoseconds). The converse is not true. For more information, see [“Time Syntax” on page 61](#).

- ◆ The Driver for JDBC is unable to parse proprietary database time stamp formats.

Some databases, such as Sybase* and DB2*, have proprietary time stamp formats that the `java.sql.Timestamp` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html>) class can't parse.

When synchronizing time stamp columns from these databases, the Driver for JDBC, by default, assumes time stamp values placed in the event log table are in ODBC canonical format (that is, `yyyy-mm-dd hh:mm:ss.ffffffffff`).

The recommended method for enabling the Driver for JDBC to handle proprietary database time stamp formats is to implement a custom `DBTimestampTranslator` class. This interface is documented in the Javadoc Tool that ships with the Driver for JDBC. Using this approach avoids the problem of reformatting time stamps in the database before they are inserted into the event log table or reformatted them in style sheets. The Driver for JDBC ships with default implementations for the native DB2 time stamp format and the Sybase style 109 time stamp format.

- ◆ Statements executed against the database server might block indefinitely.

Typically, blocking is caused by a database resource being exclusively locked. Because the locking mechanisms and locking SQL vary by database, the general solution to this problem is to implement a custom `DBLockStatementGenerator` class. For additional information, see [“Lock Statement Generator Class” on page 76](#). The driver for JDBC ships with a default implementation for Oracle.

Many factors can cause blocking. To mitigate the likelihood of blocking, we recommend that you do not set the [Transaction Isolation Level](#) parameter to a level greater than `read committed`.

The JDBC interface defines a method `java.sql.Statement.setQueryTimeout(int):void` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html>) that allows a statement to time out after a specified number of seconds. Unfortunately, implementations of this method between third-party JDBC drivers range from not being implemented to having bugs. For this reason, this method was deemed unsuitable as a general-purpose solution.

2.1.3 Limitations

- ◆ The Driver for JDBC does not support the use of delimited (quoted) database identifiers (for example, “names with spaces”).
- ◆ JDBC 2 data types are not supported, with the exception of Large Object data types (LOBs) such as CLOB and BLOB.
- ◆ JDBC 3 data types are not supported.
- ◆ PostgreSQL does not support `<check-object-password>` events. Authentication is controlled by manually inserting entries into the `pg_hba.conf` file.

2.1.4 Placing Jar Files

The following tables identify the paths to place JDBC driver jar files on an Identity Manager or Remote Loader server assuming default installation paths.

Identity Manager File Paths

The following table identifies where to place JDBC driver jar files on an Identity Management server, by platform.

Table 2-1 Locations for jar Files: Identity Manager Server

Platform	Directory Path
NetWare®	sys:\system\lib
Solaris, Linux, or AIX	/usr/lib/dirxml/classes (pre-eDirectory 8.8) /opt/novell/eDirectory/lib/dirxml/classes (eDirectory 8.8)
Windows NT/2000	novell\NDS\lib

Remote Loader File Paths

The following table identifies where to place JDBC driver jar files on a Remote Loader server, by platform.

Table 2-2 Locations for jar Files: Remote Loader

Platform	Directory Path
Solaris, Linux, or AIX	/usr/lib/dirxml/classes (pre-eDirectory 8.8) /opt/novell/eDirectory/lib/dirxml/classes (eDirectory 8.8)
Windows NT/2000	novell\RemoteLoader\lib

2.2 Installing the Driver Separately

This section assumes that you have already installed the Metadirectory engine (and, most likely, other drivers) on the server and need to install only the JDBC driver. For instructions on installing Identity Manager, see “[Installing Identity Manager](#)” in the *Identity Manager 3.5.1 Installation Guide*.

Typically, an Identity Manager installation installs all drivers, including the JDBC driver, at the same time that the Metadirectory engine is installed. If the JDBC driver wasn’t installed at that time, you can install the driver separately. The schema won’t be extend during this driver install because the Identity Manager installation already extended it when the Metadirectory engine was installed.

The installation process installs the driver on the server where Identity Manager is installed or on a server where you will use the Remote Loader to run the driver. It also installs the driver information and policies on the iManager server.

Table 2-3 *Where to Find Install Instructions*

Platform	Where to Find Install Instructions
Linux	See “Installing Identity Manager through the GUI Interface on UNIX/Linux Platforms” in the <i>Identity Manager 3.5.1 Installation Guide</i>
NetWare	See “Installing Identity Manager on NetWare” in the <i>Identity Manager 3.5.1 Installation Guide</i>
Windows	See “Installing Identity Manager on Windows” in the <i>Identity Manager 3.5.1 Installation Guide</i> .

2.3 Setting Up a Remote Loader

Using a Remote Loader is optional. It isn’t required unless you want the JDBC driver to run in a connected system.

- 1 If a Remote Loader isn’t already installed, install one.

See “Deciding Whether to Use the Remote Loader” in the *Novell Identity Manager 3.5.1 Administration Guide*.

- 2 Copy the appropriate third-party JDBC driver jar files onto the Remote Loader server.

2a For information on third-party JDBC driver filenames and where to get them, refer to “Supported Third-Party JDBC Drivers” on page 191.

2b For information on file installation paths, refer to “Placing Jar Files” on page 31.

- 3 Configure the remote driver.

In the *Remote Driver Configuration* parameters, set the *Driver* parameter to `com.novell.nds.dirxml.driver.jdbc.JDBCDSDriverShim`.

- 4 Configure other remote loader parameters. See “Deciding Whether to Use the Remote Loader” in the *Novell Identity Manager 3.5.1 Administration Guide*.

2.4 Installing and Configuring Database Objects

Install and configure database objects (for example, tables, triggers, and indexes) for synchronization with the sample driver configuration. If you don’t configure database objects, the sample configuration file won’t work.

2.4.1 SQL Script Conventions

The following table lists default locations for SQL scripts:

Table 2-4 *Default Locations for SQL Scripts*

Platform	Default Location
Windows	<code>c:\novell\NDS\jdbc\sql\database-abbreviation</code>
UNIX or Linux	<code>/usr/lib/dirxml/rules/jdbc/database-abbreviation</code>

For example, when installed on a SuSE Linux Enterprise Server with eDirectory, the DB2 scripts are found in `/usr/lib/ dirxml/rules/jdbc/db2/*`.

All SQL scripts use the same conventions, regardless of the database.

The maximum size of a DB2 identifier is 18 characters. This least common denominator length defines the upper bound of database identifier length across all SQL scripts. Because of this restricted length, abbreviations are used. The following table summarizes identifier abbreviations and their meaning:

Table 2-5 *Identifier Abbreviations and Meanings*

Abbreviation	Interpretation
proc_ ¹	stored procedure/function
idx_	index
trg_	trigger
_i	on insert trigger
_u	on update trigger
_d	on delete trigger
chk_	check constraint
pk_	view primary key constraint
fk_	view foreign key constraint
mv_	view multi-valued column
sv_	view single-valued column (implicit default)

¹ The more common abbreviation is `sp_`. This prefix is reserved for system-stored procedures on Microsoft* SQL Server. Also, this prefix forces lookup of a procedure first in the master database before evaluating any qualifiers (for example, database or owner). To maximize procedure lookup efficiency, this prefix has been deliberately avoided.

The following table indicates identifier naming conventions for indexes, triggers, stored procedures, functions, and constraints:

Table 2-6 *Identifier Naming Conventions*

Database Object	Naming Convention	Examples
stored procedure/ function	<code>proc_<i>procedure-or-function-name</i></code>	<code>proc_idu</code>
index	<code>idx_<i>unqualified-table-name</i>_<i>sequence-number</i></code>	<code>idx_indirectlog _1</code>
trigger	<code>tgr_<i>unqualified-table-name</i>_<i>triggering-statement-type</i>_<i>sequence-number</i></code>	<code>tgr_usr_i_1</code>
primary key constraint	<code>pk_<i>unqualified-table-name</i>_<i>column-name</i></code>	<code>pk_usr_idu</code>

Database Object	Naming Convention	Examples
foreign key constraint	<code>fk_unqualified-table-name_column-name</code>	<code>fk_usr_idu</code>
check constraint	<code>chk_unqualified-table-name_column-name</code>	<code>chk_usr_idu</code>

Other conventions:

- ◆ All database identifiers are lowercase.
This is the most commonly used case convention between databases.
- ◆ String field lengths are 64 characters.
Fields of this length can hold most eDirectory™ attribute values. You might want to refine field lengths to enhance storage efficiency.
- ◆ For performance reasons, primary key columns use native, scalar numeric types whenever possible (such as `BIGINT` as opposed to `NUMERIC`).
- ◆ The `record_id` column in event log tables has the maximum numeric precision permitted by each database to avoid overflow.
- ◆ Identity columns and sequence objects do not cache values. Some databases throw away cached values when a rollback occurs. This action can cause large gaps in identity column or sequence values.

2.4.2 Installing IBM DB2 Universal Database (UDB)

IMPORTANT: For IBM* DB2, you must manually create operating system user accounts before running the provided SQL scripts.

Because the process to create user accounts differs between operating systems, Step 1 below is OS-specific. These instructions are for a Windows NT operating environment. If you rerun the SQL scripts, repeat only Steps 2 through 5.

The directory context for DB2 is `install-dir\jdbc\sql\db2_udb\install`

- 1 Create user accounts for users `idm`, `indirect` and `direct`.

Use `novell` as the password in *User Manager for Domains*.

Remember to deselect *User Must Change Password at Next Login* for this account.

You might want to also select *Password Never Expires*.

NOTE: The remaining instructions are OS-independent.

- 2 Adjust the file path to `idm_db2.jar` in the `1_install.sql` installation script. The file path to `idm_db2.jar` should reflect the location of this file on your client machine.
- 3 Execute the `1_install.sql` script from the Command Line Processor (CLP.)

For example:

```
db2 -f 1_install.sql
```

IMPORTANT: The scripts won't execute in the Command Center interface beyond version 7. The scripts use the '\ line continuation character. Later versions of the Command Center don't recognize this character.

- 4 For versions 8 or later, execute the `2_install_8.sql` script.

For example:

```
db2 -f 2_install_8.sql
```

2.4.3 Installing Informix Dynamic Server (IDS)

IMPORTANT: For Informix* Dynamic Server, you must manually create an operating system user account before running the provided SQL scripts.

Because the process of creating user accounts differs between operating systems, Step 1 below is OS-specific. These instructions are for a Windows NT operating environment. If you rerun the SQL scripts, you should repeat only Steps 2 through 4.

The directory context for Informix SQL scripts is `install-dir\jdbc\sql\informix_ids\install`.

- 1 In Windows NT, create a user account for user `idm`.

Use `novell` as the password in *User Manager for Domains*.

Remember to deselect *User Must Change Password at Next Login* for this account.

You might want to also select *Password Never Expires*.

NOTE: The remaining instructions are OS-independent.

- 2 Start a client such as SQL Editor or DBAccess.
- 3 Log in to your server as the `informix` user or another user with DBA (database administrator) privileges.

By default, the password for the `informix` user is `informix`.

NOTE: If you execute scripts as a user other than `informix`, change all references to `informix` in the scripts prior to execution.

- 4 Open and execute `1_install_9.sql` from either the `ansi` (transactional, ANSI-compliant), `log` (transactional, non-ANSI-compliant), or `no_log` (non-transactional, non-ANSI-compliant) subdirectory, depending upon which type of database you want to create.
- 5 For version 10 or later, open and execute `2_install_10.sql` from either the `ansi` (transactional, ANSI-compliant), `log` (transactional, non-ANSI-compliant), or `no_log` (non-transactional, non-ANSI-compliant) subdirectory, depending upon which type of database you want to create.

2.4.4 Installing Microsoft SQL Server

The directory context for Microsoft SQL Server scripts is *install-dir\jdbc\sql\mssql\install*.

- 1 Start a client such as Query Analyzer (7, 2000) or Microsoft SQL Server Management Studio (2005).
- 2 Log in to your database server as the *sa* user.
By default, the *sa* user has no password.
- 3 Execute the installation script.
For version 7, execute *1_install_7.sql*.
For version 2000 (8), execute *1_install_2k.sql*.
For version 2005 (9), execute *1_install_2005.sql*.

NOTE: The execute hotkey in Query Analyzer is F5.

2.4.5 Installing MySQL

The directory context for MySQL* SQL scripts is *install-dir\jdbc\sql\mysql\install*.

- 1 From a MySQL client, such as *mysql*, log in as *root* user or another user with administrative privileges.

For example, from the command line, execute

```
mysql -u root -p
```


By default, the *root* user has no password.
- 2 Execute the installation script *1_install_innodb.sql* or *1_install_myisam.sql*, depending upon which table type you wish to use. For MySQL 3 or 4, use the scripts in subdirectory *3or4*. For version 5 or later, use the scripts in subdirectory *5*.

For example:

```
mysql> \. c:\1_install_innodb.sql
```

TIP: Don't use a semicolon to terminate this statement.

2.4.6 Installing Oracle

The directory context for Oracle SQL scripts is *install-dir\jdbc\sql\oracle\install*.

- 1 From an Oracle client, such as SQL Plus, log in as the *SYSTEM* user.

By default, the password for *SYSTEM* is *MANAGER*.

NOTE: If you execute scripts as a user other than *SYSTEM* with password *MANAGER*, change all references to *SYSTEM* in the scripts prior to execution.

- 2 Execute the installation script *1_install.sql*.

For example:
SQL> @c:\1_install.sql

2.4.7 Installing PostgreSQL 7

The directory context for PostgreSQL scripts is *install-dir\jdbc\sql\postgres\install*. The directory context for executing Postgres commands is *postgres-install-dir/pgsql/bin*.

- 1 Create the database *idm*.

For example, from the UNIX* command line, execute the command *createdb*:

```
./createdb idm
```

- 2 Install the *plpgsql* procedural language to database *idm*.

For example, from the UNIX command line, execute the command *createlang*:

```
./createlang plpgsql idm
```

- 3 From a Postgres client such as *psql*, log on as user *postgres* to the *idm* database.

For example, from the UNIX command line, execute the command *psql*:

```
./psql -d idm postgres
```

By default, the Postgres user has no password.

- 4 From inside *psql*, execute the script *1_install_7.sql*.

For example:
idm=# \i 1_install_7.sql

- 5 Update the *pg_hba.conf* file.

For example, add entries for the *idm* database user. Adjust the IP-ADDRESS and IP-MASK as necessary:

```
# TYPE DATABASE USER IP-ADDRESS IP-MASK
METHOD# allow driver user idm to connect to database idm
host idm idm 255.255.255.255 255.255.255.0
password
```

- 6 Restart the Postgres server to effect changes made to the *pg_hba.conf* file.

2.4.8 Installing PostgreSQL 8

The directory context for PostgreSQL scripts is *install-dir\jdbc\sql\postgres\install*. The directory context for executing Postgres commands is *postgres-install-dir/pgsql/bin*.

- 1 From a Postgres client such as *psql*, log in as user *postgres* to the *idm* database.

For example, from the UNIX command line, execute the command *psql*:

```
./psql -d idm postgres
```

By default, the Postgres user has no password.

- 2 From inside *psql*, execute the script *1_install_8.sql*.

For example:
idm=# \i 1_install_8.sql

- 3 Update the *pg_hba.conf* file.

As of version 8, this can be done through pgAdminIII. After you start, go to *Tools > Server Configuration > pg_hba.conf*. In the pgAdminIII *pg_hba.conf* editor, the IP-ADDRESS and IP-MASK columns in the file are combined into a single field: IP-Address. Place both the IP-ADDRESS and IP-MASK values in that field, separated by a single whitespace character.

For example, add entries for the *idm* database user. Adjust the IP-ADDRESS and IP-MASK as necessary:

```
# TYPE DATABASE USER IP-ADDRESS IP-MASK
METHOD# allow driver user idm to connect to database idm
host idm idm 255.255.255.255 255.255.255.0
password
```

- Restart the Postgres server to effect changes made to the *pg_hba.conf* file.

If you are using pgAdminIII, from the *pg_hba.conf* editor press the disk icon (save file) in the toolbar. When prompted, press *Yes*.

2.4.9 Installing Sybase Adaptive Server Enterprise (ASE)

IMPORTANT: Ensure that you have JDBC metadata support installed on the database server. This is usually an issue for versions earlier than 12.5 only.

The directory context for Sybase SQL scripts is *install-dir\jdbc\sql\sybase_ase\install*.

- From a Sybase client, such as *isql*, log in as the *sa* user and execute the *1_install.sql* installation script.

For example, from the command line, execute:

```
isql -U sa -P -i 1_install.sql
```

By default, the *sa* account has no password.

2.5 Test Scripts

Test scripts for each database are located in the following directories:

Table 2-7 Location of Database Scripts

Database	Test SQL Scripts Location
IBM DB2 Universal Database	<i>install-dir\jdbc\sql\db2_udb\test</i>
Informix Dynamic Server	<i>install-dir\jdbc\sql\informix_ids\log\test</i> <i>install-dir\jdbc\sql\informix_ids\no_log\test</i> Informix ANSI test scripts are located in the <i>log\test</i> subdirectory.
Microsoft SQL Server	<i>install-dir\jdbc\sql\mssql\test</i>
MySQL	<i>install-dir\jdbc\sql\mysql\test</i>
Oracle	<i>install-dir\jdbc\sql\oracle\test</i>

Database	Test SQL Scripts Location
PostgreSQL	<code>install-dir\jdbc\sql\postgres\test</code>
Sybase Adaptive Server Enterprise	<code>install-dir\jdbc\sql\sybase_ase\test</code>

We recommend that you try the test scripts before starting the sample driver.

2.6 Troubleshooting

- ◆ See [Section 17.1, “Recognizing Publication Events,”](#) on page 215.
- ◆ See [Section 17.2, “Executing Test Scripts,”](#) on page 215.

Uninstalling the IDM Driver for JDBC

3

- ◆ [Section 3.1, “Deleting Identity Manager Driver Objects,” on page 41](#)
- ◆ [Section 3.2, “Running the Product Uninstaller,” on page 41](#)
- ◆ [Section 3.3, “Executing Database Uninstallation Scripts,” on page 41](#)

IMPORTANT: Novell® recommends that you install and uninstall preconfigured drivers and database scripts as a unit. To prevent unintentional mismatching, database scripts and preconfigured drivers contain headers with a version number, the target database name, and the database version.

3.1 Deleting Identity Manager Driver Objects

When deleting Novell Identity Vault objects, you must delete all child objects before you can delete a parent object. For example, you must delete all rules and style sheets on the Publisher channel before you can delete the Publisher object. Similarly, you must delete both the Publisher and Subscriber objects before you can delete the Driver object.

To remove a driver object from an Identity Vault:

- 1 In Novell iManager, click *Identity Manager > Identity Manager Overview*.
- 2 Select a driver set.
- 3 From the Identity Manager Overview page, click *Delete Driver*.
- 4 Select the driver that you want to delete, then click *OK*.

3.2 Running the Product Uninstaller

Uninstallation procedures vary by platform.

To uninstall the Identity Manager Driver for JDBC on Windows, use *Add or Remove Programs* in the *Control Panel*.

3.3 Executing Database Uninstallation Scripts

This section provides helps you execute database uninstallation SQL scripts.

- ◆ [Section 3.3.1, “IBM DB2 Universal Database \(UDB\) Uninstallation,” on page 42](#)
- ◆ [Section 3.3.2, “Informix Dynamic Server \(IDS\) Uninstallation,” on page 42](#)
- ◆ [Section 3.3.3, “Microsoft SQL Server Uninstallation,” on page 42](#)
- ◆ [Section 3.3.4, “MySQL Uninstallation,” on page 43](#)
- ◆ [Section 3.3.5, “Oracle Uninstallation,” on page 43](#)
- ◆ [Section 3.3.6, “PostgreSQL Uninstallation,” on page 43](#)
- ◆ [Section 3.3.7, “Sybase Adaptive Server Enterprise \(ASE\) Uninstallation,” on page 44](#)

3.3.1 IBM DB2 Universal Database (UDB) Uninstallation

The directory context for DB2 is `install-dir\jdbc\sql\db2_udb1\install`.

- 1 Drop the `idm`, `indirect` and `direct` operating system user accounts.
- 2 If you haven't already done so, change the name of the administrator account name and password in the installation scripts.
- 3 Using the Command Line Processor (CLP), execute script `uninstall.sql`.

For example:

```
db2 -f uninstall.sql
```

IMPORTANT: This script won't execute in the Command Center interface beyond version 7. It uses the `\` line continuation character. Later versions of the Command Center don't recognize this character.

- 4 Delete the `idm_db2.jar` file.

3.3.2 Informix Dynamic Server (IDS) Uninstallation

The directory context for Informix SQL scripts is `install-dir\jdbc\sql\informix_ids\install`.

- 1 Drop the `idm` operating system user account.
- 2 Start a client such as SQL Editor.
- 3 Log on to your server as user `informix` or another user with DBA (database administrator) privileges.

By default, the password for `informix` is `informix`.

If you execute scripts as a user other than `informix`, change all references to `informix` in the install scripts prior to execution.

- 4 If you aren't using the `informix` account with the default password, change the name of the DBA account name and password in the installation scripts if you haven't already done so.
- 5 Open and execute `uninstall.sql` from the `ansi` (transactional, ANSI-compliant), `log` (transactional, non-ANSI-compliant), or `no_log` (non-transactional, non-ANSI-compliant) subdirectory, depending upon which type of database you installed.

3.3.3 Microsoft SQL Server Uninstallation

The directory context for Microsoft SQL Server scripts is `install-dir\jdbc\sql\mssql\install`.

- 1 Start a client such as Query Analyzer.
- 2 Log on to your database server as user `sa`.
By default, the `sa` user has no password.
- 3 Open and execute the first installation script `uninstall.sql`.

The execute hotkey in Query Analyzer is F5.

3.3.4 MySQL Uninstallation

The directory context for MySQL SQL scripts is *install-dir\jdbc\sql\mysql\install*.

- 1 From a MySQL client, such as `mysql`, log on as user `root` or another user with administrative privileges.

For example, from the command line execute
`mysql -u root -p`

By default, the `root` user has no password.

- 2 Execute the uninstallation script `uninstall.sql`.

For example:

```
mysql> \. c:\uninstall.sql
```

Don't use a semicolon to terminate this statement.

3.3.5 Oracle Uninstallation

The directory context for Oracle SQL scripts is *install-dir\jdbc\sql\oracle\install*.

- 1 From an Oracle client, such as SQL Plus, log on as user `SYSTEM`.

By default, the password for `SYSTEM` is `MANAGER`.

If you execute scripts as a user other than `SYSTEM` with password `MANAGER`, change all references to `SYSTEM` in the scripts prior to execution.

- 2 Execute the uninstallation script `uninstall.sql`.

For example:

```
SQL> @c:\uninstall.sql
```

3.3.6 PostgreSQL Uninstallation

The directory context for PostgreSQL scripts is *install-dir\jdbc\sql\postgres\install*. The directory context for executing Postgres commands is *postgres-install-dir/pgsql/bin*.

- 1 From a Postgres client such as `psql`, log on as user `postgres` to the `idm` database.

For example, from the UNIX command line, execute
`./psql -d idm postgres`

By default, the Postgres user has no password.

- 2 From inside `psql`, execute the script `uninstall.sql`.

For example:

```
idm=# \i uninstall.sql
```

- 3 Drop the database `idm`.

For example, from the UNIX command line, execute
`./dropdb idm`

- 4 Remove or comment out entries for the `idm` user from the `pg_hba.conf` file.

For example:

```
#host      idm          idm      255.255.255.255    255.255.255.0
```

5 Restart the Postgres server to effect changes made to the `pg_hba.conf` file.

3.3.7 Sybase Adaptive Server Enterprise (ASE) Uninstallation

The directory context for Sybase SQL scripts is `install-dir\jdbc\sql\sybase_ase\install`.

- 1** From a Sybase client, such as `isql`, log on as user `sa`.
- 2** Execute the installation script `uninstall.sql`.

For example, from the command line, execute
`isql -U sa -P -i uninstall.sql`

By default, the `sa` account has no password.

Upgrading the JDBC Driver

4

If you are upgrading from Identity Manager 3.5.0 to Identity Manager 3.5.1, skip this section. The following information does not apply.

Identity Manager 3.5 introduced a new architecture for how policies reference one another. To take advantage of this architecture, the driver configuration file provided for earlier versions of the JDBC driver must be upgraded. For more information on the new architecture, see “[Upgrading Identity Manager Policies](#)” in *Understanding Policies for Identity Manager 3.5.1*.

You can upgrade the Identity Manager Driver for JDBC at the same time that you install Identity Manager, or you can upgrade after Identity Manager is installed.

- ◆ [Section 4.1, “Upgrading While Installing Identity Manager 3.5.1,” on page 45](#)
- ◆ [Section 4.2, “Upgrading after Identity Manager Is Installed,” on page 46](#)

4.1 Upgrading While Installing Identity Manager 3.5.1

During an Identity Manager 3.5.1 installation, you can install the Driver for JDBC (along with other Identity Manager drivers) at the same time that Identity Manager 3.5.1 is installed. You can upgrade from DirXML 1.1a, Identity Manager 2, or Identity Manager 3 to Identity Manager 3.5.1.

The Identity Manager Driver for JDBC 3.5.1 won’t run on Identity Manager earlier than Identity Manager 3.0.

The Identity Manager Driver for JDBC 2.1 won’t run on Identity Manager earlier than Identity Manager 3.0. To use the Driver for JDBC 2.1, you must upgrade to Identity Manager 3.x.

The Identity Manager Driver for JDBC 2.0 runs on Identity Manager 2.

To upgrade from the Identity Manager Driver for JDBC 1.5 or later to 3.5.1, install the Driver for JDBC 3.5.1. This task replaces only binaries.

You can upgrade from DirXML 1.1a or Identity Manager 2 to Identity Manager 3.

The following tables summarize upgrade paths:

Table 4-1 *Upgrading to the Identity Manager Driver for JDBC 3.5.1*

If You Are Running This Version	Upgrade to This Version	Before Upgrading to This Version
Driver for JDBC earlier than 1.5	Driver for JDBC 1.51	Driver for JDBC 2.0
Driver for JDBC 1.5 or later	None	Driver for JDBC 2.0

Table 4-2 Upgrading to the Identity Manager Driver for JDBC 2.1 or later

If You Are Running This Version	Upgrade to This Version	Before Upgrading to This Version
Driver for JDBC earlier than 1.5	Driver for JDBC 1.51	Driver for JDBC 2.1
Driver for JDBC 1.5 or later	None	Driver for JDBC 3.5.1

For Identity Manager Driver for JDBC versions earlier than 1.5, you must first upgrade to version 1.5. Refer to the *DirXML Driver 1.5 for JDBC Implementation Guide* (<http://www.novell.com/documentation/ig/dirxmldrivers/index.html>). Be sure to use the 2.1 Association Utility. It supersedes all previous versions.

4.1.1 Backward Incompatibilities

- ◆ The driver now requires a minimum of two database connections for bidirectional synchronization. For additional information, refer to “[Use Minimal Number of Connections?](#)” on page 69.
- ◆ The driver now returns schema qualifiers (when available) for logical database class names (parent table or view names). This change doesn’t affect existing configurations unless class names are remapped in Schema Mapping policies. If class names are remapped, all references to class names in existing policy need to be schema-qualified.
- ◆ Slightly alter existing configurations that use views. Set the parameter Enable Meta-Identifier Support to Boolean False. See “[Enable Meta-Identifier Support?](#)” on page 77.
- ◆ Slightly alter existing configurations that reference the `com.novell.nds.dirxml.driver.jdbc.util.MappingPolicy` class. Methods in this class no longer edit the source document. Instead, they return node sets that must be copied into the destination document. The example `.xml` driver configuration file includes examples of how to do this.
- ◆ Slightly alter existing configurations deployed against DB2/AS400 or other legacy databases that do not implement or support column position. Add and set the *Sort Column Names By* parameter. To sort column names by string collation order, see “[Sort Column Names By](#)” on page 80. The default behavior has been changed to sort column names by hexadecimal value.

4.2 Upgrading after Identity Manager Is Installed

If Identity Manager 3.5.1 is already installed, you can use either Designer for Identity Manager or iManager to upgrade the JDBC driver.

4.2.1 Upgrading the Driver by Using Designer

- 1 Make sure that you have updated your driver with all the patches for the version you are currently running.

We recommend this step for all drivers, to help minimize upgrade issues.

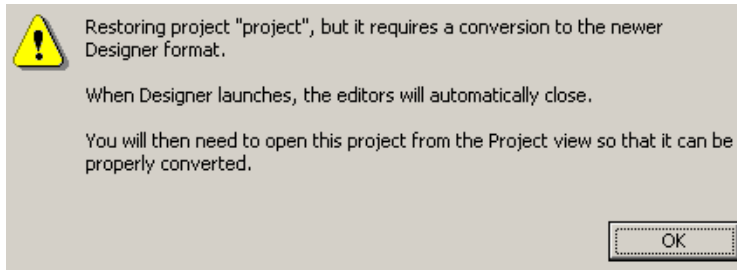
- 2 Back up the driver.

See [Chapter 18, “Backing Up the JDBC Driver,”](#) on page 223.

- 3 Install Designer version 2.1 or later, then launch Designer.


If you had a project open in Designer when you upgraded Designer, proceed to **Step 4**. If you didn't have a project open in Designer when you upgraded Designer, skip to **Step 5**.

- 4 If you had a project open when upgrading Designer, read the warning message, then click *OK*.

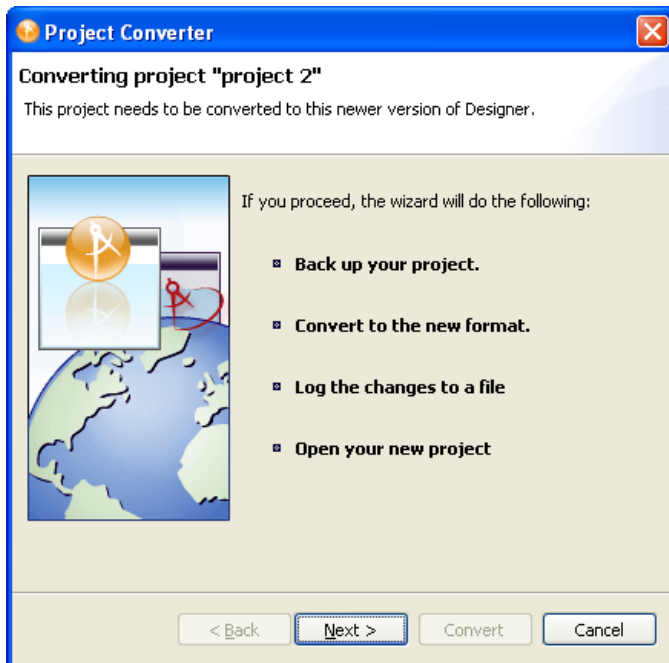


Designer closes the project to preform the upgrade.

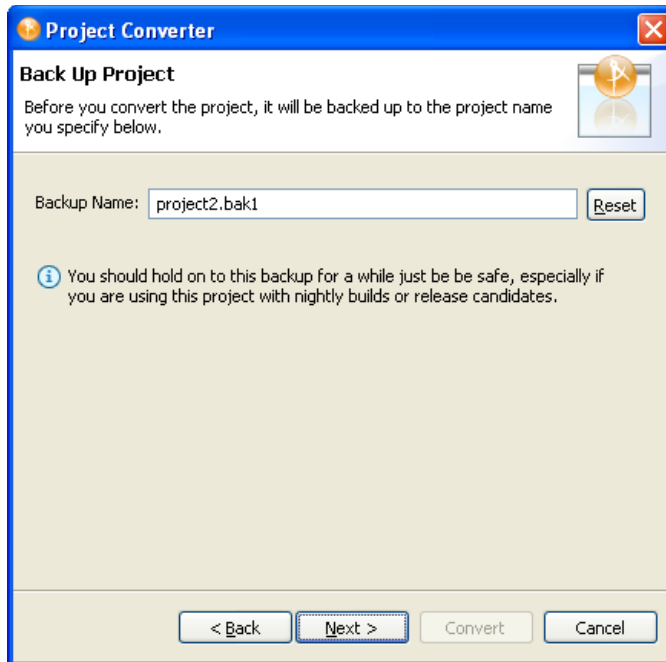
- 5 To open and convert the project, double-click *System Model* in the Project view.



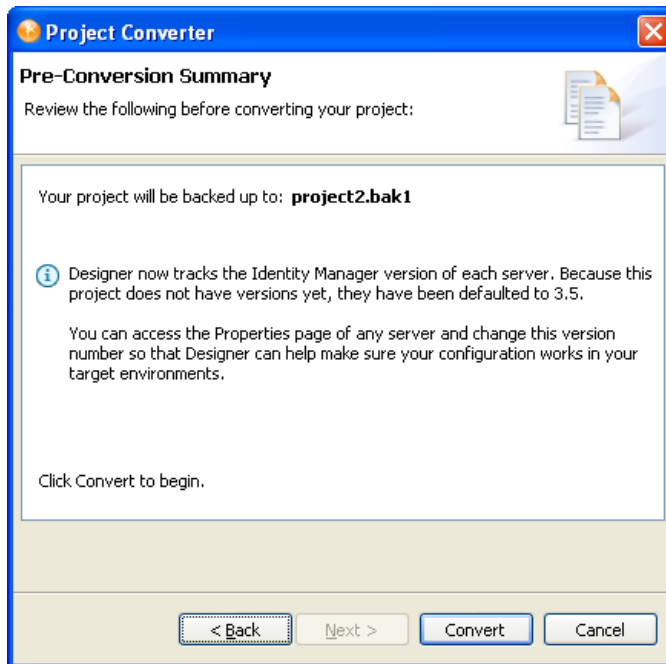
- 6 Read the tasks listed in the Project Converter message, then click *Next*.



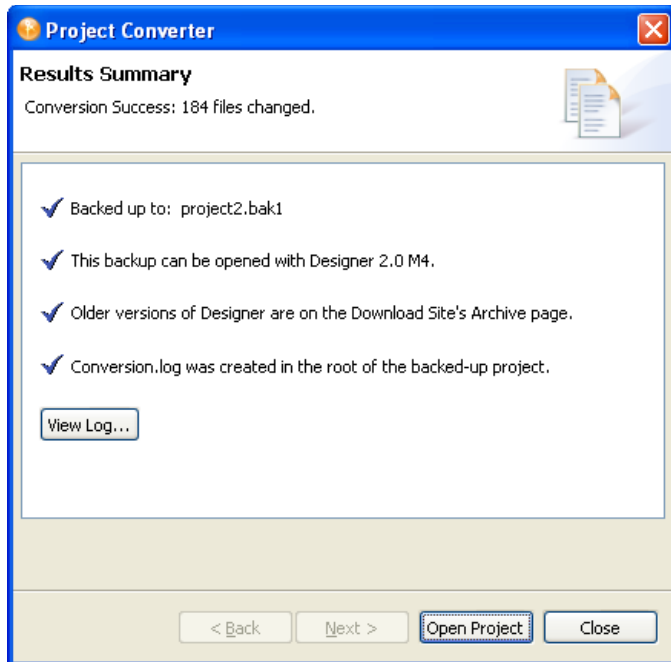
- 7 Specify the name of the backup project name, then click *Next*.



- 8 Read the project conversion summary, then click *Convert*.



- 9 Read the project conversion result summary, then click *Open Project*.



To view the log file that is generated, click *View Log*.

4.2.2 Upgrading the Driver by Using iManager


- 1 Make sure that you have updated your driver with all the patches for the version you are currently running.

To help minimize upgrade issues, we recommend that you complete this step on all drivers.

- 2 Back up the driver.

See [Chapter 18, “Backing Up the JDBC Driver,” on page 223](#).

- 3 Verify that Identity Manager 3.5.1 has been installed and that you have the current plug-ins installed.
- 4 Launch iManager.
- 5 Click *Identity Manager > Identity Manager Overview*.
- 6 Click *Search* to find the Driver Set object, then click the driver that you want to upgrade.
- 7 Read the message that is displayed, then click *OK*.



IMPORTANT: The example configuration file for the updated driver changed for the Identity Manager 3.0 release. If your current configuration meets your requirements, you don't need to import this example configuration. If you *do* import the new sample configuration, you will see an

additional driver for Delimited Text with a new name, a new Identity Vault container specified in the placement rule, and new rule names.

Importing an Example JDBC Configuration File

5

The Identity Manager Driver for JDBC includes an example configuration file that you can use as a starting point for creating the Driver object. When you import this file, Designer for Identity Manager or iManager creates and configures the objects and policies needed to make the driver work properly.

- ♦ [Section 5.1, “Using Designer to Import,” on page 51](#)
- ♦ [Section 5.2, “Using iManager to Import,” on page 51](#)
- ♦ [Section 5.3, “JDBC Driver Settings,” on page 52](#)

5.1 Using Designer to Import

The example `.xml` configuration file creates and configures the Identity Manager objects needed for the sample driver to work properly. The configuration file also includes example policies that you can customize.

The following procedure explains one of several ways to import the example configuration file:

- 1 Open a project in Designer.
- 2 In the Modeler, right-click the Driver Set object, then select *New > Driver*.
- 3 From the drop-down list, select *IDM Driver for JDBC 3.5.1*, then click *Run*.
- 4 Configure the driver by filling in the fields.
Specify information specific to your environment. See [Table 5-1 on page 52](#).
- 5 After specifying parameters, click *Finish* to import the driver.
- 6 Test the driver.
- 7 Deploy the driver into the Identity Vault.
See “[Deploying a Driver to an Identity Vault](#)” in the *Designer 2.1 for Identity Manager 3.5.1* guide.

5.2 Using iManager to Import

Identity Manager provides an example configuration file. You installed this file when you installed the Identity Manager Web components on an iManager server. Think of the example configuration file as a template that you import and customize or configure for your environment.

- 1 In iManager, select *Identity Manager Utilities > New Driver*.
- 2 Select a driver set, then click *Next*.

Where do you want to place the new drivers?

- In an existing driver set

snati_drset.novell  

- In a new driver set

If you place this driver in a new driver set, you must specify a driver set name, context, and associated server.

3 Select *JDBC-IDM3_5-v1.xml*, then click *Next*.

4 Configure the driver by filling in the configuration parameters.

For information on the settings, see [Table 5-1 on page 52](#).

5 Define security equivalences by using a User object that has the rights that the driver needs to have on the server

The Admin user object is most often used for this task. However, you might want to create a DriversUser (for example) and assign security equivalence to that user. Whatever rights that the driver needs to have on the server, the DriversUser object must have the same security rights.

6 Identify all objects that represent administrative roles and exclude them from replication.

Exclude the security-equivalence object (for example, DriversUser) that you specified in Step 2. If you delete the security-equivalence object, you have removed the rights from the driver. Therefore, the driver can't make changes to Identity Manager.

7 Click *Finish*.

Configuration File Conventions

- ◆ Database usernames are the surname of a user concatenated with the corresponding numeric primary key value. For example, John Doe's username could be Doe1.
- ◆ Initial passwords are the surname of a user. For example, John Doe's password would be Doe.

Sybase passwords must be at least 6 characters long. When shorter than 6 characters, last names are padded with the character "p." For example, John Doe's password would be Doepppp. The padding character can be adjusted in the Subscriber Command Transformation policies.

5.3 JDBC Driver Settings

Table 5-1 JDBC Driver Settings

Setting	Description
<i>Driver name</i>	The name that you want to display in the driver set.
<i>Target database</i>	That database that the driver writes to.
<i>Driver is local/remote</i>	Specifies whether the driver runs locally or remotely on a Remote Loader.

Setting	Description
<i>Synchronization model</i>	Specifies whether the driver uses views to synchronize directly to existing tables of arbitrary structure or synchronize to intermediate staging tables of a particular structure.
<i>Third-party JDBC implementation</i>	The third-party implementation that the driver uses.
<i>Data flow</i>	Specifies whether the authoritative source of data is the database, Identity Manager, or bi-directional (both the database and Identity Manager).
<i>Database host IP address</i>	The IP address of the database host.
<i>Database port</i>	Specifies the port that the driver uses to communicate with the database. If you don't provide a port number, the Driver Configuration Wizard provides a default port number for the database that you selected at install time.
<i>User container DN</i>	The Distinguished Name (complete context) of the container where the database users are published.
<i>Group container DN</i>	The Distinguished Name (complete context) of the container where the database groups are published.
<i>Publication mode</i>	Specifies whether publication is triggered (default) or triggerless

Configuring the JDBC Driver

6

- ◆ [Section 6.1, “Smart Configuration,” on page 55](#)
- ◆ [Section 6.2, “Configuration Parameters,” on page 57](#)
- ◆ [Section 6.3, “Driver Parameters,” on page 59](#)
- ◆ [Section 6.4, “Subscription Parameters,” on page 81](#)
- ◆ [Section 6.5, “Publication Parameters,” on page 89](#)
- ◆ [Section 6.6, “Trace Levels,” on page 98](#)
- ◆ [Section 6.7, “Configuring Third-Party JDBC Drivers,” on page 99](#)

6.1 Smart Configuration

The JDBC driver can recognize the supported set of third-party JDBC drivers and databases. Also, the driver can dynamically and automatically configure the majority of driver compatibility parameters. These features alleviate the need for the end user to understand and explicitly set such parameters.

These features are implemented via the following four types of XML descriptor files, which describe a third-party JDBC driver or database to the JDBC driver.

- ◆ Third-party JDBC driver
- ◆ Third-party JDBC driver import
- ◆ Database
- ◆ Database import

In addition to pre-defined descriptor files, you can create custom descriptor files for a database or third-party JDBC driver.

Specifying Custom Descriptor Files

You can force the driver to use a custom descriptor file for a database or third-party JDBC driver. To specify a custom database descriptor file, see [“Database Descriptor Filename” on page 72](#). To specify a custom third-party driver descriptor file, see [“JDBC Driver Descriptor Filename” on page 71](#). This is useful when multiple descriptor files exist for the same database or third-party JDBC driver. For the custom descriptor file to take effect, set the driver parameter as the jdbc-driver-descriptor.

Reserved Filenames for Descriptor Files

Descriptor filenames that ship with the driver begin with the underscore character (_). Such filenames are reserved to ensure that descriptor files that ship with the driver do not conflict with custom descriptor files. Obviously, custom descriptor filenames must not begin with the underscore character.

Import Descriptor Files

Import descriptor files allow multiple, nonimport descriptor files to share content. This functionality reduces the size of nonimport descriptor files, minimizes the need for repetition of content, and increases maintainability. Import files cannot be imported across major types. That is, JDBC driver descriptors cannot import database imports, and database descriptors cannot import JDBC driver imports.

Furthermore, custom nonimport descriptors cannot import reserved descriptor imports. For example, if a custom third-party JDBC driver descriptor file named `custom.xml` tries to import a reserved third-party JDBC driver descriptor named `_reserved.xml`, an error is issued. These limitations accomplish the following:

- ◆ Ensure that no dependencies exist between reserved and custom import files
- ◆ Allow extension of existing reserved descriptor files in later versions of the driver

Descriptor File Locations

Descriptor files must be located in a `jar` file whose name begins with the prefix “jdbc” (case-insensitive) and resides in the runtime classpath.

The following table identifies where to place descriptors within a descriptor `jar` file:

Table 6-1 *Where to Place Descriptors*

Descriptor Type	Directory Path
Third-party JDBC driver	<code>com/novell/nds/dirxml/driver/jdbc/db/descriptor/driver</code>
Third-party JDBC driver import	<code>com/novell/nds/dirxml/driver/jdbc/db/descriptor/driver/import</code>
Database	<code>com/novell/nds/dirxml/driver/jdbc/db/descriptor/db</code>
Database import	<code>com/novell/nds/dirxml/driver/jdbc/db/descriptor/db/import</code>

Reserved descriptor files are located in the `JDBConfig.jar` file. To ensure that these reserved files are not overwritten when the Driver for JDBC is updated, place custom descriptors in a different `jar` file.

Precedence

Parameters explicitly specified through a management console, such as iManager, always have precedence over parameters specified through descriptor files. Descriptor file parameters only take effect when a parameter is not set through the management console.

Parameters and other information specified in a nonimportable descriptor file always have precedence over that specified in descriptor import files. If a parameter or other information is duplicated within a descriptor file, the first instance of the parameter or information takes precedence over subsequent instances.

Between import files, precedence is determined by import order. Import files declared earlier in the import list take precedence over those that follow.

Custom Descriptor Best Practices

- ❑ Do not begin custom descriptor files name with the underscore (`_`) character.
- ❑ Place custom descriptor files in a jar file other than `JDBCConfig.jar`, and begin the filename with the prefix “jdbc” (case-insensitive).
- ❑ Do not use custom descriptors to import reserved import files (filenames that begin with the underscore character).

Descriptor File DTDs

The following appendixes contain DTDs for all descriptor file types. These DTDs can help you construct custom descriptor files.

Table 6-2 *Where to Find Descriptor DTDs*

Descriptor Type	Appendix
Third-party JDBC driver	Appendix F, “Third-Party JDBC Driver Descriptor DTD,” on page 241
Third-party JDBC driver import	Appendix G, “Third-Party JDBC Driver Descriptor Import DTD,” on page 243
Database	Appendix H, “Database Descriptor DTD,” on page 245
Database import	Appendix I, “Database Descriptor Import DTD,” on page 247

6.2 Configuration Parameters

- ◆ [Section 6.2.1, “Viewing Driver Parameters,” on page 57](#)
- ◆ [Section 6.2.2, “Deprecated Parameters,” on page 57](#)
- ◆ [Section 6.2.3, “Authentication Parameters,” on page 58](#)

6.2.1 Viewing Driver Parameters

- 1 In iManager, click *Identity Manager > Identity Manager Overview*.
- 2 Locate the driver set containing the driver, then click the driver’s icon.
- 3 From the Identity Manager Driver Overview, click the driver object.
iManager displays the driver’s configuration parameters.

6.2.2 Deprecated Parameters

The following parameters have been deprecated since version 1.6:

Table 6-3 *Deprecated Parameters*

Tag Name	Justification
connection-tester-class	The driver now dynamically creates a connection tester class at runtime, based upon information in XML descriptor files. This parameter is still operable, to ensure backwards compatibility. Its continued use, however, is discouraged.
connection-test-stmt	The driver now dynamically creates a connection tester class at runtime, based upon information in XML descriptor files. This parameter is still operable, to ensure backwards compatibility. Its continued use, however, is discouraged.
reconnect-interval	The reconnect interval is now fixed at 30 seconds on both channels.

6.2.3 Authentication Parameters

After you import the driver, provide authentication information for the target database.

Authentication ID

An Authentication ID is the name of the driver's database user/login account. The installation SQL script for each database provides information on the database privileges required for this account to authenticate to a supported database. The scripts are located in the *install-dir\tools\sql\abbreviated-database-name\install* directory.

This value can be referenced in the *Connection Properties* parameter value via the token `{ $username }`. See [“Connection Properties” on page 70](#).

The default value for the sample configuration is `idm`.

Authentication Context

The authentication context is the JDBC URL of the target database.

URL format and content are proprietary. They differ between third-party JDBC drivers. However, they have some similarities in content. Each URL, whatever the format, usually includes an IP address or DNS name, port number, and a database identifier. For the exact syntax and the content requirements of your driver, consult your third-party driver documentation.

For a list of JDBC URL syntaxes for supported third-party drivers, see [“JDBC URL Syntaxes” on page 192](#).

IMPORTANT: Changing anything in this value other than URL properties forces a resynchronization of all objects when triggerless publication is used.

Application Password

An application password is the password for the driver's database user/login account. The default value for the sample driver configuration is `novell`.

This value can be referenced in the *Connection Properties* parameter value via the token `{ $password }`. See [“Connection Properties” on page 70](#).

6.3 Driver Parameters

The following table summarizes all driver-level parameters and their properties:

Table 6-4 *Driver Parameters and Properties*

Display Name	Tag Name	Sample Value	Default Value	Required
Third-Party JDBC Driver Class Name	jdbc-class	oracle.jdbc.driver.OracleDriver	(none)	yes
Time Syntax	time-syntax	1 (integer)	1 (integer)	no
Synchronization Filter	sync-filter	schema (include by schema membership)	(none)	no
Schema Name	sync-schema	indirect	(none)	yes ¹
Include Filter Expression	include-table-filter	IDM_.*	(none)	no
Exclude Filter Expression	exclude-table-filter	BIN\\$.{22}==\ \$0	(none)	no
Table/View Names	sync-tables	usr	(none)	yes ¹
Connection Initialization Statements	connection-init	USE idm	(none)	no
Use Minimal Number of Connections?	use-single-connection	0 (no)	(dynamic ³)	no
Connection Properties	connection-properties	USER=\${username}; PASSWORD=\${password}	(dynamic ³)	no
State directory	state-dir	. (current directory)	. (current directory)	no
JDBC Driver Descriptor Filename	jdbc-driver-descriptor	ora_client_thin.xml	(none)	no
Database Descriptor Filename	database-descriptor	ora_10g.xml	(none)	no
Use Manual Transactions?	use-manual-transactions	1 (yes)	(dynamic ²)	no
Transaction Isolation Level	transaction-isolation-level	read committed	(dynamic ³)	no
Reuse Statements?	reuse-statements	1 (reuse)	(dynamic ³)	no
Number of Returned Result Sets	handle-stmt-results	one	(dynamic ³)	no
Enable Statement-Level Locking?	enable-locking	1 (yes)	0 (no)	no
Lock Statement Generator Class	lock-generator-class	com.novell.nds.dirxml.driver.jdbc.db.lock.OraLockGenerator	(dynamic ³)	no

Display Name	Tag Name	Sample Value	Default Value	Required
Enable Referential Attribute Support?	enable-refs	1 (yes)	1 (yes)	no
Enable Meta-Identifier Support?	enable-meta-identifiers	1 (yes)	1 (yes)	no
Force Username Case	force-username-case	upper (to uppercase)	(none)	no
Left Outer Join Operator	left-outer-join-operator	(+)	(dynamic ³)	no
Retrieve Minimal Metadata	minimal-metadata	0 (no)	(dynamic ³)	no
Function Return Method	function-return-method	result set	(dynamic ³)	no
Supports Schemas in Metadata Retrieval?	supports-schemas-in-metadata-retrieval	1 (yes)	(dynamic ³)	no
Sort Column Names By	column-position-comparator	com.novell.nds.dirxml.driver.jdbc.util.config.comp.StringByteComparator (hexadecimal value)	(dynamic ³)	no

¹ One of these mutually-exclusive parameters must be present if the Synchronization Filter parameter is not present. See [“Synchronization Filter” on page 64](#).

² This default is derived dynamically at runtime from descriptor files and database metadata.

³ This default is derived dynamically from descriptor files at runtime.

Driver parameters fall into the following subcategories:

- ◆ [Section 6.3.1, “Uncategorized Parameters,” on page 60](#)
- ◆ [Section 6.3.2, “Database Scoping Parameters,” on page 64](#)
- ◆ [Section 6.3.3, “Connectivity Parameters,” on page 69](#)
- ◆ [Section 6.3.4, “Compatibility Parameters,” on page 71](#)

6.3.1 Uncategorized Parameters

- ◆ [“Third-Party JDBC Driver Class Name” on page 60](#)
- ◆ [“Time Syntax” on page 61](#)
- ◆ [“State Directory” on page 63](#)

Third-Party JDBC Driver Class Name

This parameter is the fully-qualified Java class name of your third-party JDBC driver.

The following table lists the properties of this parameter:

Table 6-5 *Third-Party JDBC Driver Class Name: Properties*

Property	Value
Tag Name	jdbc-class
Required?	yes
Case-Sensitive?	yes
Sample Value	oracle.jdbc.driver.OracleDriver
Default Value	(none)

For a list of supported third-party JDBC driver classnames, see [“JDBC Driver Class Names” on page 192](#).

Time Syntax

The Time Syntax parameter specifies the format of time-related data types that the driver returns. The format can be any of the following options:

- ◆ [“Return Database Time, Date, and Timestamp Values as 32-Bit Integers” on page 61](#)
- ◆ [“Return Database Time, Date, and Timestamp Values as Canonical Strings” on page 61](#)
- ◆ [“Return database Time, Date, and Timestamp Values in their Java String Representation as Returned by the Method toString\(\):java.lang.String” on page 62](#)

Return Database Time, Date, and Timestamp Values as 32-Bit Integers

This is the default.

eDirectory Time and Timestamp syntaxes are composed of unsigned, 32-bit integers that express the number of whole seconds that have elapsed since 12:00 a.m., January 1st, 1970 UTC. The maximum range of this data type is approximately 136 years. When interpreted as unsigned integers (as originally intended), these syntaxes are capable of expressing dates and times to the second in the range of 1970 to 2106. When interpreted as a signed integer, these syntaxes are capable of expressing dates and times to the second in the range of 1901 to 2038.

This option has two problems:

- ◆ Identity Vault Time and Timestamp syntaxes cannot express as large a date range as database Date or Timestamp syntaxes.
- ◆ Identity Vault Time and Timestamp syntaxes are granular to the second. Database Timestamp syntaxes are often granular to the nanosecond.

The second and third options overcome these two limitations.

NOTE: Map the database Time, Date, and Timestamp values to eDirectory attributes of type Time or Timestamp.

Return Database Time, Date, and Timestamp Values as Canonical Strings

The following table shows abstract database data types and their corresponding canonical string representations:

Table 6-6 Database Types and Canonical String Representations

JDBC Data Type	Canonical String Format ¹
java.sql.Time	HHMMSS
java.sql.Date	CCYYMMDD
ava.sql.Timestamp	CCYYMMDDHHMMSSNNNNNNNNNN

¹ C = century, Y = year, M = month D = day, H = hour, M= minute, S = second, N = nano

These fixed-length formats collate in chronological order on any platform in any locale. Even though the precision of nanoseconds varies by database, the length of Timestamps does not.

NOTE: Map the database Time, Date, and Timestamp values to attributes of type Numeric String.

Return database Time, Date, and Timestamp Values in their Java String Representation as Returned by the Method toString():java.lang.String

The following table shows abstract database data types and their corresponding Java String representations:

Table 6-7 Database Types and Java String Formats

JDBC Data Type	Java String Format ¹
java.sql.Time	hh:mm:ss
java.sql.Date	yyyy-mm-dd
java.sql.Timestamp	yyyy-mm-dd hh:mm:ss.fffffffff

¹ y= year, m= month, d= day, h= hour, m= minute, s= second, f= nano

These fixed-length formats collate in chronological order on any platform in any locale. The precision of nanoseconds, and hence the length of Timestamps, varies by database.

NOTE: Map the database Time, Date, and Timestamp values to attributes of type Case Ignore/Case Exact String.

The following table lists the properties of the Time Syntax parameter:

Table 6-8 Time Syntax: Properties

Property	Value
Tag Name	time-syntax
Required?	no
Default Value	1 (integer)

Property	Value
Legal Values	1 (integer) 2 (canonical string) 3 (java string)
Schema-Dependent?	True

State Directory

The State Directory parameter specifies where a driver instance should store state data. State data is currently used for triggerless publication. See [“Triggerless Publication Parameters” on page 95](#). State data might be used to store additional state information in the future.

Each driver instance has two state files. State filenames follow the formats `jdbc_driver-instance-guid.db` and `jdbc_driver-instance-guid.lg`. For example, `jdbc_bd2a3dd5-d571-4171-a195-28869577b87e.db` and `jdbc_bd2a3dd5-d571-4171-a195-28869577b87e.lg` are state filenames.

State files are named to be unique. These names are not intuitive. The names begin with `jdbc_` and end in `.lg` or `.db`. The rest of the filename is a GUID value that must be looked up by using a directory browser that can display it.

Defunct state files (those belonging to deleted drivers) in the state directory are deleted each time a driver instance with the same state directory is started.

Changes That Can Force Triggerless Publisher Resynchronization

If you delete state files, the triggerless publisher will build new state files by resynchronizing. If you move the JDBC driver without moving the state files, the triggerless publisher builds new state files by resynchronizing. Changing to and from the Remote Loader is a move. Therefore, if you move the JDBC driver using triggerless publication and want to avoid a full resync, also move all `jdbc_*.lg` and `jdbc_*.db` files in the state directory.

If more than two files exist in the specified state directory, you must look up the GUID to know which files belong to the driver instance being moved. To identify a driver instance's state files, you can use `DSTrace` or `DSBrowse`. For convenience, the IDM engine traces each driver's GUID in `DSTrace` on startup. You can use `Dsbrowse` to find the GUID.

If no value is provided for the state directory parameter, or the value is `'.'`, the state directory is the current directory. The current directory depends upon the following:

- ◆ The platform that the driver is running on
- ◆ Whether the driver is running locally or remotely

When a process is started, a default directory in the file system is assigned to it. The default directory is the current directory `"."` If you don't supply a value, the default State Directory is the current directory (the one that the process is running in).

Table 6-9 *Default Directories*

Platform or Environment	Default Directory
Windows, for the Remote Loader	<code>novell\remoteloader</code>

Platform or Environment	Default Directory
Windows, for Identity Manager (local; not on the Remote Loader)	c:\novell\nds\dibfiles
NetWare (local)	_netware

The current directory might be different for a custom installation.

No data is lost when resynchronization occurs, although additional data might remain. For example, because deletes are not captured, users that were deleted in the database during the move will not be disabled/deleted (depending upon policy).

Moving the driver is not to be undertaken whimsically. As a rule of thumb, don't move the driver unless you must do so.

Properties

The following table lists the properties of the State Directory parameter:

Table 6-10 *State Directory: Properties*

Property	Value
Tag Name	state-dir
Required?	no
Case-Sensitive?	platform-dependent
Sample Value	c:\novell\nds\DIBFiles
Default Value	. (current directory)

6.3.2 Database Scoping Parameters

- ◆ [“Synchronization Filter” on page 64](#)
- ◆ [“Schema Name” on page 66](#)
- ◆ [“Include Filter Expression” on page 67](#)
- ◆ [“Exclude Filter Expression” on page 67](#)
- ◆ [“Table/View Names” on page 68](#)

Synchronization Filter

The Synchronization Filter parameter determines which database objects, such as tables and views, are members of the synchronization schema (the set of tables/views visible to the driver at runtime). With the addition of this parameter, the driver can now run in two modes: schema-aware or schema-unaware.

Schema-Unaware Mode. When the Synchronization Filter parameter is present and set to empty (exclude all tables/views), the driver is schema-unaware. It does not retrieve table/view metadata on startup. Therefore, no metadata methods are required. See [Appendix D, “java.sql.DatabaseMetaData Methods,” on page 233](#).

When schema-unaware, the synchronization schema can be empty. Both the Schema Name and Sync Tables/Views parameters are completely ignored. Neither is required. Both can be absent, present, valued or valueless. See [“Schema Name” on page 66](#) and [“Table/View Names” on page 68](#).

In schema-unaware mode, the driver acts as a pass-through agent for embedded SQL. In this state, standard XDS events (for example, Add, Modify, and Delete) are ignored. See [Chapter 13, “Embedded SQL Statements in XDS Events,” on page 153](#). Also, triggered or triggerless publication no longer work.

Schema-Aware Mode. When the Synchronization Filter parameter is not present or set to a value other than empty (exclude all tables/views), the driver is schema-aware. It retrieves table/view metadata on a limited number of tables/view to facilitate data synchronization. You can cache metadata on all tables/views owned by a single database user (include by schema membership), or cache metadata on an explicit list of table/view names (include by table/view name). When schema-aware, the driver retrieves database table/view metadata on startup. For a list of required metadata methods, see [Appendix D, “java.sql.DatabaseMetaData Methods,” on page 233](#).

When schema-aware, parameter Schema Name or Table/View Names must be present and have a value. Because these two parameters are mutually exclusive, only one parameter can have a value. See [“Schema Name” on page 66](#) and [“Table/View Names” on page 68](#).

The following table lists parameters that require the driver to be schema-aware. When the driver is schema-unaware, these parameters do not have any effect on driver behavior.

Table 6-11 *Schema-Dependent Parameters*

Parameter
Lock Statement Generator Class
Enable Referential Attribute Support?
Enable Meta-Identifier Support?
Left Outer Join Operator
Retrieve Minimal Metadata
Supports Schemas in Metadata Retrieval?
Sort Column Names By
Disable Statement-Level Locking
Check Update Counts?
Add Default Values on Insert?
Generation/Retrieval Method (Table-Global)
Retrieval Timing (Table-Global)
Retrieval Timing
Disable Publisher?
Disable Statement-Level Locking?
Publication Mode

Parameter

Enable Future Event Processing?

Event Log Table Name

Delete Processed Rows?

Allow Loopback?

Startup Option

Polling Interval (In Seconds)

Publication Time of Day

Post Polling Statements

Batch Size

The following table lists the properties of this parameter:

Table 6-12 *Synchronization Filter: Properties*

Property	Value
Tag Name	sync-filter
Required?	no
Case-Sensitive?	no
Sample Value	indirect
Legal Values	empty (exclude all tables/views) schema (include by schema membership) list (include by table/view name)
Default Value:	(none)

Schema Name

The Schema Name parameter identifies the database schema being synchronized. A database schema is analogous to the name of the owner of the tables or views being synchronized. For example, to synchronize two tables, `usr` and `grp`, each belonging to database user `idm`, you enter `idm` as this parameter's value.

When using this parameter instead of Table/View Names, names of database objects are implicitly schema-qualified by the driver. As such, parameters referencing stored procedure, function, or table names do not need to be schema-qualified unless they reside in a schema other than the one specified here. In particular, Method and Timing (Table-Local) and Event Log Table Name are affected. See [“Table/View Names” on page 68](#), [“Method and Timing \(Table-Local\)” on page 85](#), and [“Event Log Table Name” on page 93](#).

The following table lists the properties of this parameter:

Table 6-13 *Schema Name: Properties*

Property	Value
Tag Name	sync-schema
Required?	yes ¹
Case-Sensitive?	See “Undelimited Identifier Case Sensitivity” on page 181 .
Sample Value	indirect
Default Value:	(none)

¹ When the Schema Name parameter is used without the Synchronization Filter parameter, the Table/View Names parameter must be left empty or omitted from a configuration. See [“Synchronization Filter” on page 64](#) and [“Table/View Names” on page 68](#).

IMPORTANT: Changing the value of the Schema Name parameter forces a resync of all objects when triggerless publication is used.

Include Filter Expression

The Include Filter Expression parameter is only operative when the Schema Name parameter is used. See [“Schema Name” on page 66](#).

The following table lists the properties of this parameter:

Table 6-14 *Include Filter Expression: Properties*

Property	Value
Tag Name	include-table-filter
Required?	no
Case-Sensitive?	yes
Sample Value	idm_*. (all table/view names starting with “idm_”)
Default Value	(none)
Legal Values	(any legal Java regular expression)

Exclude Filter Expression

This parameter is only operative when the Schema Name parameter is used. See [“Schema Name” on page 66](#).

The following table lists the properties of this parameter:

Table 6-15 Exclude Filter Expression: Properties

Property	Value
Tag Name	exclude-table-filter
Required?	no
Case-Sensitive?	yes
Sample Value	bin*. (all table/view names starting with "bin")
Default Value	(none)
Legal Values	(any legal Java regular expression)

Table/View Names

The Table/View Names parameter allows you to create a logical database schema by listing the names of the logical database classes to synchronize. Logical database class names are the names of parent tables and views. It is an error to list child table names.

This parameter is particularly useful for synchronizing with databases that do not support the concept of schema, such as MySQL, or when a database schema contains a large number of tables or views of which only a few are of interest. Reducing the number of table/view definitions cached by the driver can shorten startup time as well as reduce runtime memory utilization.

When using this parameter instead of Schema Name, you likely need to schema-qualify other parameters that reference stored procedure, function, or table names. In particular, the Method and Timing (Table-Local) and Event Log Table Name parameters are affected. See [“Schema Name” on page 66](#), [“Method and Timing \(Table-Local\)” on page 85](#) and [“Event Log Table Name” on page 93](#).

The following table lists the properties of this parameter:

Table 6-16 Table/View Names: Properties

Property	Value
Tag Name	sync-tables
Required?	yes ¹
Case-Sensitive?	See “Undelimited Identifier Case Sensitivity” on page 181 .
Delimiters	semicolon, white space, comma
Sample Value	indirect.usr; indirect.grp
Default Value	(none)

¹When this parameter is used without the Synchronization Filter parameter, the Schema Name parameter must be left empty or omitted from a configuration. See [“Synchronization Filter” on page 64](#) and [“Schema Name” on page 66](#).

IMPORTANT: Changing anything in the Table/View Name parameter other than URL properties forces a resynchronization of all objects when triggerless publication is used.

6.3.3 Connectivity Parameters

- ♦ “Use Minimal Number of Connections?” on page 69
- ♦ “Connection Initialization Statements” on page 69
- ♦ “Connection Properties” on page 70

Use Minimal Number of Connections?

The Use Minimal Number of Connections? parameter specifies whether the driver should use two instead of three database connections.

By default, the driver uses three connections: one for subscription, and two for publication. The Publisher channel uses one of its two connections to query for events and the other to facilitate query-back operations.

When this parameter is set to Boolean True, the number of required database connections is reduced to two. One connection is shared between the Subscriber and Publisher channels. It is used to process subscription and publication query-back events. The other is used to query for publication events.

In previous versions, the driver was able to support bidirectional synchronization by using a single connection. The publication algorithm was redesigned to increase performance, enable support for future event processing, and to overcome limitations of the previous algorithm at the expense of requiring an additional connection.

Table 6-17 Use Minimal Number of Connections?: Properties

Property	Value
Tag Name	use-single-connection
Required?	no
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	False

¹This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is Boolean False.

NOTE: Setting this parameter to Boolean True reduces performance.

Connection Initialization Statements

The Connection Initialization Statements parameter specifies what SQL statements, if any, should be executed immediately after connecting to the target database. Connection initialization

statements are useful for changing database contexts and setting session properties. These statements are executed each time the driver, irrespective of channel, connects or reconnects to the target database.

The following table lists the properties of this parameter:

Table 6-18 *Connection Initialization Statements: Properties*

Property	Value
Tag Name	connection-init
Required?	no
Case-Sensitive?	See “Undelimited Identifier Case Sensitivity” on page 181 .
Delimiters	semicolon
Sample Value	USE idm; SET CHAINED OFF
Default Value	(none)
Schema-Dependent	False

Connection Properties

The Connection Properties parameter specifies authentication properties. This parameter is useful for specifying properties that cannot be set via the JDBC URL specified in the Authentication Context parameter. See [“Authentication Context” on page 58](#).

The primary purpose of this parameter is to enable encrypted transport for third-party JDBC drivers. For a list of relevant connection properties, see [“Sybase Adaptive Server Enterprise JConnect JDBC Driver” on page 201](#) and [“Oracle Thin Client JDBC Driver” on page 197](#).

Connection properties are specified as key-value pairs. The key is specified as the value to the left of the “=” character. The value is the value to the right of the “=” character. You can specify multiple key-value pairs, but each pair must be delimited by the “;” character.

When you use the Connection Properties parameter, authentication information can be passed via the JDBC URL specified in the Authentication Context parameter or here. See [“Authentication Context” on page 58](#).

If specified as connection properties, value tokens can be used as placeholders for the database username specified in the Authentication ID parameter and the password specified in the Application Password parameter. See [“Authentication ID” on page 58](#) and [“Application Password” on page 58](#). For username, the token is { \$username }. For password, the token is { \$password }.

The following table lists the properties of this parameter:

Table 6-19 *Connection Properties: Properties*

Property	Value
Tag Name	connection-properties

Property	Value
Required?	no
Case-Sensitive?	third-party JDBC driver-dependent
Delimiters	semicolon
Sample Value	user=\${username}; password=\${password}; oracle.jdbc.defaultNChar=true
Default Value	(none)
Schema-Dependent	False

6.3.4 Compatibility Parameters

- ◆ “JDBC Driver Descriptor Filename” on page 71
- ◆ “Database Descriptor Filename” on page 72
- ◆ “Use Manual Transactions?” on page 72
- ◆ “Transaction Isolation Level” on page 73
- ◆ “Reuse Statements?” on page 74
- ◆ “Number of Returned Result Sets” on page 75
- ◆ “Enable Statement-Level Locking?” on page 75
- ◆ “Lock Statement Generator Class” on page 76
- ◆ “Enable Referential Attribute Support?” on page 76
- ◆ “Enable Meta-Identifier Support?” on page 77
- ◆ “Force Username Case” on page 77
- ◆ “Left Outer Join Operator” on page 78
- ◆ “Retrieve Minimal Metadata” on page 78
- ◆ “Function Return Method” on page 79
- ◆ “Supports Schemas in Metadata Retrieval?” on page 79
- ◆ “Sort Column Names By” on page 80

JDBC Driver Descriptor Filename

The JDBC Driver Descriptor Filename parameter specifies the third-party JDBC descriptor file to use. Descriptor file names must not be prefixed with the underscore character (for example, `_mysql_jdriver.xml`) because such filenames are reserved. Place descriptor files in a jar file beginning with the case-insensitive prefix “jdbc” (for example, `JDBCCustomConfig.jar`) and in the jar file’s `com/novell/nds/dirxml/driver/jdbc/db/descriptor/driver` directory.

The following table lists the properties of this parameter:

Table 6-20 *JDBC Driver Descriptor Filename: Properties*

Property	Value
Tag Name	jdbc-driver-descriptor
Required?	no
Case-Sensitive?	platform-dependent
Sample Value	my_custom_jdbc_driver_descriptor.xml
Default Value	(none)
Schema-Dependent	False

Database Descriptor Filename

The Database Descriptor Filename parameter specifies the database descriptor file to use. Do not use the underscore character in prefixes to Descriptor filenames (for example, `_mysql.xml`). Such names are reserved. Place Descriptor files in a jar file beginning with the case-insensitive prefix “jdbc” (for example, `JDBCCustomConfig.jar`). Also, place Descriptor files in the jar file’s `com/novell/nds/dirxml/driver/jdbc/db/descriptor/db` directory.

The following table lists the properties of this parameter:

Table 6-21 *Database Descriptor Filename: Properties*

Property	Value
Tag Name	jdbc-driver-descriptor
Required?	no
Case-Sensitive?	platform-dependent
Sample Value	my_custom_database_descriptor.xml
Default Value	(none)
Schema-Dependent	False

Use Manual Transactions?

The Use Manual Transactions? parameter specifies whether to use manual or user-defined transactions.

This parameter is primarily used to enable interoperability with MySQL MyISAM table types, which do not support transactions.

When set to Boolean True, the driver uses manual transactions. When set to Boolean False, each statement executed by the driver is executed autonomously (automatically).

The following table lists the properties of this parameter:

Table 6-22 *Use Manual Transactions?: Properties*

Property	Value
Tag Name	use-manual-transactions
Required?	no
Case-Sensitive?	no
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	False

¹This default is derived dynamically from descriptor files and database metadata at runtime.

NOTE: To ensure data integrity, set this parameter to Boolean True whenever possible.

Transaction Isolation Level

The Transaction Isolation Level parameter sets the transaction isolation level for connections that the driver uses. Six values exist:

- ◆ `unsupported`
- ◆ `none`
- ◆ `read uncommitted`
- ◆ `read committed`
- ◆ `repeatable read`
- ◆ `serializable`

Five of the values correspond to the public constants defined in the [java.sql Interface Connection](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html>).

Because some third-party drivers do not support setting a connection's transaction isolation level to `none`, the driver also supports the additional non-standardized value of `unsupported`.

[PostgreSQL online documentation](http://www.postgresql.org/docs/current/static/transaction-iso.html) (<http://www.postgresql.org/docs/current/static/transaction-iso.html>) has one of the better, concise primers on what each isolation level actually means.

IMPORTANT: The list of supported isolation levels varies by database. For a list of supported transaction isolation levels for supported databases, see “[Supported Transaction Isolation Levels](#)” on [page 181](#).

We recommend using a transaction isolation level of `read committed` because it is the minimum isolation level that prevents the driver from seeing uncommitted changes (dirty reads).

The following table lists the properties of this parameter:

Table 6-23 *Transaction Isolation Level: Properties*

Property	Value
Tag Name	transaction-isolation-level
Required?	no
Case-Sensitive?	no
Default Value	(dynamic ¹)
Legal Values	unsupported none read uncommitted read committed repeatable read serializable
Schema-Dependent	False

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is read committed.

Reuse Statements?

The Reuse Statements? parameter specifies whether one or more `java.sql.Statement` items are active at a time on a given connection. See [java.sql.Statement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html>).

This parameter is primarily used to enable interoperability with **Microsoft SQL Server 2000 Driver for JDBC**.

When set to Boolean True, the driver allocates a Java SQL Statement once and then reuses it. When set to Boolean False, the driver allocates/deallocates statement objects each time they are used, ensuring that no more than one statement is active at a time on a given connection.

The following table lists the properties of this parameter:

Table 6-24 *Reuse Statements?: Properties*

Property	Value
Tag Name	reuse-statements
Required?	no
Case-Sensitive?	no
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	False

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is Boolean True.

NOTE: Setting this parameter to Boolean False degrades performance.

Number of Returned Result Sets

The Number of Returned Result Sets parameter specifies how many `java.sql.Result` objects can be returned from an arbitrary SQL statement. See `java.sql.ResultSet` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSet.html>).

This parameter is primarily used to avoid infinite loop conditions in “Oracle Thin Client JDBC Driver” on page 197 when evaluating the results of arbitrary SQL statements.

The following table lists the properties of this parameter:

Table 6-25 *Number of Returned Result Sets: Properties*

Property	Value
Tag Name	handle-stmt-results
Required?	no
Sample Value	one
Default Value	(dynamic ¹)
Legal Values	none, no (none) single, one (one) multiple, many, yes (multiple)
Schema-Dependent	False

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is multiple, many, or yes.

Enable Statement-Level Locking?

The Enable Statement-Level Locking? parameter specifies whether the driver explicitly locks database resources before executing SQL statements.

The following table lists the properties of this parameter:

Table 6-26 *Enable Statement-Level Locking?: Properties*

Property	Value
Tag Name	enable-locking
Required?	no
Default Value	0 (no)

Property	Value
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Lock Statement Generator Class

The Lock Statement Generator Class parameter specifies which DBLockStatementGenerator implementation to use to generate the SQL statements necessary to explicitly lock database resources for a pending SQL statement. Information on the DBLockStatementGenerator interface is in the Java documents that ship with the driver.

The following table lists the properties of this parameter:

Table 6-27 Lock Statement Generator Class: Properties

Property	Value
Tag Name	lock-generator-class
Required?	no
Sample Value	com.novell.nds.dirxml.driver.jdbc.db.lock.OraLockGenerator
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is com.novell.nds.dirxml.driver.jdbc.db.lock.DBLockGenerator.

Enable Referential Attribute Support?

The Enable Referential Attribute Support? parameter toggles whether the driver recognizes foreign key constraints between logical database classes. These are used to denote containment. Foreign key constraints between parent and child tables within a logical database class are unaffected.

When set to Boolean True, foreign key columns are interpreted as referential. When set to Boolean False, foreign key columns are interpreted as non-referential.

The primary purpose of this parameter is to ensure backward compatibility with the 1.0 version of the driver. For 1.0 compatibility, set this parameter to Boolean False.

The following table lists the properties of this parameter:

Table 6-28 Enable Referential Attribute Support?: Properties

Property	Value
Tag Name	enable-refs

Property	Value
Required?	no
Default Value	1 (yes)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Enable Meta-Identifier Support?

The Enable Meta-Identifier Support? parameter toggles whether the driver interprets view column name prefixes such as “pk_” and “fk_” strictly as metadata. When interpreted as metadata, such prefixes are not considered part of the view column name.

For example, when meta-identifier support is enabled, column “pk_idu” has an effective column name of “idu,” prohibiting the existence of another column with the same effective name in the same view. When meta-identifier support is disabled, column “pk_idu” has the effective column name of “pk_idu,” allowing the existence of another column named “idu.” Furthermore, when meta-identifier support is enable, a view with a primary key named “pk_idu” would conflict with a table having a primary key column named “idu.” When meta-identifier support is disabled, they would not conflict.

When set to Boolean True, view column prefixes are interpreted as metadata. When set to Boolean False, view column name prefixes are interpreted as part of the column name proper.

The primary purpose of this parameter is to ensure backward compatibility with the 1.5 version of the driver. For 1.5 compatibility, set this parameter to Boolean False.

The following table lists the properties of this parameter:

Table 6-29 *Enable Meta-Identifier Support?: Properties*

Property	Value
Tag Name	enable-meta-identifiers
Required?	no
Default Value	1 (yes)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Force Username Case

The Force Username Case parameter changes the case of the driver’s username used to authenticate to the target database.

The primary purpose of this parameter is to enable interoperability with the Informix JDBC Driver when used against ANSI-compliant databases. See [“Informix JDBC Driver” on page 194](#).

The following table lists the properties of this parameter:

Table 6-30 *Force Username Case: Properties*

Property	Value
Tag Name	force-username-case
Required?	no
Default Value	(don't force)
Legal Values	lower (to lowercase) mixed (to mixed case) upper (to uppercase)
Schema-Dependent	False

Left Outer Join Operator

The Left Outer Join Operator parameter specifies the left outer join operator used in the triggerless publication query. It might be used for other purposes in the future.

The following table lists the properties of this parameter:

Table 6-31 *Left Outer Join Operator: Properties*

Property	Value
Tag Name	left-outer-join-operator
Required?	no
Default Value	(dynamic ¹)
Legal Values	*= (+) LEFT OUTER JOIN
Schema-Dependent	True

¹This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is LEFT OUTER JOIN.

Retrieve Minimal Metadata

When set to Boolean True, the driver calls only required metadata methods. When set to Boolean False, the driver calls required and optional metadata methods. For a list of required and optional metadata methods, refer to [Appendix D, “java.sql.DatabaseMetaData Methods,” on page 233](#).

Optional metadata methods are required for multivalue and referential attribute synchronization.

Table 6-32 Retrieve Minimal Metadata: Properties

Property	Value
Tag Name	minimal-metadata
Required?	no
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is Boolean False.

NOTE: Setting this value to Boolean True improves startup time and third-party JDBC driver compatibility at the expense of functionality.

Function Return Method

The Function Return Method parameter specifies how data is retrieved from database functions.

The primary purpose of this parameter is to enable interoperability with the Informix JDBC driver. See [“Informix JDBC Driver” on page 194](#).

When set to `result set`, function results are retrieved through a result set. When set to `return value`, the function result is retrieved as a single, scalar return value.

Table 6-33 Function Return Method: Properties

Property	Value
Tag Name	function-return-method
Required?	no
Default Value	(dynamic ¹)
Legal Values	result set return value (scalar return value)
Schema-Dependent	False

¹ This default is derived dynamically from descriptor files at runtime.

Supports Schemas in Metadata Retrieval?

The Supports Schemas in Metadata Retrieval? parameter specifies whether schema names should be used when retrieving database metadata.

The primary purpose of this parameter is to enable interoperability with the Informix JDBC Driver when used against ANSI-compliant databases. See [“Informix JDBC Driver” on page 194](#).

When set to Boolean True, schema names are used. When set to Boolean False, they are not.

Table 6-34 *Supports Schemas in Metadata Retrieval?: Properties*

Property	Value
Tag Name	supports-schemas-in-metadata-retrieval
Required?	no
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	False

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is Boolean True.

Sort Column Names By

The Sort Column Names By parameter specifies how column position is to be determined for legacy databases that do not support sorting by column names.

The primary purpose of this parameter is to enable interoperability with legacy databases, such as DB2/AS400.

Sorting columns names by hexadecimal value ensures that if a driver instance is relocated to a different server, it continues to function without modification. Sorting column names by platform or locale string collation order is more intuitive, but might require configuration changes if a driver instance is relocated to a different server. In particular, log table column order and compound column name order might change. In the case of the latter, Schema-Mapping policies and object association values might need to be updated. In the case of the former, log table columns might have to be renamed.

It is also possible to specify any fully-qualified Java class name as long as the following occur:

- ◆ The Java class name implements the [java.util.Comparator](http://java.sun.com/j2se/1.5.0/docs/api/java/util/Comparator.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/util/Comparator.html>) interface.
- ◆ The Java class name accepts [java.lang.String](http://java.sun.com/j2se/1.5.0/docs/api/java/lang/String.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/lang/String.html>) arguments.
- ◆ The class is in the runtime classpath.

Table 6-35 *Sort Column Names By: Properties*

Property	Value
Tag Name	column-position-comparator
Required?	no
Default Value	(dynamic ¹)

Property	Value
Legal Values	com.novell.nds.dirxml.driver.jdbc.util.config.comp.StringByteComparator (hexadecimal value) com.novell.nds.dirxml.driver.jdbc.util.config.comp.StringComparator (string collation order) (any java.util.Comparator that accepts java.lang.String arguments)
Schema-Dependent	True

¹ This default is derived dynamically from descriptor files at runtime. Otherwise, the default value is com.novell.nds.dirxml.driver.jdbc.util.config.comp.StringByteComparator.

IMPORTANT: After you set this parameter for a given configuration, don't change the parameter.

6.4 Subscription Parameters

The following table summarizes Subscriber-level parameters and their properties:

Table 6-36 *Subscriber-Level Parameters and Properties*

Display Name	Tag Name	Sample Value	Default Value	Required
Disable Subscriber?	disable	1 (yes)	0 (no)	no
Generation/Retrieval Method (Table-Global)	key-gen-method	auto	none (subscription event)	
Retrieval Timing (Table-Global)	key-gen-timing	after (after row insertion)	before (before row insertion)	no
Method and Timing (Table-Local)	key-gen	usr("=?=indirect.proc_idu()", before)	(none)	no
Disable Statement-Level Locking?	disable-locking	1 (yes)	0 (no)	no
Check Update Counts?	check-update-count	0 (no)	1 (yes)	no
Add Default Values on Insert?	add-default-values-on-view-insert	0 (no)	(dynamic ¹)	no

¹ This default is derived dynamically from descriptor files at runtime.

Subscription parameters are in two subcategories:

- ◆ [Section 6.4.1, “Uncategorized Parameters,” on page 82](#)
- ◆ [Section 6.4.2, “Primary Key Parameters,” on page 84](#)

6.4.1 Uncategorized Parameters

- ◆ “Disable Subscriber?” on page 82
- ◆ “Disable Statement-Level Locking?” on page 82
- ◆ “Check Update Counts?” on page 83
- ◆ “Add Default Values on Insert?” on page 83

Disable Subscriber?

The Disable Subscriber? parameter specifies whether the Subscriber channel is disabled.

When this parameter is set to Boolean True, the Subscriber channel is disabled. When the parameter is set to Boolean False, the Subscriber channel is active.

Table 6-37 *Disable Subscriber?: Properties*

Property	Value
Tag Name	disable
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	False

Disable Statement-Level Locking?

The Disable Statement-Level Locking? parameter specifies whether database resources are explicitly locked on this channel before each SQL statement is executed. This parameter is active only if [Enable Statement-Level Locking?](#) is set to Boolean True.

When this parameter is set to Boolean True, database resources are explicitly locked. When this parameter is set to Boolean False, database resources are not explicitly locked.

Table 6-38 *Disable Statement-Level Locking?: Properties*

Property	Value
Tag Name	disable-locking
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Check Update Counts?

The Check Update Counts? parameter specifies whether the Subscriber channel checks to see if a table was actually updated when INSERT, UPDATE, and DELETE statements executed against a table.

When set to Boolean True, update counts are checked. If nothing is updated, an exception is thrown. When set to Boolean False, update counts are ignored.

When statements are redefined in before-trigger logic, set his parameter to Boolean False

When using Microsoft SQL Server, use the default value, because errors in trigger logic (that might roll back a transaction) are not propagated back to the Subscriber channel.

Table 6-39 *Check Update Counts?: Properties*

Property	Value
Tag Name	check-update-count
Required?	no
Default Value	1 (yes)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Add Default Values on Insert?

The Add Default Values on Insert? parameter specifies whether the Subscriber channel provides default values when executing an INSERT statement against a view.

The primary purpose of this parameter is to enable interoperability with Microsoft SQL Server 2000. This database requires that view columns constrained NOT NULL have a non-NULL value in an INSERT statement.

When this parameter is set to Boolean True, default values are provided for INSERT statements executed against views, and explicit values are not already available. When this parameter is set to Boolean False, default values are not provided.

Table 6-40 *Add Default Values on Insert?: Properties*

Property	Value
Tag Name	add-default-values-on-view-insert
Required?	no
Default Value	(dynamic ¹)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

¹ This default is derived dynamically from descriptor files at runtime.

6.4.2 Primary Key Parameters

- ♦ [“Generation/Retrieval Method \(Table-Global\)” on page 84](#)
- ♦ [“Retrieval Timing \(Table-Global\)” on page 85](#)
- ♦ [“Method and Timing \(Table-Local\)” on page 85](#)

When processing `<add>` events, which map to `INSERT` statements, the Subscriber channel uses primary key values to create Identity Manager associations. These parameters specify how and when the Subscriber channel obtains the primary key values necessary to construct association values. How primary key values are obtained is the primary key generation/retrieval method. The retrieval timing indicates when primary key values are retrieved.

The following table identifies the supported methods and timings:

Table 6-41 *Supported Methods and Timings*

Method	Timing: before (row insertion)	Timing: after (row insertion)
None (subscription event)	X	0 ¹
Driver (Subscriber-generated)	X	X
Auto (auto-generated/identity column) (stored procedure/function)	0 ² X	X X

¹ The Subscriber channel automatically overrides this timing to `before`.

² The Subscriber channel automatically overrides this timing to `after`.

Generation/Retrieval Method (Table-Global)

The Generation/Retrieval Method (Table-Global) parameter specifies how primary key values are generated or retrieved for all parent tables and views. The Method and Timing parameter overrides this parameter on a per-table/view basis. See [“Method and Timing \(Table-Local\)” on page 85](#).

When this parameter is set to `none`, primary key values are assumed to already exist in the subscription event. When this parameter is set to `driver`, primary key values are generated by one of the following:

- ♦ Using a `SELECT (MAX () +1)` statement if retrieval timing is set to `before`
- ♦ Using a `SELECT MAX ()` statement if retrieval timing is set to `after`

For string column types, the Subscriber channel generates a value by using the return value of `System.currentTimeMillis()`. Other data types are not supported.

When this parameter is set to `auto`, primary key values are retrieved via the `java.sql.Statement.getGeneratedKeys () : java.sql.ResultSet` method. The MySQL Connector/J JDBC driver is the only supported third-party JDBC driver that currently implements this method. See [“MySQL Connector/J JDBC Driver” on page 196](#).

Table 6-42 *Generation/Retrieval Method (Table-Global): Properties*

Property	Value
Tag Name	key-gen-method
Required?	no
Default Value	none (subscription event)
Legal Values	none (subscription event) driver (Subscriber-generated) auto (auto-generated/identity column)
Schema-Dependent	True

Retrieval Timing (Table-Global)

The Retrieval Timing (Table-Global) parameter specifies when the Subscriber channel retrieves primary key values for all parent tables and views. The parameter Method and Timing (Table-Local) overrides this parameter. See [“Method and Timing \(Table-Local\)” on page 85](#).

When this parameter is set to `before`, primary key values are retrieved before insertion. When this parameter is set to `after`, primary key values are retrieved after insertion.

Table 6-43 *Retrieval Timing (Table-Global): Properties*

Property	Value
Tag Name	key-gen-timing
Required?	no
Default Value	before (before row insertion)
Legal Values	before (before row insertion) after (after row insertion)
Schema-Dependent	True

Method and Timing (Table-Local)

The Method and Timing (Table-Local) parameter specifies the primary key generation/retrieval method and retrieval timing on a per parent table/view basis. It essentially maps a generation/retrieval method and retrieval timing to a table or view name. The syntax for this parameter mirrors a procedural programming language method call with multiple arguments (such as, *method-name(argument1, argument2)*).

When using the [Table/View Names](#) parameter, you probably need to explicitly schema-qualify any tables, views, stored procedures or functions referenced in this parameter’s value. When you use the [Schema Name](#) parameter, tables, views, stored procedures, or functions referenced in this parameter’s value are implicitly schema-qualified with that schema name. If tables, views, stored procedures, or functions referenced in this parameter’s value are located in a different schema other than the implicit schema, they must be schema-qualified.

BNF

The BNF (Backus Naur Form (<http://cui.unige.ch/db-research/Enseignement/analyseinfo/AboutBNF.html>)) notation for this parameter's value is the following:

```
<key-gen> ::= <table-or-view-name> " (" <generation-retrieval-method>,  
  
                                <retrieval-timing> ")"  
{[<delimiter>] <key-gen>}  
  
<generation-retrieval-method> ::= none | driver | auto |  
                                "" <procedure-signature> "" |  
                                "" <function-signature> ""  
  
<table-or-view-name> ::= <legal-undelimited-database-table-or-view-  
                           identifier>  
  
<delimiter> ::= ";" | "," | <white-space>  
  
<procedure-signature> ::= <schema-qualifier> "." <stored-routine-  
                           name> "(" <argument-list> ")"  
  
<function-signature> ::= "?" <procedure-signature>  
  
<schema-qualifier> ::= <legal-undelimited-database-username-  
                           identifier>  
  
<stored-routine-name> ::= <legal-undelimited-database-stored-routine-  
                           identifier>  
  
<argument-list> ::= <column-name>{"," <column-name>}  
  
<column-name> ::= <column-from-table-or-view-name-previously-  
                           specified>
```

Generation or Retrieval Method

The generation or retrieval method specifies how primary key values are to be generated, if necessary, and retrieved. The possible methods are None, Driver, Auto, and Stored Procedure/Function:

None

By default, the Subscriber channel assumes that the Identity Vault is the authoritative source of primary key values and that the requisite values are already present in a given <add> event. If this is the case, no primary values need to be generated because they already exist. They only need to be retrieved from the current <add> event. This method is desirable when an eDirectory attribute, such as GUID, is explicitly schema-mapped to a parent table or view's primary key column.

Assuming the existence of a table named `usr` and a view named `view_usr` where the Identity Vault is the authoritative source of primary key values, this parameter's value would look something like the following:

```
usr (none); view_usr (none)
```

When you use this method, we recommend mapping GUID rather than CN to a parent table or view's primary key column.

Driver

This method assumes that the database is the authoritative source of primary key values for the specified parent table or view.

When prototyping or in the initial stages of deployment, it is often desirable to have the Subscriber channel generate primary key values before a stored procedure or function is written. You can also use this method against databases that do not support stored procedures or functions. When you use this method in a production environment, however, all SQL statements generated by an `<add>` event should be contained in a serializable transaction. For additional information, refer to [“Transaction Isolation Level” on page 73](#).

Instead of making all transactions serializable, you can also set individual transaction isolation levels by using embedded SQL attributes. For additional information, refer to [Section 13.6, “Transaction Isolation Level,” on page 160](#).

For any numeric column types, the Subscriber channel uses the following to generate primary key values:

- ◆ A simple `SELECT (MAX+1)` statement for `before` timing
- ◆ A `SELECT MAX()` statement for `after` timing

For string column types, the Subscriber channel generates a value by using the return value of `System.currentTimeMillis()`. Other data types are not supported.

Assuming the existence of a table named `usr` and a view named `view_usr`, where the database is the authoritative source of primary key values, this parameter’s value would look something like the following:

```
usr(driver); view_usr(driver)
```

When you use this method, we recommend that you omit primary key columns from Schema Mapping policies and channel filters.

Auto

This method assumes that the database is the authoritative source of primary key values for the specified parent table or view.

Some databases support identity columns that automatically generate primary key values for inserted rows. This method retrieves auto-generated primary key values through the JDBC 3 interface method

```
java.sql.Statement.getGeneratedKeys():java.sql.ResultSet.
```

The MySQL Connector/J JDBC driver is the only supported third-party JDBC driver that currently implements this method. See [“MySQL Connector/J JDBC Driver” on page 196](#).

Assuming the existence of a table named `usr` and a view named `view_usr`, where the database is the authoritative source of primary key values, this parameter’s value would look something like the following:

```
usr(auto); view_usr(auto)
```

When you use this method, we recommend that you omit primary key columns from Schema Mapping policies and channel filters.

Stored-Procedure/Function:

This method assumes that the database is the authoritative source of primary key values for the specified parent table or view.

Assuming

- ◆ The existence of a table named `usr` with a primary key column named `idu`
- ◆ A view named `view_usr` with a primary key values named `pk_idu`
- ◆ The existence of a database function `func_last_usr_idu` and stored procedure `sp_last_view_usr_pk_idu` that both return the last generated primary key value for their respective table/view

this parameter's value would look something like the following:

```
usr("?=func_last_usr_idu()");  
view_usr("sp_last_view_usr_pk_idu(pk_idu)")
```

In the previous examples, a parameter is passed to the stored procedure. Parameters can also be passed to functions, but this is not usually necessary. Unlike functions, stored procedures usually return values through parameters. For stored procedures, primary key columns must be passed as `IN OUT` parameters. Non-key columns must be passed as `IN` parameters.

For both stored procedures and functions, parameter order, number and data type must correspond to the order, number and data type of the parameters expected by the procedure or function.

When you use this method, we recommend that you omit primary key columns from Schema Mapping policies and channel filters.

Retrieval Timing

The Retrieval Timing parameter specifies when primary key values are retrieved.

An `<add>` event always results in at least one `INSERT` statement against a parent table or view. This portion of this parameter specifies when primary key values are to be retrieved relative to the initial `INSERT` statement.

Before

This is the default setting. When this setting is specified, primary key values are retrieved before the initial `INSERT` statement.

IMPORTANT: This retrieval timing is supported for all generation/retrieval methods except `auto`. Retrieval timing is required for the `none` method.

After

When this setting is specified, primary key values are retrieved after the initial `INSERT` statement.

IMPORTANT: This retrieval timing is supported for all generation/retrieval methods except `none`. Retrieval timing is required for the `auto` method.

The following examples augment the previous ones by adding retrieval timing information:

```
usr(none, before); view_usr(none, before)  
usr(driver, before); view_usr(driver, after)  
usr(auto, after); view_usr(auto, after)  
  
usr("?=func_last_usr_idu()", before);  
view_usr("sp_last_view_usr_pk_idu(pk_idu)", after)
```


The following table lists the properties of this parameter:

Table 6-44 Retrieval Timing: Properties

Property	Value
Tag Name	key-gen
Required?	no
Case-Sensitive?	See “ Undelimited Identifier Case Sensitivity ” on page 181.
Sample Value	usr(“?=proc_idu()”, before)
Default Value	(none)
Legal Values	(any string adhering to the BNF)
Schema-Dependent	True

6.5 Publication Parameters

The following table summarizes publisher-level parameters and their properties:

Table 6-45 Publisher-Level Parameters and Properties

Display Name	Tag Name	Sample Value	Default Value	Required
Disable Publisher?	disable	1 (yes)	0 (no)	no
Disable Statement-Level Locking?	disable-locking	1 (yes)	0 (no)	no
Publication Mode	publication-mode	2 (triggerless)	1 (triggered)	no
Event Log Table Name	log-table	indirect_process	(none)	yes ¹
Delete Processed Rows?	delete-from-log	0 (no)	1 (yes)	no
Allow Loopback?	allow-loopback	1 (yes)	0 (no)	no
Enable Future Event Processing?	handle-future-events	1 (yes)	0 (no)	no
Startup Option	startup-option			no
Polling Interval (In Seconds)	polling-interval	60	10	no ²
Publication Time of Day	time-of-day	15:30:00	(none)	no ²
Post Polling Statements	post-poll-stmt	DELETE FROM direct.direct_process	(none)	no
Batch Size	batch-size	16	1	no
Heartbeat Interval (In Minutes)	pub-heartbeat-interval	10	0	no

¹ Required for triggered publication mode.

² These parameters are mutually exclusive.

Publication parameters fall into four major subcategories:

- ◆ [Section 6.5.1, “Uncategorized Parameters,” on page 90](#)
- ◆ [Section 6.5.2, “Triggered Publication Parameters,” on page 93](#)
- ◆ [Section 6.5.3, “Triggerless Publication Parameters,” on page 95](#)
- ◆ [Section 6.5.4, “Polling Parameters,” on page 95](#)

6.5.1 Uncategorized Parameters

- ◆ [“Disable Publisher?” on page 90](#)
- ◆ [“Disable Statement-Level Locking?” on page 90](#)
- ◆ [“Publication Mode” on page 91](#)
- ◆ [“Enable Future Event Processing?” on page 91](#)

Disable Publisher?

The `Disable Publisher?` parameter specifies whether the Publisher channel is disabled. When disabled, the Publisher channel does not query for database events. Unlike with the [Disable Subscriber?](#) parameter, you can still issue database queries on the Publisher channel to facilitate alternative publication algorithms.

When this parameter is set to Boolean True, the Publisher channel is disabled. When this parameter is set to Boolean False, the Publisher channel is active.

Table 6-46 *Disable Publisher?: Properties*

Property	Value
Tag Name	disable
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Disable Statement-Level Locking?

The `Disable Statement-Level Locking?` parameter specifies whether database resources should be explicitly locked on this channel before each SQL statement is executed. This parameter is only active if the [Enable Statement-Level Locking?](#) parameter is set to Boolean True.

When this parameter is set to Boolean True, database resources are explicitly locked. When this parameter is set to Boolean False, database resources are not explicitly locked.

Table 6-47 *Disable Statement-Level Locking?: Properties*

Property	Value
Tag Name	disable-locking
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

Publication Mode

The Publication Mode parameter specifies which publication algorithm is used.

When set to 1 (triggered), the Publisher channel polls the event log table for events. When set to 2 (triggerless), the Publisher channel dredges all tables/views in the synchronization schema for changes, and synthesizes events.

The following table lists the properties of this parameter:

Table 6-48 *Publication Mode: Properties*

Property	Value
Tag Name	publication-mode
Required?	no
Default Value	1 (triggered)
Legal Values	1 (triggered) 2 (triggerless)
Schema-Dependent	True

Enable Future Event Processing?

For triggered publication, Enable Future Event Processing? specifies whether rows in the event log table are ordered and processed by insertion order (the `record_id` column) or chronologically (the `event_time` column).

When this parameter is set to Boolean False, rows in the event log table are published by order of insertion. When this parameter is set to Boolean True, rows in the event log table are published chronologically.

For triggerless publication, Enable Future Event Processing specifies whether database local time is published with each event. This additional information can be used to force a retry of future-dated events. In order for this to work, a column specifying when an event should be processed must be part of each logical database class utilizing this feature and placed in the Publisher filter as a notification-only attribute.

Database local time is published as an attribute on each XDS event (for example, add, modify, delete). The attribute name is `jdbc:database-local-time`, where the `jdbc` namespace prefix is bound to `urn:dirxml:jdbc`. The format is the Java string representation of a `java.sql.Timestamp`: `yyyy-mm-dd hh:mm:ss.ffffffffff`. Depending upon the value of the Time Syntax parameter, the value indicating when an event should be processed can be published as an integer, as a canonical string, or as a Java string. See “Time Syntax” on page 61.

Regardless of the publication syntax, this value can be parsed and compared to the database local time value. The following table maps the time syntax to the appropriate parse method.

Table 6-49 Mapping Time Syntax to Parse Methods

Time Syntax	Parse Method
integer	<code>java.sql.Timestamp(long)</code> (http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html)
canonical string	<code>com.novell.nds.dirxml.driver.jdbc.db.DSTime(java.lang.String, java.lang.String, java.lang.String, java.lang.String)</code>
java string	<code>java.sql.Timestamp.valueOf(java.lang.String):java.sql.Timestamp</code> (http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html)

After both time values are in a common Timestamp object representation, they can be compared by using the following methods:

- ◆ `com.novell.nds.dirxml.driver.jdbc.db.TimestampUtil.before(java.sql.Timestamp, java.sql.Timestamp):boolean`
- ◆ `com.novell.nds.dirxml.driver.jdbc.db.TimestampUtil.after(java.sql.Timestamp, java.sql.Timestamp):boolean`

An example policy is provided in [Appendix J, “Policy Example: Triggerless Future Event Processing,”](#) on page 249.

When this parameter is set to Boolean True, local database time is published with each event. When this parameter is set to Boolean False, this information is omitted.

The following table lists the properties of this parameter:

Table 6-50 Enable Future Event Processing?: Properties

Property	Value
Tag Name	handle-future-events
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

6.5.2 Triggered Publication Parameters

The Driver for JDBC can use any of four triggered publication parameters.

- ♦ “Event Log Table Name” on page 93
- ♦ “Delete Processed Rows?” on page 93
- ♦ “Allow Loopback?” on page 94

Event Log Table Name

The Event Log Table Name parameter specifies the name of the event log table where publication events are stored.

The table specified here must conform to the definition of [Chapter 12, “The Event Log Table,” on page 143](#).

When using [“Table/View Names” on page 68](#), you’ll probably need to explicitly schema-qualify this table name. When you use [“Schema Name” on page 66](#), this table name is implicitly schema-qualified with that schema name. If this table is located in a schema other than the implicit schema, it must be schema-qualified.

The following table lists the properties of this parameter:

Table 6-51 *Event Log Table Name: Properties*

Property	Value
Tag Name	log-table
Required?	no ¹
Case-Sensitive?	See “Undelimited Identifier Case Sensitivity” on page 181 .
Sample Value	eventlog
Default Value	(none)
Schema-Dependent	True

¹ This parameter is required if [“Publication Mode” on page 91](#) is set to 1 (triggered publication).

Delete Processed Rows?

The Delete Processed Rows? parameter specifies whether processed rows are deleted from the event log table.

When this parameter is set to a Boolean True, processed rows are deleted. When this parameter is set to Boolean False, processed row’s `status` field values are updated.

To mitigate the performance hit caused when processed rows remain in the event log table, we recommend periodically moving the rows into a history table. Do one of the following:

- ♦ Call a clean-up stored procedure via the parameter [“Post Polling Statements” on page 97](#).

- ◆ Place a before-delete trigger on the event log table to intercept delete events executed against the event log table and to move deleted rows to a history table before they are deleted from the event log table.

The following table lists the properties of this parameter:

Table 6-52 *Delete Processed Rows?: Properties*

Property	Value
Tag Name	delete-from-log
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

NOTE: Setting this parameter to Boolean False degrades publication performance unless processed rows are periodically removed from the event log table.

Allow Loopback?

The Allow Loopback? parameter specifies whether events caused by the driver's database user account should be published.

When this parameter is set to Boolean True, loopback events are published. When this parameter is set to Boolean False, loopback events are ignored.

The following table lists the properties of this parameter:

Table 6-53 *Allow Loopback?: Properties*

Property	Value
Tag Name	allow-loopback
Required?	no
Default Value	0 (no)
Legal Values	1, yes, true (yes) 0, no, false (no)
Schema-Dependent	True

NOTE: Setting this parameter to Boolean True might degrade performance because extraneous events might be published.

6.5.3 Triggerless Publication Parameters

The Startup Option parameter specifies what happens when a triggerless publisher starts.

Startup Option

Table 6-54 *Startup Option: Settings and Results*

Setting	Result
1	All objects are assumed to have changed and are republished.
2	Past and present changes are ignored.
3	All past and present changes are published.

The following table lists the properties of this parameter:

Table 6-55 *Startup Option: Properties*

Property	Value
Tag Name	startup-option
Required?	no
Default Value	1 (process all changes)
Legal Values	1 (resync all objects) 2 (process future changes only) 3 (process all changes)
Schema-Dependent	True

IMPORTANT: The following configuration changes can force a full resynchronization:

- ♦ Changing anything in the **Authentication Context** parameter other than URL properties forces a resynchronization of all objects when triggerless publication is used.
- ♦ Changing the value of the **Schema Name** parameter or the **Table/View Names** parameter forces a resynchronization of all objects when triggerless publication is used.
- ♦ Changing the **State Directory** parameter value.
- ♦ Moving or deleting state files. See “**Changes That Can Force Triggerless Publisher Resynchronization**” on page 63.
- ♦ Changing table/view structure in the database (in particular, changing the position or type of key columns).

6.5.4 Polling Parameters

- ♦ “**Polling Interval (In Seconds)**” on page 96
- ♦ “**Publication Time of Day**” on page 96

- ◆ [“Post Polling Statements” on page 97](#)
- ◆ [“Batch Size” on page 97](#)
- ◆ [“Heartbeat Interval \(In Minutes\)” on page 98](#)

Polling Interval (In Seconds)

The Polling Interval (In Seconds) parameter specifies how many seconds of inactivity elapse between polling cycles.

The following table lists the properties of this parameter:

Table 6-56 *Polling Interval (In Seconds): Properties*

Property	Value
Tag Name	polling-interval
Required?	no
Default Value	10 (seconds)
Legal Values	1-604800 (1 week)
Schema-Dependent	True

NOTE: We recommend that you set this value to no less than 10 seconds.

Publication Time of Day

The Publication Time of Day parameter specifies at what time, each day, publication begins. Time is understood to mean server local time (the time on the server where the driver is running).

The following table lists the properties of this parameter:

Table 6-57 *Publication Time of Day: Properties*

Property	Value
Tag Name	time-of-day
Required?	no
Sample Value	13:00:00 (1PM)
Default Value	(none)
Legal Values	hh:mm:ss (h = hour, m = minute, s = second)
Schema-Dependent	True

NOTE: This parameter overrides the parameter Polling Interval (In Seconds). See [“Polling Interval \(In Seconds\)” on page 96](#).

Post Polling Statements

The Post Polling Statements parameter specifies the SQL statements that are executed at the end of each active polling cycle. An active polling cycle is one where some publication activity has occurred.

The primary purpose of this parameter is to allow cleanup of the event log table following publication activity.

You probably need to explicitly schema-qualify any database objects (for example, tables, stored procedures, and functions) referenced in these statements.

The following table lists the properties of this parameter:

Table 6-58 *Post Polling Statements: Properties*

Property	Value
Tag Name	post-poll-stmt
Required?	no
Case-Sensitive?	See “Undelimited Identifier Case Sensitivity” on page 181 .
Delimiters	semicolon
Sample Value	DELETE FROM direct.direct_process
Default Value	(none)
Legal Values	(any set of legal SQL statements)
Schema-Dependent	True

Batch Size

The Batch Size parameter specifies how many events are sent in a single publication document.

Basically, the larger the batch, the better the performance.

- ◆ Larger batches necessitate fewer trips across the network in both directions.
- ◆ More events in a single document require fewer trips from the Publisher channel to the Identity Manager engine (assuming that query-back events are not being used).
- ◆ Larger batches minimize the number of trips from the Publisher channel to the database (assuming that the third-party JDBC driver and database support batch processing).
- ◆ Larger batches require fewer commits to state files in the local file system.

Commits can also be costly.

This parameter defines an upper bound. The Publisher channel might override the specified value under certain conditions. The upper bound of 128 was chosen to minimize the likelihood of overflowing the Java heap and to mitigate delaying termination of the Publisher thread on driver shutdown.

The following table lists the properties of this parameter:

Table 6-59 *Batch Size: Properties*

Property	Value
Tag Name	batch-size
Required?	no
Default Value	1
Legal Values	1 to 128
Schema-Dependent	True

Heartbeat Interval (In Minutes)

The Heartbeat Interval (In Minutes) parameter specifies how many minutes the Publisher channel can be inactive before it sends a heartbeat document. In practice, more than the number of minutes specified can elapse. That is, this parameter defines a lower bound. The Publisher channel sends a heartbeat document only if the Publisher channel has been inactive for the specified number of minutes. Any publication document sent is, in effect, a heartbeat document.

The following table lists the properties of this parameter:

Table 6-60 *Heartbeat Interval (In Minutes): Properties*

Property	Value
Tag Name	pub-heartbeat-interval
Required?	no
Default Value	0
Legal Values	0 to 2,147,483,647 (java.lang.Integer.MAX_VALUE)
Schema-Dependent	False

6.6 Trace Levels

To see debugging output from the driver, add a DirXML-DriverTraceLevel attribute value from 1 to 7 on the driver set containing the driver instance. This attribute is commonly confused with the DirXML-XSL TraceLevel attribute. For more information on driver set trace levels, refer to “[Viewing Identity Manager Processes](#)” in the *Novell Identity Manager 3.5.1 Administration Guide*.

The driver supports the following seven trace levels:

Table 6-61 *Supported Trace Levels*

Level	Description
1	Minimal tracing
2	Database properties

Level	Description
3	Connection status, SQL statements, event log records
4	Verbose output
5	Database resource allocation/deallocation; state file contents
6	JDBC API (invoked methods, passed arguments, returned values, etc.)
7	Third-party driver

Levels 6 and 7 are particularly useful for debugging third-party drivers.

6.7 Configuring Third-Party JDBC Drivers

The following guidelines help you configure third-party drivers. For specific configuration instructions, refer to your third-party driver's documentation.

- ◆ Use the latest version of the driver.
- ◆ Third-party driver behavior might be configurable.
In many cases, incompatibility issues can be resolved by adjusting the driver's JDBC URL properties.
- ◆ When you work with international characters, you often must explicitly specify to third-party drivers the character encoding that the database uses.

Do this by appending a property string to the end of the driver's JDBC URL.

Properties usually consist of a property keyword and character encoding value (for example, `jdbc:odbc:mssql;charSet=Big5`). The property keyword might vary among third-party drivers.

The possible character encoding values are defined by Sun. For more information, refer to [Sun's Supported Encoding Web site \(http://java.sun.com/j2se/1.5.0/docs/guide/intl/encoding.doc.html\)](http://java.sun.com/j2se/1.5.0/docs/guide/intl/encoding.doc.html).

The following table lists the recommended settings for maximum driver compatibility. These settings are useful when you use an unsupported third-party driver during initial configuration.

Table 6-62 Recommended Settings for Third-Party JDBC Drivers

Parameter Name	Compatibility Value
Synchronization filter	empty
Reuse statements?	0 (no)
Use manual transactions?	0 (no)
Use minimal number of connections?	yes
Retrieve minimal metadata?	1 (yes)
Number of returned result sets	one

Activating the JDBC Driver

7

Activate the driver within 90 days of installation. Otherwise, the driver won't work.

For information on activation, see “[Activating Novell Identity Manager Products](#)” in the *Identity Manager 3.5.1 Installation Guide*.

- ♦ [Section 8.1, “Starting, Stopping, or Restarting the JDBC Driver,” on page 103](#)
- ♦ [Section 8.2, “Migrating and Resynchronizing Data,” on page 103](#)
- ♦ [Section 8.3, “Using the DirXML Command Line Utility,” on page 104](#)
- ♦ [Section 8.4, “Viewing Driver Versioning Information,” on page 104](#)
- ♦ [Section 8.5, “Reassociating a Driver Set Object with a Server Object,” on page 109](#)
- ♦ [Section 8.6, “Changing the Driver Configuration,” on page 110](#)
- ♦ [Section 8.7, “Storing Driver Passwords Securely with Named Passwords,” on page 110](#)
- ♦ [Section 8.8, “Adding a Driver Heartbeat,” on page 117](#)

8.1 Starting, Stopping, or Restarting the JDBC Driver

To start the JDBC driver in Designer for Identity Manager:

- 1 Open a project in the Modeler, then right-click the driver line.
- 2 Select *Live > Start Driver, Stop Driver, or Restart Driver*.

To start the JDBC driver in iManager:

- 1 If you changed default data locations during configuration, ensure that the new locations exist before you start the driver.
- 2 Click *Identity Manager > Identity Manager Overview*.
- 3 Browse to the driver set where the driver exists, then click *Search*.
- 4 Click the driver status indicator in the upper right corner of the driver icon, then click *Start driver, Stop driver, or Restart driver*.

If a change log is available, the driver processes all the changes in the change log. To force an initial synchronization, see [“Migrating and Resynchronizing Data” on page 103](#).

8.2 Migrating and Resynchronizing Data

Identity Manager synchronizes data as it changes. If you want to synchronize all data immediately, you can choose from the following options:

- ♦ **Migrate Data from the Identity Vault:** Allows you to select containers or objects you want to migrate from an Identity Vault to an JDBC server. When you migrate an object, the Metadirectory engine applies all of the Matching, Placement, and Create policies, as well as the Subscriber filter, to the object.
- ♦ **Migrate Data into the Identity Vault:** Allows you to define the criteria that Identity Manager uses to migrate objects from an JDBC server into an Identity Vault. When you migrate an object, the Metadirectory engine applies all of the Matching, Placement, and Create policies, as well as the Publisher filter, to the object. Objects are migrated into the Identity Vault by using the order you specify in the Class list.

- ♦ **Synchronize:** Identity Manager looks in the Subscriber class filter and processes all objects for those classes. Associated objects are merged. Unassociated objects are processed as Add events.

To use one of the options:

- 1 In iManager, click *Identity Manager > Identity Manager Overview*.
- 2 Browse to and select the driver set where the driver exists, then click *Search*.
- 3 Click the driver icon.
- 4 Click the appropriate migration button.

For more information, see [Chapter 9, “Synchronizing Objects,” on page 119](#).

8.3 Using the DirXML Command Line Utility

The DirXML Command Line utility provides command line access to manage the driver. This utility is not a replacement for iManager or Designer. The primary use of this utility is to allow you to create platform-specific scripts to manage the driver.

For example, you could create a shell script on Linux to check the status of the driver. See [Appendix M, “The DirXML Command Line Utility,” on page 257](#) for information about the DirXML Command Line utility. For daily tasks, use iManager or Designer.



8.4 Viewing Driver Versioning Information

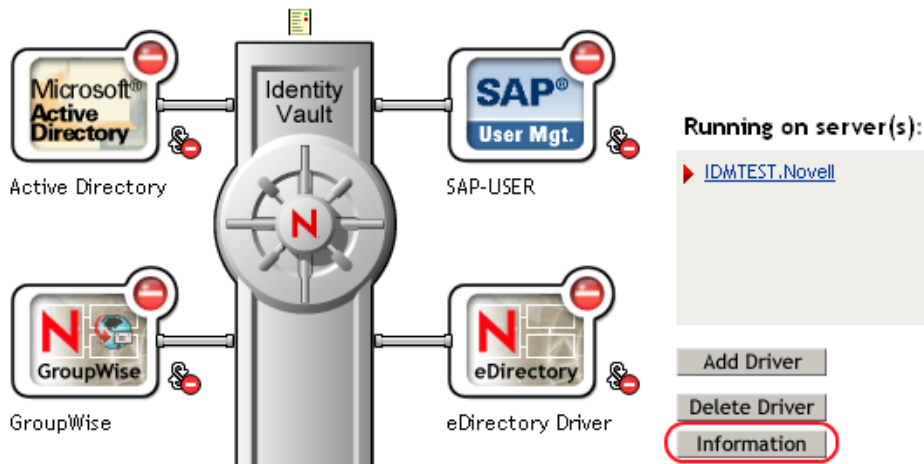
The Versioning Discovery tool only exists in iManager.

- ♦ [Section 8.4.1, “Viewing a Hierarchical Display of Versioning Information,” on page 104](#)
- ♦ [Section 8.4.2, “Viewing the Versioning Information As a Text File,” on page 106](#)
- ♦ [Section 8.4.3, “Saving Versioning Information,” on page 108](#)

8.4.1 Viewing a Hierarchical Display of Versioning Information

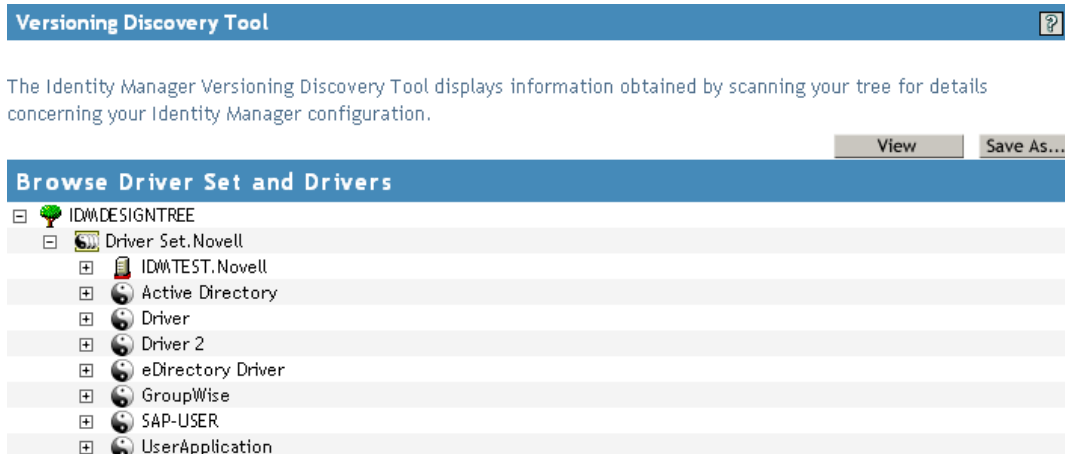
- 1 To find your Driver Set object in iManager, click *Identity Manager > Identity Manager Overview*, then click *Search*.
- 2 In the Identity Manager Overview, click *Information*.

Driver Set: Driver Set.Novell  Activation required by: November 1, 2006 



You can also select *Identity Manager Utilities > Versions Discovery*, browse to and select the Driver Set object, then click *OK*.

- 3 View a top-level or unexpanded display of versioning information.




The unexpanded hierarchical view displays the following:

- ◆ The eDirectory™ tree that you are authenticated to
- ◆ The Driver Set object that you selected
- ◆ Servers that are associated with the Driver Set object

If the Driver Set object is associated with two or more servers, you can view Identity Manager information on each server.

- ◆ Drivers

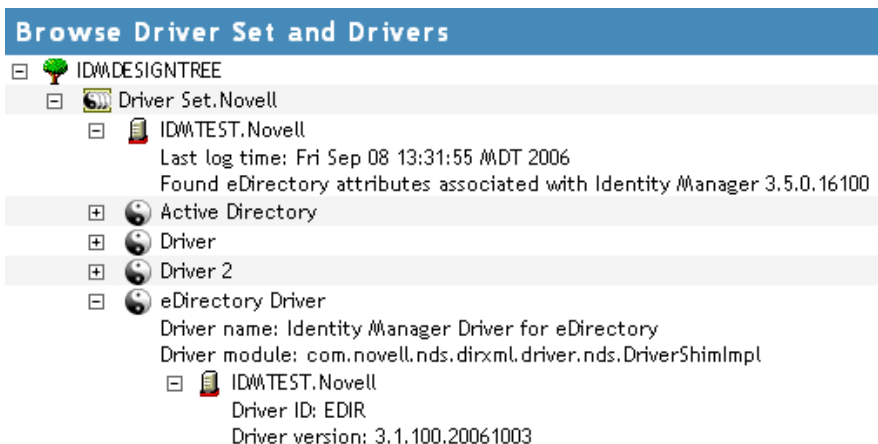
- 4 View versioning information related to servers by expanding the server icon.



The expanded view of a top-level server icon displays the following:

- ◆ Last log time
- ◆ Version of Identity Manager that is running on the server

5 View versioning information related to drivers by expanding the driver icon.



The expanded view of a top-level driver icon displays the following:

- ◆ The driver name
- ◆ The driver module (for example, `com.novell.nds.dirxml.driver.delimitedtext.DelimitedTextDriver`)

The expanded view of a server under a driver icon displays the following:

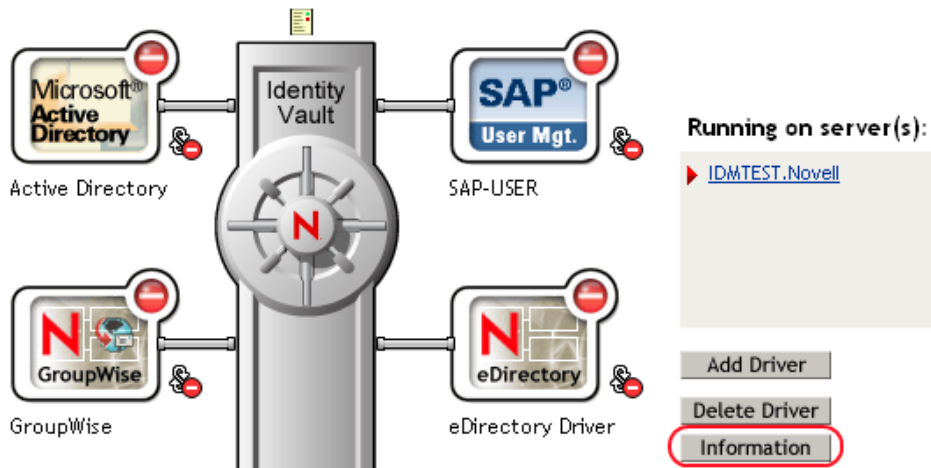
- ◆ The driver ID
- ◆ The version of the instance of the driver running on that server

8.4.2 Viewing the Versioning Information As a Text File

Identity Manager publishes versioning information to a file. You can view this information in text format. The textual representation is the same information contained in the hierarchical view.

- 1** To find your Driver Set object in iManager, click *Identity Manager > Identity Manager Overview*, then click *Search*.
- 2** In the Identity Manager Overview, click *Information*.

Driver Set: Driver Set.Novell  Activation required by: November 1, 2006 



You can also select *Identity Manager Utilities > Versioning Discovery*, browse to and select the Driver Set object, then click *Information*.

Versioning Discovery Tool

The Identity Manager Versioning Discovery Tool displays information obtained by scanning your tree for details concerning your Identity Manager configuration.

View

Save As...

Browse Driver Set and Drivers

IDMDESIGNTREE

- Driver Set.Novell
 - IDMTEST.Novell
 - Active Directory
 - Driver
 - Driver 2
 - eDirectory Driver
 - GroupWise
 - SAP-USER
 - UserApplication

3 In the Versioning Discovery Tool dialog box, click *View*.

The information is displayed as a text file in the Report Viewer window.

Versioning Discovery Tool - Report Viewer

```
Identity Manager Version Discovery Tool v2.0
Novell, Inc. Copyright 2003, 2004

Version Query started Saturday, January 20, 2007 11:02:52 AM MST

Parameter Summary:
  Default server's DN:  IDMTEST.Novell
  Default server's IP address:  137.65.151.208
  Logged in as admin, context Novell
  Tree name:  IDMDSIGNTREE
  Found 7 Identity Manager Drivers


Driver Set:  Driver Set.Novell
  Driver Set running on Identity Vault:  IDMTEST.Novell
  Last log time:  Fri Sep 08 13:31:55 MDT 2006
  Found eDirectory attributes associated with Identity Manager 3.5.0.1
Driver:  Active Directory.Driver Set.Novell
  Driver name:  Identity Manager Driver for Active Directory and Excha
  Driver module:  addriver.dll
  Driver Set running on Identity Vault:  IDMTEST.Novell
  Didn't find any DirXML-DriverVersion attributes associated w:
  This may mean the Metadirectory engine is older than
  It does not indicate anything about the version of t
Driver:  Driver.Driver Set.Novell
  Driver name:  Identity Manager Driver for Peoplesoft
  Driver module:  NPSShim.dll
  Driver Set running on Identity Vault:  IDMTEST.Novell
```

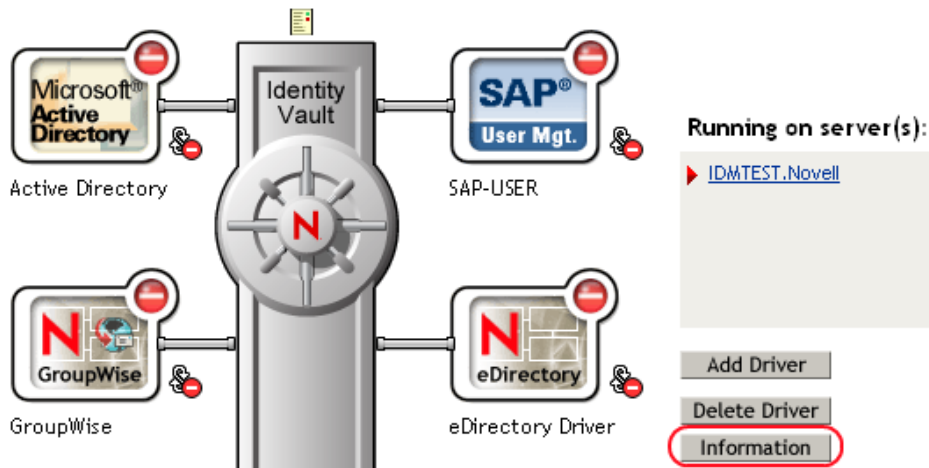
OK

8.4.3 Saving Versioning Information

You can save versioning information to a text file on your local or network drive.

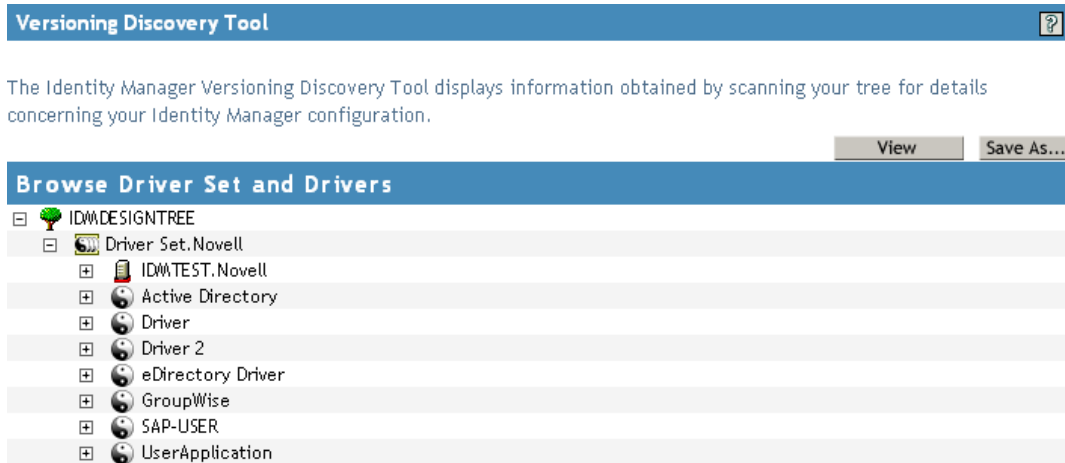
- 1 To find the Driver Set object in iManager, click *Identity Manager > Identity Manager Overview*, then click *Search*.
- 2 In the Identity Manager Overview, click *Information*.

Driver Set: Driver Set.Novell  [Activation required by: November 1, 2006](#) 



You can also select *Identity Manager Utilities > Versions Discovery*, browse to and select the Driver Set object, then click *Information*.

- 3 In the Versioning Discovery Tool dialog box, click *Save As*.



- 4 In the File Download dialog box, click *Save*.
- 5 Navigate to the desired directory, type a filename, then click *Save*.
Identity Manager saves the data to a text file.

8.5 Reassociating a Driver Set Object with a Server Object

The driver set object should always be associated with a server object. If the driver set is not associated with a server object, none of the drivers in the driver set can start.

If the link between the driver set object and the server object becomes invalid, you see one of the following conditions:

- ◆ When upgrading eDirectory your Identity Manager server, you get the error UniqueSPIException error -783.
- ◆ No server is listed next to the driver set in the Identity Manager Overview window.
- ◆ A server is listed next to the driver set in the Identity Manager Overview window, but the name is garbled text.

To resolve this issue, disassociate the driver set object and the server object, then reassociate them.

- 1 In iManager click *Identity Manager > Identity Manager Overview*, then click *Search* to find the Driver Set object that the driver should be associated with.
- 2 Click the *Remove server* icon, then click *OK*.
- 3 Click the *Add server* icon, then browse to and select the Server object.
- 4 Click *OK*.

8.6 Changing the Driver Configuration

If you need to change the driver configuration, Identity Manager allows you to make the change through Designer or iManager.

To change the driver configuration in Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties*.

To change the driver configuration in iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties*.

For a listing of all of the configuration fields, see [Appendix N, “Properties of the JDBC Driver,” on page 271](#).

8.7 Storing Driver Passwords Securely with Named Passwords

Identity Manager allows you to store multiple passwords securely for a particular driver. This functionality is referred to as Named Passwords. Each different password is accessed by a key, or name.

You can also use the Named Passwords feature to store other pieces of information securely, such as a user name.

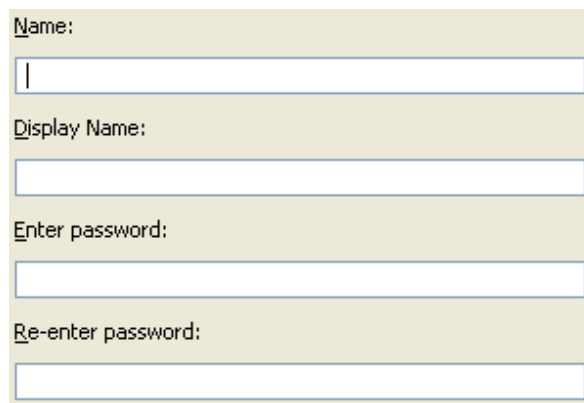
To use a named password in a driver policy, you refer to it by the name of the password, instead of using the actual password, and the Metadirectory engine sends the password to the driver. The

method described in this section for storing and retrieving named passwords can be used with any driver without making changes to the driver shim.

- ♦ [Section 8.7.1, “Using Designer to Configure Named Passwords,” on page 111](#)
- ♦ [Section 8.7.2, “Using iManager to Configure Named Passwords,” on page 111](#)
- ♦ [Section 8.7.3, “Using Named Passwords in Driver Policies,” on page 113](#)
- ♦ [Section 8.7.4, “Configuring Named Passwords by Using the DirXML Command Line Utility,” on page 114](#)

8.7.1 Using Designer to Configure Named Passwords

- 1 Right-click the Driver object, then select *Properties*.
- 2 Select *Named Password*, then click *New*.



The screenshot shows a dialog box for configuring a named password. It has a light beige background and contains four input fields. The first field is labeled "Name:" and is empty. The second field is labeled "Display Name:" and is empty. The third field is labeled "Enter password:" and is empty. The fourth field is labeled "Re-enter password:" and is empty.

- 3 Specify the *Name* of the named password.
- 4 Specify the *Display name* of the named password.
- 5 Specify the named password, then re-enter the password.
- 6 Click *OK* twice.

8.7.2 Using iManager to Configure Named Passwords

- 1 Click *Identity Manager > Identity Manager Overview*.
- 2 Click *Search* to search for the driver set that is associated with the driver.
- 3 In the Identity Manager Overview, click the upper right corner of the driver icon, then click *Edit properties*.
- 4 On the Identity Manager tab, click *Named Passwords*.

The Named Passwords page appears, listing the current named passwords for this driver. If you have not set up any named passwords, the list is empty.



- 5 To add a named password, click *Add*, complete the fields, then click *OK*.

Named Password

Named Passwords lets you securely store multiple passwords for a driver. Instead of including a password in clear text in a driver policy, you can configure the policy to request a Named Password.

Name:

Display name:

Enter password:

Reenter password:

OK *Cancel*

- 6 Specify a name, display name, and a password, then click *OK* twice.
 You can use this feature to store other kinds of information securely, such as a username.
- 7 Click *OK* to restart the driver and have the changes take effect.

To remove a Named Password, select the password name, then click *Remove*. The password is removed without prompting you to confirm the action.

8.7.3 Using Named Passwords in Driver Policies

- ♦ “Making a Call to a Named Password” on page 113
- ♦ “Referencing a Named Password” on page 113


Making a Call to a Named Password

Policy Builder allows you to make a call to a named password. Create a new rule and select Named Password as the condition, then set an action, depending upon if the Named Password is available or not available.

- 1 In Designer, launch Policy Builder, right-click, then click *New > Rule*.
- 2 Specify the name of the rule, then click *Next*.
- 3 Select the condition structure, then click *Next*.
- 4 Select *named password* for the *Condition*.
- 5 Browse to and select the named password that is stored on the driver.
In this example, the named password is *userinfo*.
- 6 Select whether the Operator is available or not available.
- 7 Select an action for the *Do* field.
In this example, the action is *veto*.

The example indicates that if the *userinfo* named password is not available, then the event is vetoed.

Figure 8-1 A Policy Using Named Passwords



Referencing a Named Password

The following example shows how a named password can be referenced in a driver policy on the Subscriber channel in XSLT:

```
<xsl:value-of  
select="query:getNamedPassword($srcQueryProcessor, 'mynamedpassword') "  
xmlns:query="http://www.novell.com/java/  
com.novell.nds.dirxml.driver.XdsQueryProcessor/>
```

8.7.4 Configuring Named Passwords by Using the DirXML Command Line Utility

- ♦ [“Creating a Named Password in the DirXML Command Line Utility” on page 114](#)
- ♦ [“Removing a Named Password by Using the DirXML Command Line Utility” on page 115](#)

Creating a Named Password in the DirXML Command Line Utility

- 1 Run the DirXML Command Line utility.

For information, see [Appendix M, “The DirXML Command Line Utility,” on page 257](#).

- 2 Enter your username and password.

The following list of options appears.

```
DirXML commands
 1: Start driver
 2: Stop driver
 3: Driver operations...
 4: Driver set operations...
 5: Log events operations...
 6: Get DirXML version
 7: Job operations...
99: Quit
Enter choice:
```

- 3 Enter 3 for driver operations.

A numbered list of drivers appears.

- 4 Enter the number for the driver you want to add a named password to.

The following list of options appears.

```
Select a driver operation for:
driver_name
 1: Start driver
 2: Stop driver
 3: Get driver state
 4: Get driver start option
 5: Set driver start option
 6: Resync driver
 7: Migrate from application into DirXML
 8: Submit XDS command document to driver
 9: Submit XDS event document to driver
10: Queue event for driver
11: Check object password
12: Initialize new driver object
13: Passwords operations
14: Cache operations
99: Exit
Enter choice:
```

- 5 Enter 13 for password operations.

The following list of options appears.

```
Select a password operation
```

```
1: Set shim password
2: Reset shim password
3: Set Remote Loader password
4: Clear Remote Loader password
5: Set named password
6: Clear named password(s)
7: List named passwords
8: Get passwords state
99: Exit
Enter choice:
```

6 Enter 5 to set a new named password.

The following prompt appears:

```
Enter password name:
```

7 Enter the name by which you want to refer to the named password.

8 Enter the actual password that you want to secure at the following prompt:

```
Enter password:
```

The characters you type for the password are not displayed.

9 Confirm the password by entering it again at the following prompt:

```
Confirm password:
```

10 After you enter and confirm the password, you are returned to the password operations menu.

11 After completing this procedure, you can use the 99 option twice to exit the menu and quit the DirXML Command Line Utility.

Removing a Named Password by Using the DirXML Command Line Utility

This option is useful if you no longer need named passwords that you previously created.

1 Run the DirXML Command Line utility.

For information, see [Appendix M, “The DirXML Command Line Utility,” on page 257](#).

2 Enter your username and password.

The following list of options appears.

```
DirXML commands
1: Start driver
2: Stop driver
3: Driver operations...
4: Driver set operations...
5: Log events operations...
6: Get DirXML version
7: Job operations
99: Quit
Enter choice:
```

3 Enter 3 for driver operations.

A numbered list of drivers appears.

4 Enter the number for the driver you want to remove named passwords from.

The following list of options appears.

Select a driver operation for:
driver_name
1: Start driver
2: Stop driver
3: Get driver state
4: Get driver start option
5: Set driver start option
6: Resync driver
7: Migrate from application into DirXML
8: Submit XDS command document to driver
9: Submit XDS event document to driver
10: Queue event for driver
11: Check object password
12: Initialize new driver object
13: Passwords operations
14: Cache operations
99: Exit
Enter choice:

5 Enter 13 for password operations.

The following list of options appears.

Select a password operation
1: Set shim password
2: Reset shim password
3: Set Remote Loader password
4: Clear Remote Loader password
5: Set named password
6: Clear named password(s)
7: List named passwords
8: Get passwords state
99: Exit
Enter choice:

6 (Optional) Enter 7 to see the list of existing named passwords.

The list of existing named passwords is displayed.

This step can help you make sure you are removing the correct password.

7 Enter 6 to remove one or more named passwords.

8 Enter No to remove a single named password at the following prompt:

Do you want to clear all named passwords? (yes/no):

9 Enter the name of the named password you want to remove at the following prompt:

Enter password name:

After you enter the name of the named password you want to remove, you are returned to the password operations menu:

Select a password operation
1: Set shim password
2: Reset shim password
3: Set Remote Loader password
4: Clear Remote Loader password
5: Set named password

```
6: Clear named password(s)
7: List named passwords
8: Get passwords state
99: Exit
Enter choice:
```

10 (Optional) Enter 7 to see the list of existing named passwords.

This step lets you verify that you have removed the correct password.

11 After completing this procedure, you can use the 99 option twice to exit the menu and quit the DirXML Command Line utility.

8.8 Adding a Driver Heartbeat

The driver heartbeat is a feature of the Identity Manager drivers that ship with Identity Manager 2 and later. Its use is optional. The driver heartbeat is configured by using a driver parameter with a time interval specified. If a heartbeat parameter exists and has an interval value other than 0, the driver sends a heartbeat document to the Metadirectory engine if no communication occurs on the Publisher channel for the specified interval of time.

The intent of the driver heartbeat is to give you a trigger to allow you to initiate an action at regular intervals, if the driver does not communicate on the Publisher channel as often as you want the action to occur. To take advantage of the heartbeat, you must customize your driver configuration or other tools. The Metadirectory engine accepts the heartbeat document but does not take any action because of it.

For most drivers, a driver parameter for heartbeat is not used in the sample configurations, but you can add it.

A custom driver that is not provided with Identity Manager can also provide a heartbeat document, if the driver developer has written the driver to support it.

To configure the heartbeat:

- 1** In iManager, click *Identity Manager > Identity Manager Overview*.
- 2** Browse to and select your driver set object, then click *Search*.
- 3** In the Identity Manager Overview, click the upper right corner of the driver icon, then click *Edit properties*.
- 4** On the Identity Manager tab, click *Driver Configuration*, scroll to *Publisher Settings*, then locate *Heartbeat interval (in minutes)*.

If a driver parameter already exists for heartbeat, you can change the interval and save the changes. Configuration is then complete.

The value of the interval cannot be less than 1. A value of 0 means that the feature is turned off.

- 5** (Conditional) If a driver parameter does not exist for heartbeat, click *Edit XML*.
- 6** (Conditional) Add a driver parameter entry similar to the following example, as a child of `<publisher-options>`.

```
<pub-heartbeat-interval display-name="Heart Beat">10</pub-heartbeat-interval>
```

TIP: If the driver does not produce a heartbeat document after being restarted, check the placement of the driver parameter in the XML.

7 Save the changes, then make sure the driver is stopped and restarted.

After you add the driver parameter, you can edit the time interval by using the graphical view. Another option is to create a reference to a global configuration value (GCV) for the time interval. Like other global configuration values, the driver heartbeat can be set at the driver set level instead of on each individual Driver object. If a driver does not have a particular GCV, and the Driver Set object does have it, the driver inherits the value from the Driver Set object.

Synchronizing Objects

9

This section explains driver and object synchronization in DirXML[®] 1.1a, Identity Manager 2.0, and Identity Manager 3.x. Driver synchronization was not available for DirXML 1.0 and DirXML 1.1.

After the driver is created, instead of waiting for objects to be modified or created, the data between the two connected systems can be sent through the synchronization process.

- ♦ [Section 9.1, “What Is Synchronization?” on page 119](#)
- ♦ [Section 9.2, “When Does Synchronization Occur?” on page 119](#)
- ♦ [Section 9.3, “How Does the Metadirectory Engine Decide Which Object to Synchronize?” on page 120](#)
- ♦ [Section 9.4, “How Synchronization Works,” on page 121](#)

9.1 What Is Synchronization?

The actions commonly referred to as “synchronization” in Identity Manager refer to several different but related actions:

- ♦ Synchronization (or merging) of attribute values of an object in the Identity Vault with the corresponding attribute values of an associated object in a connected system.
- ♦ Migration of all Identity Vault objects and classes that are included in the filter on the Subscriber channel.
- ♦ Generation of the list of objects to submit to the driver’s Subscriber channel for synchronization or migration in response to a user request (a manual synchronization).
- ♦ Generation of the list of objects to submit to the driver’s Subscriber channel for synchronization or migration in response to enabling a formerly disabled driver, or in response to a cache error.

9.2 When Does Synchronization Occur?

The Metadirectory engine synchronizes objects or merges them in the following circumstances:

- ♦ A `<sync>` event element is submitted on the Subscriber or Publisher channel.
- ♦ A `<sync>` event element is submitted on the Subscriber channel in the following circumstances:
 - ♦ The state of the object’s association value is set to “manual” or “migrate.” (This causes an eDirectory™ event, which in turn causes the Identity Manager caching system to queue an object synchronization command in the affected driver’s cache.)
 - ♦ An object synchronization command is read from the driver’s cache.
- ♦ A `<sync>` event element is submitted on the Publisher channel in the following circumstances:
 - ♦ A driver submits a `<sync>` event element. No known driver currently does this.

- ♦ The Metadirectory engine submits a <sync> event element for each object found as the result of a migrate-into-NDS query. The engine submits these <sync> events by using the Subscriber thread, but processes them by using the Publisher channel filter and policies.
- ♦ An <add> event (real or synthetic) is submitted on a channel, and the channel Matching policy finds a matching object in the target system.
- ♦ An <add> event with an association is submitted on the Subscriber channel. This normally occurs only in exceptional cases, such as the bulk load of objects into eDirectory with DirXML-Associations attribute values.
- ♦ An <add> event is submitted on the Publisher channel, and an object is found in eDirectory that already has the association value reported with the <add> event.

The Metadirectory engine generates synchronization requests for zero or more objects in the following cases:

- ♦ The user issues a manual driver synchronization request. This corresponds to the *Resync* button in the Driver Set property page in ConsoleOne®, or to the *Synchronize* button on the iManager Identity Manager Driver Overview page.
- ♦ The Metadirectory engine encounters an error with the driver's cache and cannot recover from the cache error. The driver's cache is deleted, and the engine generates object synchronization commands as detailed in [Section 9.3, "How Does the Metadirectory Engine Decide Which Object to Synchronize?"](#) on page 120.

9.3 How Does the Metadirectory Engine Decide Which Object to Synchronize?

The Metadirectory engine processes both manually initiated and automatically initiated synchronization requests in the same manner. The only difference in the processing of manually initiated versus automatically initiated driver synchronization requests is the starting filter time used to filter objects being considered for synchronization.

The starting filter time is used to filter objects that have modification or creation times that are older than the starting time specified in the synchronization request.

For automatically initiated driver synchronization, the starting filter time is obtained from the time stamps of cached eDirectory events. In particular, the starting filter time is the earliest time for the cached events that haven't yet been successfully processed by the driver's Subscriber channel.

For manually initiated driver synchronization, the default starting filter time is the earliest time in the eDirectory database. In Identity Manager 2 and Identity Manager 3, an explicit starting filter time can also be set. DirXML 1.1a has no facility to set the starting filter time value for synchronization when manually initiating driver synchronization.

The Metadirectory engine creates a list of objects to be synchronized on the Subscriber channel in the following manner:

1. It finds all objects that:
 - ♦ Have an entry modification time stamp greater than or equal to the starting filter time and
 - ♦ Exist in the filter on the Subscriber channel.

2. It finds all objects that have an entry creation time stamp greater than or equal to the starting filter time.
3. It adds a `synchronize object` command to the following:
 - ♦ The driver cache for each unique object found that has an entry modification time stamp greater than or equal to the starting filter time
 - ♦ All objects and classes that are in the Subscriber filter channel in the driver being synchronized

9.4 How Synchronization Works

After the Metadirectory engine determines that an object is to be synchronized, the following processes occur:

1. Each system (the Identity Vault and the connected system) is queried for all attribute values in the appropriate filters.
 - ♦ eDirectory is queried for all values in the Subscriber filter, and for values that are marked for synchronization in Identity Manager 2.x and Identity Manager 3.x.
 - ♦ The connected system is queried for all values in the Publisher filter, and for values that are marked for synchronization in Identity Manager 2.x and Identity Manager 3.x.
2. The returned attribute values are compared, and modification lists are prepared for the Identity Vault and the connected system according to [Table 9-1 on page 122](#), [Table 9-2 on page 124](#), and [Table 9-3 on page 125](#).

In the tables the following pseudo-equations are used:

- ♦ “Left = Right” indicates that the left side receives all values from the right side.
- ♦ “Left = Right[1]” indicates that the left side receives one value from the right side. If there is more than one value, it is indeterminate.
- ♦ “Left += Right” indicates that the left side adds the right side values to the left side’s existing values.
- ♦ “Left = Left + Right” indicates that the left side receives the union of the values of the left and right sides.

Identity Manager has three different combinations of selected items in the filter, and each one creates a different output.

- ♦ [Section 9.4.1, “Scenario One,” on page 121](#)
- ♦ [Section 9.4.2, “Scenario Two,” on page 123](#)
- ♦ [Section 9.4.3, “Scenario Three,” on page 124](#)

9.4.1 Scenario One

The attribute is set to *Synchronize* on the Publisher and Subscriber channels, and the merge authority is set to *Default*.

Figure 9-1 Scenario One

Class Name: User
Attribute Name: Facsimile Telephone Num

Publish

Synchronize
 Ignore
 Notify
 Reset

Subscribe

Synchronize
 Ignore
 Notify
 Reset

Merge Authority

Default
 Identity Vault
 Application
 None

Optimize modifications to Identity Vault

Yes
 No

The following table contains the values that the Metadirectory engine synchronizes when the attribute is sent through a filter that is set to the configuration for Scenario One. The table shows different outputs depending upon the following:

- ◆ Whether the attribute comes from the Identity Vault or the Application
- ◆ If the attribute is single-valued or multi-valued, and if the attribute is empty or non-empty.
- ◆ If the attribute is empty or non-empty

Table 9-1 Output of Scenario One

	Identity Vault single-valued empty	Identity Vault single-valued non-empty	Identity Vault multi-valued empty	Identity Vault multi-valued non-empty
Application single-valued empty	No change	App = Identity Vault	No change	App = Identity Vault[1]
Application single-valued non-empty	Identity Vault = App	App = Identity Vault	Identity Vault = App	Identity Vault + = App
Application multi-valued empty	No change	App = Identity Vault	No change	App = Identity Vault

	Identity Vault single-valued empty	Identity Vault single-valued non-empty	Identity Vault multi-valued empty	Identity Vault multi-valued non-empty
Application multi-valued non-empty	Identity Vault = App[1]	App + = Identity Vault	Identity Vault = App	App = App + Identity Vault Identity Vault = App + Identity Vault

9.4.2 Scenario Two

The attribute is set to *Synchronize* only on the Subscriber channel, or it is set to *Synchronize* on both the Subscriber and Publisher channels. The merge authority is set to *Identity Vault*.

Figure 9-2 Scenario Two

Class Name: User

Attribute Name: Description

Publish

Synchronize

Ignore

Notify

Reset

Subscribe

Synchronize

Ignore

Notify

Reset

Merge Authority

Default

Identity Vault

Application

None

Optimize modifications to Identity Vault

Yes

No

The following table contains the values that the Metadirectory engine synchronizes when the attribute is sent through a filter that is set to the configuration for Scenario Two. The table shows different outputs depending upon the following:

- ◆ Whether the attribute comes from the Identity Vault or the Application
- ◆ If the attribute is single-valued or multi-valued
- ◆ If the attribute is empty or non-empty

Table 9-2 Output of Scenario Two

	Identity Vault single-valued empty	Identity Vault single-valued non-empty	Identity Vault multi-valued empty	Identity Vault multi-valued non-empty
Application single-valued empty	No change	App = Identity Vault	No change	App = Identity Vault[1]
Application single-valued empty	App = empty	App = Identity Vault	Identity Vault = App	App = Identity Vault[1]
Application multi-valued empty	No change	App = Identity Vault	No change	App = Identity Vault
Application multi-valued non-empty	App = empty	App = Identity Vault	App = empty	App = Identity Vault

9.4.3 Scenario Three

The attribute is set to *Synchronize* on the Publisher channel, or the merge authority is set to *Application*.

Figure 9-3 Scenario Three

Class Name: User
Attribute Name: DirXML-ADAliasName

Publish

Synchronize
 Ignore
 Notify
 Reset

Subscribe

Synchronize
 Ignore
 Notify
 Reset

Merge Authority

Default
 Identity Vault
 Application
 None

Optimize modifications to Identity Vault

Yes
 No

The following table contains the values that the Metadirectory engine synchronizes when the attribute is sent through a filter that is set to the configuration for Scenario Three. The table shows different outputs depending upon the following:

- ◆ Whether the attribute comes from the Identity Vault or the Application
- ◆ If the attribute is single-valued or multi-valued
- ◆ If the attribute is empty or non-empty

Table 9-3 *Output of Scenario Three*

	Identity Vault single-valued empty	Identity Vault single-valued non-empty	Identity Vault multi-valued empty	Identity Vault multi-valued non-empty
Application single-valued empty	No change	Identity Vault = empty	No change	Identity Vault = empty
Application single-valued non-empty	Identity Vault = App	Identity Vault = App	Identity Vault = App	Identity Vault = App
Application multi-valued empty	No change	Identity Vault = empty	No change	Identity Vault = empty
Application multi-valued non- empty	Identity Vault = App[1]	Identity Vault = App[1]	Identity Vault = App	Identity Vault = App

- ♦ Section 10.1, “High-Level View,” on page 127
- ♦ Section 10.2, “Logical Database Classes,” on page 127
- ♦ Section 10.3, “Indirect Synchronization,” on page 127
- ♦ Section 10.4, “Direct Synchronization,” on page 135
- ♦ Section 10.5, “Synchronizing Primary Key Columns,” on page 138
- ♦ Section 10.6, “Synchronizing Multiple Classes,” on page 138
- ♦ Section 10.7, “Mapping Multivalue Attributes to Single-Value Database Fields,” on page 139

10.1 High-Level View

The following table shows a high-level view of how the driver maps Novell® Identity Vault objects to database objects.

Table 10-1 Mapping Identity Vault Objects to Database Objects

Identity Vault Object	Database Object
Tree	Schema
Class	Table/View
Attribute	Column
Association	Primary Key

10.2 Logical Database Classes

A logical database class is the set of tables or the view used to represent an eDirectory™ class in a database. A logical database class can consist of a single view or one parent table and zero or more child tables.

The name of a logical database class is the name of the parent table or view.

10.3 Indirect Synchronization

In an indirect synchronization model, the driver maps the following:

Table 10-2 Mappings in Indirect Synchronization

Identity Vault Object	Database Object
Classes	Tables
Attributes	Columns
1 Class	1 parent table and 0 or more child tables
Single-value attribute	Parent table column
Multivalue attribute	Parent table column (holding delimited values) or Child table column (preferred)

- ◆ [Section 10.3.1, “Mapping eDirectory Classes to Logical Database Classes,” on page 128](#)
- ◆ [Section 10.3.2, “Parent Tables,” on page 130](#)
- ◆ [Section 10.3.3, “Parent Table Columns,” on page 130](#)
- ◆ [Section 10.3.4, “Child Tables,” on page 131](#)
- ◆ [Section 10.3.5, “Referential Attributes,” on page 132](#)
- ◆ [Section 10.3.6, “Single-Value Referential Attributes,” on page 132](#)
- ◆ [Section 10.3.7, “Multivalue Referential Attributes,” on page 133](#)

10.3.1 Mapping eDirectory Classes to Logical Database Classes

In the following example, the logical database class `usr` consists of the following:

- ◆ One parent table `usr`
- ◆ Two child tables: `usr_phone` and `usr_faxno`.

Logical class `usr` is mapped to the eDirectory class `User`.

```
CREATE TABLE indirect.usr
(
  idu          INTEGER NOT NULL,
  fname       VARCHAR2(64),
  lname       CHAR(64),
  pwdminlen   NUMBER(4),
  pwdeptime  DATE,
  disabled    NUMBER(1),
  username    VARCHAR2(64),
  loginame    VARCHAR2(64),
  photo       LONG RAW,
  manager     INTEGER,
  CONSTRAINT pk_usr_idu      PRIMARY KEY (idu),
  CONSTRAINT fk_usr_manager FOREIGN KEY (manager)
```



```

        REFERENCES indirect.usr(idu)
    )
CREATE TABLE indirect.usr_phone
(
    idu          INTEGER          NOT NULL,
    phoneno     VARCHAR2(64)     NOT NULL,
    CONSTRAINT fk_phone_idu FOREIGN KEY (idu)
        REFERENCES indirect.usr(idu)
)
CREATE TABLE indirect.usr_fax
(
    idu          INTEGER          NOT NULL,
    faxno       VARCHAR2(64)     NOT NULL,
    CONSTRAINT fk_fax_idu FOREIGN KEY (idu)
        REFERENCES indirect.usr(idu)
)
<rule name="Schema Mapping Rule">
    <attr-name-map>
        <class-name>
            <nds-name>User</nds-name>
            <app-name>indirect.usr</app-name>
        </class-name>
        <attr-name class-name="User">
            <nds-name>Given Name</nds-name>
            <app-name>fname</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>Surname</nds-name>
            <app-name>lname</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>Password Expiration Time</nds-name>
            <app-name>pwdexptime</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>jpegPhoto</nds-name>
            <app-name>photo</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>manager</nds-name>
            <app-name>manager</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>Password Minimum Length</nds-name>
            <app-name>pwdminlen</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>Facsimile Telephone Number</nds-name>
            <app-name>usr_fax.faxno</app-name>
        </attr-name>
        <attr-name class-name="User">
            <nds-name>Telephone Number</nds-name>
            <app-name>usr_phone.phoneno</app-name>
    </attr-name-map>

```

```

    </attr-name>
    <attr-name class-name="User">
      <nds-name>Login Disabled</nds-name>
      <app-name>disabled</app-name>
    </attr-name>
  </attr-name-map>
</rule>

```

10.3.2 Parent Tables

Parent tables are tables with an explicit primary key constraint that contains one or more columns. In a parent table, an explicit primary key constraint is required so that the driver knows which fields to include in an association value.

```

CREATE TABLE indirect.usr
(
  idu  INTEGER  NOT NULL,
  -- ...
  CONSTRAINT pk_usr_idu PRIMARY KEY (idu)
)

```

The following table contains sample data for table `indirect.usr`.

idu	fname	lname
1	John	Doe

The resulting association for this row is

```
idu=1, table=usr, schema=indirect
```

NOTE: The case of database identifiers in association values is determined dynamically from database metadata at runtime.

10.3.3 Parent Table Columns

Parent table columns can contain only one value. As such, they are ideal for mapping single-value eDirectory attributes, such as mapping the single-value eDirectory attribute Password Minimum Length to the single-value parent table column `pwdminlen`.

Parent table columns are implicitly prefixed with the schema name and name of the parent table. It is not necessary to explicitly table-prefix parent table columns. For example, `indirect.usr.fname` is equivalent to `fname` for schema mapping purposes.

```

<rule name="Schema Mapping Rule">
  <attr-name-map>
    <class-name>
      <nds-name>User</nds-name>
      <app-name>indirect.usr</app-name>
    </class-name>
    <attr-name class-name="User">
      <nds-name>Given Name</nds-name>
      <app-name>fname</app-name>
    </attr-name>
  </attr-name-map>
</rule>

```

```
</attr-name-map>
</rule>
```

Large binary and string data types should usually be mapped to parent table columns. To map to a child table column, data types must be comparable in SQL statements. Large data types usually cannot be compared in SQL statements.

Large binary and string data types can be mapped to child table columns if the following occur:

- ◆ Each `<remove-value>` event on these types is transformed in a policy into a `<remove-all-values>` element
- ◆ An `<add-value>` element follows each `<remove-value>` event

10.3.4 Child Tables

A child table is a table that has a foreign key constraint on its parent table's primary key, linking the two tables together. The columns that comprise the child table's foreign key can have different names than the columns in the parent table's primary key.

The following example shows the relationship between parent table `usr` and child tables `usr_phone` and `usr_faxno`:

```
CREATE TABLE indirect.usr
(
    idu    INTEGER    NOT NULL,
    -- ...
    CONSTRAINT pk_usr_idu    PRIMARY KEY (idu)
)
CREATE TABLE indirect.usr_phone
(
    idu        INTEGER        NOT NULL,
    phoneno    VARCHAR2(64)    NOT NULL,
    CONSTRAINT fk_phone_idu    FOREIGN KEY (idu)
        REFERENCES indirect.usr(idu)
)
CREATE TABLE indirect.usr_fax
(
    idu        INTEGER        NOT NULL,
    faxno      VARCHAR2(64)    NOT NULL,
    CONSTRAINT fk_fax_idu      FOREIGN KEY (idu)
        REFERENCES indirect.usr(idu)
)
```

NOTE: In a child table, constrain all columns NOT NULL.

The first constrained column in a child table identifies the parent table. In the above example, the constrained column in child table `usr_phone` is `idu`. The only purpose of this column is to relate tables `usr_phone` and `usr`. Because constrained columns do not contain any useful information, omit them from publication triggers and Schema Mapping policies.

The unconstrained column is the column of interest. It represents a single, multivalued attribute. In the above example, the unconstrained columns are `phoneno` and `faxno`. Because unconstrained columns can hold multiple values, they are ideal for mapping multivalued eDirectory attributes (for

example, mapping the multivalue eDirectory attribute Telephone Number to `usrphone.phoneno`).

The following table contains sample data for `indirect.usr_phone`.

Table 10-3 *Sample Data*

idu	phoneno
1	111-1111
1	222-2222

Like parent table columns, child table columns are implicitly schema-prefixed. Unlike parent table columns, however, a child table column name must be explicitly prefixed with the child table name (for example, `usr_phone.phoneno`). Otherwise, the driver implicitly interprets column `phoneno` (the parent table column) as `usr.phoneno`, not the child table column `usr_phone.phoneno`.

```
<rule name="Schema Mapping Rule">
  <attr-name-map>
    <class-name>
      <nds-name>User</nds-name>
      <app-name>indirect.usr</app-name>
    </class-name>
    <attr-name class-name="User">
      <nds-name>Facsimile Telephone Number</nds-name>
      <app-name>usr_fax.faxno</app-name>
    </attr-name>
    <attr-name class-name="User">
      <nds-name>Telephone Number</nds-name>
      <app-name>usr_phone.phoneno</app-name>
    </attr-name>
  </attr-name-map>
</rule>
```

NOTE: Map each multivalue eDirectory attribute to a different child table.

10.3.5 Referential Attributes

You can represent referential containment in the database by using foreign key constraints. Referential attributes are columns within a logical database class that refer to the primary key columns of parent tables in the same logical database class or those of other logical database classes.

10.3.6 Single-Value Referential Attributes

You can relate two parent tables through a single-value parent table column. This column must have a foreign key constraint pointing to the other parent table's primary key. The following example relates a single parent table `usr` to itself:

```
CREATE TABLE indirect.usr
(
  idu          INTEGER NOT NULL,
```

```

-- ...
manager INTEGER,
CONSTRAINT pk_usr_idu PRIMARY KEY (idu),
CONSTRAINT fk_usr_manager FOREIGN KEY (manager)
REFERENCES indirect.usr(idu)
)

```

NOTE: Single-valued referential columns should be nullable.

```

<rule name="Schema Mapping Rule">
  <attr-name-map>
    <class-name>
      <nds-name>User</nds-name>
      <app-name>indirect.usr</app-name>
    </class-name>
    <attr-name class-name="User">
      <nds-name>manager</nds-name>
      <app-name>manager</app-name>
    </attr-name>
  </attr-name-map>
</rule>

```

The interpretation of the above example is that each user can have only one manager who himself is a user.

10.3.7 Multivalue Referential Attributes

You can relate two parent tables through a common child table. This child table must have a column constrained by a foreign key pointing to the other parent table's primary key. The following example relates two parent tables `usr` and `grp` through a common child table `member`.

```

CREATE TABLE indirect.usr
(
  idu INTEGER NOT NULL,
  -- ...
  CONSTRAINT pk_usr_idu PRIMARY KEY (idu)
)
CREATE TABLE indirect.grp
(
  idg INTEGER NOT NULL,
  -- ...
  CONSTRAINT pk_grp_idg PRIMARY KEY (idg)
)
CREATE TABLE indirect.grp_member
(
  idg INTEGER NOT NULL,
  idu INTEGER NOT NULL,
  CONSTRAINT fk_member_idg FOREIGN KEY (idg) REFERENCES
indirect.grp(idg),
  CONSTRAINT fk_member_idu FOREIGN KEY (idu)
REFERENCES indirect.usr(idu)
)

```

NOTE: Constrain all columns in a child table NOT NULL.

```

<rule name="Schema Mapping Rule">
  <attr-name-map>
    <class-name>
      <nds-name>Group</nds-name>
      <app-name>indirect.grp</app-name>
    </class-name>
    <class-name>
      <nds-name>User</nds-name>
      <app-name>indirect.usr</app-name>
    </class-name>
    <attr-name class-name="Group">
      <nds-name>Member</nds-name>
      <app-name>grp_member.idu</app-name>
    </attr-name>
  </attr-name-map>
</rule>

```

The first constrained column in a child table determines which logical database class the child table `grp_member` belongs to. In the above example, `grp_member` is considered to be part of logical database class `grp`. `grp_member` is said to be a proper child of `grp`. The second constrained column in a child table is the multivalued referential attribute.

In the following example, the order of the constrained columns has been reversed so that `grp_member` is part of class `usr`. To more accurately reflect the relationship, table `grp_member` has been renamed to `usr_mbr_of`.

```

CREATE TABLE indirect.usr
(
  idu  INTEGER  NOT NULL,
  -- ...
  CONSTRAINT pk_usr_idu PRIMARY KEY (idu)
)
CREATE TABLE indirect.grp
(
  idg  INTEGER  NOT NULL,
  -- ...
  CONSTRAINT pk_grp_idg PRIMARY KEY (idg)
)
CREATE TABLE indirect.usr_mbr_of
(
  idu  INTEGER  NOT NULL,
  idg  INTEGER  NOT NULL,
  CONSTRAINT fk_mbr_of_idu FOREIGN KEY (idu)
    REFERENCES indirect.usr(idu) ON DELETE CASCADE,
  CONSTRAINT fk_mbr_of_idg FOREIGN KEY (idg)
    REFERENCES indirect.grp(idg) ON DELETE CASCADE
)
<rule name="Schema Mapping Rule">
  <attr-name-map>
    <class-name>
      <nds-name>Group</nds-name>
      <app-name>indirect.grp</app-name>
    </class-name>
    <class-name>

```

```

        <nds-name>User</nds-name>
        <app-name>indirect.usr</app-name>
    </class-name>
    <attr-name class-name="User">
        <nds-name>Group Membership</nds-name>
        <app-name>usr_mbr_of.idg</app-name>
    </attr-name>
</attr-name-map>
</rule>

```

In databases that have no awareness of column position (such as DB2/AS400), order is determined by sorting column names by string or hexadecimal value. For additional information, see [“Sort Column Names By” on page 80](#).

In general, it is necessary to synchronize only bidirectional, multivalued, referential attributes as part of one class or the other, not both. If you want to synchronize referential attributes for both classes, construct two child tables, one for each class. For example, if you want to synchronize eDirectory attributes Group Membership and Member, you need two child tables.

In practice, when you synchronize User and Group classes, we recommend that you synchronize the Group Membership attribute of class User instead of the Member attribute of class Group. Synchronizing the group memberships of a user is usually more efficient than synchronizing all members of a group.

10.4 Direct Synchronization

In a direct synchronization model, the driver maps the following:

Table 10-4 Mappings in Direct Synchronization

Identity Vault Object	Database Object
Classes	Views
Attributes	View Columns
Class	View
Single-value attribute	View Column
Multivalued attribute	View Column

The update capabilities of views vary between databases. Most databases allow views to be updated when they are comprised of a single base table. (That is, they do not join multiple tables.) If views are strictly read-only, they cannot be used for subscription. Some databases allow update logic to be defined on views in instead-of-triggers, which allow a view to join multiple base tables and still be updateable.

For a list of databases that support instead-of-triggers, see [“Database Features” on page 178](#). Instead-of-trigger logic can be simulated, regardless of database capability using embedded SQL. See [Section 13.4, “Virtual Triggers,” on page 158](#).

- ◆ [Section 10.4.1, “View Column Meta-Identifiers,” on page 136](#)
- ◆ [Section 10.4.2, “Primary Key Columns,” on page 138](#)

- ♦ [Section 10.4.3, “Schema Mapping,” on page 138](#)

10.4.1 View Column Meta-Identifiers

A view is a logical table. Unlike tables, views do not physically exist in the database. As such, views usually cannot have traditional primary key/foreign key constraints. To simulate these constructs, the driver for JDBC embeds constraints and other metadata in view column names. The difference between these constraints and traditional ones is that the former are not enforced at the database level. They are an application-level construct.

For example, to identify to the driver which fields to use when constructing association values, place a primary key constraint on a parent table. The corollary to this for a view is to prefix one or more column names with `pk_` (case-insensitive).

The following table lists the constraint prefixes that can be embedded in view column names.

Table 10-5 *Constraint Prefixes*

Constraint Prefixes (case-insensitive)	Interpretation
<code>pk_</code>	primary key
<code>fk_</code>	foreign key
<code>sv_</code>	single-value
<code>mv_</code>	multivalued

The following example views contain all of these constraint prefixes:

```
CREATE VIEW direct.view_usr
(
  pk_idu,           -- primary key column; implicitly single-valued
  sv_fname,        -- single-valued column
  mv_phoneno,      -- multi-valued column
  fk_idu_manager,  -- self-referential foreign key column; refers
                  -- to primary key column idu in view_usr;
                  -- implicitly single-valued
  fk_mv_idg_mbr_of -- extra-referential foreign key column; refers
                  -- to primary key column idg in view_grp;
                  -- multi-valued
)
AS
-- ...

CREATE VIEW direct.view_grp
(
  pk_idg,           -- primary key column; implicitly single-valued
  fk_mv_idu_mbr     -- extra-referential foreign key column; refers
                  -- to primary key column idu in view_usr;
                  -- multi-valued
)
AS
-- ...
```


BNF

The BNF ([Backus Naur Form \(http://cui.unige.ch/db-research/Enseignement/analyseinfo/AboutBNF.html\)](http://cui.unige.ch/db-research/Enseignement/analyseinfo/AboutBNF.html)) notation for view column meta-identifiers:

```
<view-column-name> ::= [<meta-info>] <column-name>

<column-name> ::= <legal-unquoted-database-identifier>
<meta-info> ::= <referential> | <non-referential>

<non-referential> ::= [<single-value> | <multiple-value>]

<single-value> ::= "sv_"

<multiple-value> ::= "mv_"

<referential> ::= <primary-key> | <foreign-key>

<primary-key> ::= "pk_" [<single-value>] [<column-group-id>]
                [<referenced-column-name>]
<column-group-id> ::= <non-negative-integer> "_"
<referenced-column-name> ::= "_" <column-name> "__"

<foreign-key> ::= "fk_" [<non-referential>] [<column-group-id>]
                <referenced-column-name>
```

Normalized Forms

By default, all view column names are single-valued. Therefore, explicitly specifying the `sv_` prefix in a view column name is redundant. For example, `sv_fname` and `fname` are equivalent forms of the same column name.

Also, primary key column names implicitly refer to themselves. Therefore, it is redundant to specify the referenced column name. For example, `pk_idu` is equivalent to `pk__idu__idu`.

The Driver for JDBC uses two normalized forms of view meta-identifiers:

- ◆ Database native form

Database native form is the column name as declared in the database. This form is usually much more verbose than schema mapping form, and contains all necessary meta information.

- ◆ Schema mapping form

Schema mapping form is returned when the driver returns the application schema. This form is much more concise than database native form because much of the meta information included in database native form is represented in XDS XML and not in the identifier.

The referential prefixes `pk_` and `fk_` are the only meta information preserved in schema mapping form. This limitation ensures backward compatibility.

The following table provides examples of each form:

Table 10-6 Example Normalized Forms of View Meta-Identifiers

Database Native Form	Schema Mapping Form
pk_idu	pk_idu
sv_fname	fname
mv_phoneno	phoneno
fk_mv__idg__mbr_of	fk_mbr_of

Equivalent Forms

A view column name without meta information is called its “effective” name, which is similar to a directory objects’ “effective” rights. For the driver, view column name equivalency is determined without respect to meta information by default. For example, `pk_idu` is equivalent to `idu`, and `fk_mv__idg__mbr_of` is equivalent to `mbr_of`. Any variant form of a view meta column identifier can be passed to the driver at runtime. For backward compatibility reasons, meta information can be treated as part of the effective view column name. See [“Enable Meta-Identifier Support?” on page 77](#).

10.4.2 Primary Key Columns

Primary key column names must be unique among all views in the synchronization schema.

10.4.3 Schema Mapping

Schema mapping conventions for views and view columns are equivalent to that used for parent tables and parent table columns.

10.5 Synchronizing Primary Key Columns

When the database is the authoritative source of primary key columns, generally omit the columns from the Publisher and Subscriber filters, Schema Mapping policies, and publication triggers.

When the Identity Vault is the authoritative source of primary key columns, include the columns in the Subscriber filter and Schema Mapping policies, but omit the columns from the Publisher filter and publication triggers. Also, GUID rather than CN is recommended for use as a primary key. CN is a multivalued attribute and can change. GUID has a single value and is static.

10.6 Synchronizing Multiple Classes

When synchronizing multiple eDirectory classes, synchronize each class to a different parent table or view. Each logical database class must have a unique primary key column name. The Publisher channel uses this common column name to identify all rows in the event log table pertaining to a single logical database class. For example, both the logical database classes `usr` and `grp` have a unique primary key column name.

```
CREATE TABLE usr
(
    idu    INTEGER          NOT NULL,
    lname  VARCHAR2(64)    NOT NULL,
```

```

--...
CONSTRAINT pk_usr_idu PRIMARY KEY(idu)
);
CREATE TABLE grp
(
  idg INTEGER NOT NULL,
  --...
CONSTRAINT pk_grp_idg PRIMARY KEY(idg)
);

```

10.7 Mapping Multivalue Attributes to Single-Value Database Fields

By default, the driver assumes that all eDirectory attributes mapped to parent table columns or view columns have a single value. Because the driver is unaware of the eDirectory schema, it has no way of knowing whether an eDirectory attribute has a single value or has multiple values. Accordingly, multivalue and single-value attribute mappings are handled identically.

The driver implements the Most Recently Touched (MRT) algorithm with regard to single-value parent table or view columns. An MRT algorithm ensures that the most recently added attribute value or most recently deleted attribute value is stored in the database. The algorithm is adequate if the attribute in question has a single value.

If the attribute has multiple values, the algorithm has some undesirable consequences. When a value is deleted from a multivalue attribute, the database field it is mapped to is set to `NULL` and remains `NULL` until another value is added. The preferred solution to this undesirable behavior is to extend the eDirectory schema so that only single-value attributes are mapping to parent table or view columns.

Other solutions include the following:

- ◆ For indirect synchronization, map each multivalue attribute to its own child table.
- ◆ For both direct or indirect synchronization, use a policy to delimit multiple values before inserting them into a table or view column.
- ◆ Implement a first or last value per replica policy in style sheets by using methods provided in the `com.novell.nds.indirect.driver.jdbc.util.MappingPolicy` class. Under a first-value-per-replica (FPR) policy, the first attribute value on the eDirectory replica is always synchronized. Under a last-value-per-replica (LPR) policy, the last attribute value on a replica is always synchronized.

By using global configuration values, you can configure the sample driver configuration to use either FPR or LPR mapping policies. Multivalue to single-value attribute mapping policies are contained in the Subscriber Command Transformation policy container. The sample driver configuration maps the multivalue eDirectory attributes Given Name and Surname to the single-value columns `fname` and `lname` respectively.

Mapping XDS Events to SQL Statements

11

- ◆ [Section 11.1, “Mapping XDS Events for Indirect Synchronization,” on page 141](#)

11.1 Mapping XDS Events for Indirect Synchronization

The following table summarizes how the Subscriber channel maps XDS events to DML SQL statements for indirect synchronization:

Table 11-1 *Mapping XDS Events for Indirect Synchronization*

XML Event	SQL Equivalent
<add>	0 or more select statements, depending upon the matching policy 1 parent table insert statement for all single value <add-attr> elements 0 or 1 stored procedure/function calls to retrieve primary key values before or after the parent table insert statement 1 child table insert statement for each multivalued <add-attr> element
<modify>	1 parent table update statement for each single value <add-value> or <remove-value> element 1 child table insert statement for each multivalued <add-value> element 1 child table delete statement for each <remove-value> element
<delete>	1 parent table delete statement 1 delete statement for each child table
<query>	1 parent table select statement 1 select statement for each child table
<move> <rename> <modify-password> <check-object-password>	0 statements unless bound to embedded SQL statements

The following table summarizes how the Subscriber channel maps XDS events to DML SQL statements for direct synchronization:

Table 11-2 Mapping XDS Events for Direct Synchronization

XML Event	SQL Equivalent
<add>	0 or more select statements, depending upon the matching policy 1 view insert statement for all single value <add-attr> element 0 or 1 stored procedure/function call to retrieve primary key values before or after the view insert statement 1 view insert statement for each multivalued <add-attr> element
<modify>	1 view update statement for each single value <add-value> or <remove-value> element 1 view insert statement for each multivalued <add-value> element 1 view delete statement for each <remove-value> element
<delete>	1 view delete statement
<query>	1 view select statement
<move> <rename> <modify-password> <check-object-password>	0 statements unless bound to embedded SQL statements

The Event Log Table

12

The event log table stores publication events. This section discusses the structure and capabilities of the event log table.

You can customize the name of the event log table and its columns to avoid conflicts with reserved database keywords. The order, number, and data types of its columns, however, are fixed. In databases that are unaware of column position, order is determined by the Sort Column Names By parameter. See “Sort Column Names By” on page 80.

Events in this table can be ordered either by order of insertion (the `record_id` column) or chronologically (the `event_time` column). Ordering events chronologically allows event processing to be delayed. To order publication events chronologically, set the Enable Future Event Processing parameter to Boolean True. See “Enable Future Event Processing?” on page 91.

- ♦ Section 12.1, “Event Log Columns,” on page 143
- ♦ Section 12.2, “Event Types,” on page 145

12.1 Event Log Columns

This section describes columns in the event log table. Columns are ordered by position.

1. `record_id`

The `record_id` column is used to uniquely identify rows in the event log table and order publication events. This column must contain sequential, ascending, positive, unique integer values. Gaps between `record_id` values no longer prematurely end a polling cycle.

2. `table_key`

Format values for this column exactly the same in all triggers for a logical database class. The BNF or Backus Naur Form (<http://cui.unige.ch/db-research/Enseignement/analyseinfo/AboutBNF.html>) of this parameter is defined below:

```
<table-key> ::= <unique-row-identifier> {"+"  
                <unique-row-identifier>}
```

```
<unique-row-identifier> ::= <primary-key-column-name> "=" <value>
```

For example, for the `usr` table referenced throughout this chapter, this column’s value might be `idu=1`.

For the `view_usr` view referenced throughout this chapter, this column’s value might be `pk_empno=1`.

For a hypothetical compound primary key (one containing multiple columns), this column’s value might be `pkey1=value1+pkey2=value2`.

NOTE: If primary key values placed in the `table_key` field contains any of the special characters `{, ;, ' + " = \ < > }`, where `' { ' }` contain the set of special characters, delimit the value with double quotes. You’ll also need to escape the double quote character `"` as `\` and the literal escape character `\` character as `\\` when contained inside a pair of double quotes.

For a hypothetical primary key containing special characters, this column's value might be `pkey=", ; ' + \" = \\ < >".` (Note the double quotes and escaped characters.)

NOTE: Differences in padding or formatting might result in out-of-order event processing. For performance reasons, remove any unnecessary white space from numeric values. For example, "idu=1" is preferred over "idu= 1". (Note the space in "idu= 1".)

3. `status`

The `status` column indicates the state of a given row. The following table lists permitted values:

Table 12-1 *Permitted Values for Status Columns*

Character Value	Interpretation
N	new
S	success
W	warning
E	error
F	fatal

To be processed, all rows inserted into the event log table must have a `status` value of N. The remainder of the status characters are used solely by the Publisher channel to designate processed rows. All other characters are reserved for future use.

NOTE: Status values are case sensitive.

4. `event_type`

Values in this column must be between 1 and 8. All other numbers are reserved for future use.

The following table describes each event type:

Table 12-2 *Event Types*

Event Type	Interpretation
1	insert field
2	update field
3	update field (remove all values)
4	delete row
5	insert row (query-back)
6	update row (query-back)
7	insert field (query-back)
8	update field (query-back)

For additional information on this field, see [Section 12.2, “Event Types,” on page 145](#).

5. `event_time`

This column serves as an alternative ordering column to `record_id`. It contains the effective date of the event. It must not be `NULL`. For this column to become the ordering column, set the `Enable Future Event Processing` parameter to Boolean `True`. See [“Enable Future Event Processing?” on page 91](#).

6. `perpetrator`

This column identifies the database user who instigated the event. A `NULL` value is interpreted as a user other than the driver user. As such, rows with a `NULL` value or value not equal to the driver’s database username are published. Rows with a value equal to the driver’s database username are not published unless the `Allow Loopback Publisher` parameter is set to Boolean `True`. See [“Allow Loopback?” on page 94](#).

7. `table_name`

The name of the table or view where the event occurred.

8. `column_name`

The name of the column that was changed. This column is used only for per-field (1-3, 7-8) event types. Nevertheless, it must always be present in the event log table. If it is missing, the `Publisher` channel cannot start.

9. `old_value`

The field’s old value. This column is used only for per-field, non-query-back event types (1-3). Nevertheless, it must always be present in the event log table. If it is missing, the `Publisher` channel cannot start.

10. `new_value`

The field’s new value. This column is used only by per-field, non-query-back event types (1-3). Nevertheless, it must always be present in the event log table. If it is missing, the `Publisher` channel cannot start.

12.2 Event Types

The following table describes each event type:

Table 12-3 *Event Types*

Event Type	Interpretation
1	insert field
2	update field
3	update field (remove all values)
4	delete row
5	insert row (query-back)
6	update row (query-back)
7	insert field (query-back)
8	update field (query-back)

Event types are in four major categories. Some categories overlap. The following table describes each category and indicates which event types are members:

Table 12-4 *Event Categories and Types*

Event Category	Event Types
Per-field (attribute)	1, 2, 3, 7, 8
Per-row (object)	4, 5, 6
Non-query-back	1, 2, 3, 4
Query-back	5, 6, 7, 8
Per-field, non-query-back	1, 2, 3
Per-field, query-back	7, 8
Per-row, non-query-back	4
Per-row, query-back	5, 6

In general, a combination of event types from each category yields the best trade-off in terms of space, time, implementation complexity, and performance.

Per-field event types are more granular, require more space, and are more complex to implement than per-row event types. Per-row events are less granular, require less space, and are easier to implement than per-field event types.

Query-back event types use less space but require more time to process than non-query-back event types. Non-query-back event types use more space but require less time to process than query-back event types.

Query-back event types trump their non-query-back counterparts. Non-query-back events are ignored if a query-back event is logged for the same field or object. For example, if an event of type 2 (update-field, non-query-back) and 8 (update-field, query-back) are logged on the same field, the type 2 event is ignored in favor of the type 8 event.

Furthermore, query-back row event types trump query-back field event types. For example, if an event type 8 (update field, query-back) and a event type 6 (update row query-back) are logged on the same object, the type 8 event is ignored in favor of the type 6 event.

Query-back events are ignored by the Publisher if the database object no longer exists. They are dependent upon the database object still being around at processing time. Therefore, logged query-back adds and modifies (event types 5, 6, 7, 8) have no effect once the database object they refer to is deleted.

The following table shows the basic correlation between publication event types and the XDS XML generated by the Publisher channel.

Table 12-5 Basic Correlation of Publication Event Types

Event Type	Resulting XDS
insert	<add>
update	<modify>
delete	<delete>

The following example illustrates XML that the Publisher channel generates for events logged on the `usr` table for each possible event type.

```
CREATE TABLE indirect.usr
(
  idu    INTEGER NOT NULL,
  fname  VARCHAR2(64),
  photo  LONGRAW,
  --...
  CONSTRAINT pk_usr_idu PRIMARY KEY(idu)
);
```

The following table shows the initial contents of `usr` after a new row has been inserted:

Table 12-6 An Inserted Row in the `usr` Table

idu	fname	lname	photo
1	Jack	Frost	0xAAAA

The following table shows the current contents of `usr` after the row has been updated:

Table 12-7 An Updated Row in the `usr` Table

idu	fname	lname	photo
1	John	Doe	0xB BBBB

Insert Field

The table below shows the contents of the event log table after a new row is inserted into table `usr`. The value for column `photo` has been Base64-encoded. The Base64-encoded equivalent of `0xAAAA` is `qqo=`.

Table 12-8 Event Log Table: Type 1

event_type	table	table_key	column_name	old_value	new_value
1	usr	idu=1	fname	NULL	Jack
1	usr	idu=1	lname	NULL	Frost
1	usr	idu=1	photo	NULL	qqo=

The Publisher channel generates the following XML:

```
<add class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <add-attr attr-name="fname">
    <value type="string">Jack</value>
  </add-attr>
  <add-attr attr-name="lname">
    <value type="string">Frost</value>
  </add-attr>
  <add-attr attr-name="photo">
    <value type="octet">qgo=</value>
  </add-attr>
</add>
```

Update Field

The following table shows the contents of the event log table after the row in table `usr` has been updated. The values for column `photo` has been Base64-encoded. The Base64-encoded equivalent of `0xBBBB` is `u7s=`.

Table 12-9 Event Log Table: Type 2

event_type	table	table_key	column_name	old_value	new_value
2	usr	idu=1	fname	Jack	John
2	usr	idu=1	lname	Frost	Doe
2	usr	idu=1	photo	qgo=	u7s=

The Publisher channel generates the following XML:

```
<modify class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <modify-attr attr-name="fname">
    <remove-value>
      <value type="string">Jack</value>
    </remove-value>
    <add-value>
      <value type="string">John</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="lname">
    <remove-value>
      <value type="string">Frost</value>
    </remove-value>
    <add-value>
      <value type="string">Doe</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="photo">
    <remove-value>
```

```

        <value type="octet">qqo=</value>
    </remove-value>
    <add-value>
        <value type="octet">u7s=</value>
    </add-value>
</modify-attr>
</modify>

```

Update Field (Remove-All-Values)

The following table shows the contents of the event log table after the row in table `usr` has been updated. The value for column `photo` has been Base64-encoded.

Table 12-10 Event Log Table: Type 3

event_type	table	table_key	column_name	old_value	new_value
3	usr	idu=1	fname	Jack	John
3	usr	idu=1	lname	Frost	Doe
3	usr	idu=1	photo	qqo=	u7s=

The Publisher channel generates the following XML:

```

<modify class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <modify-attr attr-name="fname">
    <remove-all-values/>
    <add-value>
      <value type="string">John</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="lname">
    <remove-all-values/>
    <add-value>
      <value type="string">Doe</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="photo">
    <remove-all-values/>
    <add-value>
      <value type="octet">u7s=</value>
    </add-value>
  </modify-attr>
</modify>

```

Delete Row

The table below shows the contents of the event log table after the row in table `usr` has been deleted.

Table 12-11 Event Log Table: Type 4

event_type	table	table_key	column_name	old_value	new_value
4	usr	idu=1	NULL	NULL	NULL

The Publisher channel generates the following XML:

```
<delete class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
</delete>
```

Insert Row (Query-Back)

The following table shows the contents of the event log table after a new row is inserted into table `usr`.

Table 12-12 Event Log Table: Type 5

event_type	table	table_key	column_name	old_value	new_value
5	usr	idu=1	NULL	NULL	NULL

The Publisher channel generates the following XML. The values reflect the current contents of table `usr`, not the initial contents.

```
<add class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <add-attr attr-name="fname">
    <value type="string">John</value>
  </add-attr>
  <add-attr attr-name="lname">
    <value type="string">Doe</value>
  </add-attr>
  <add-attr attr-name="photo">
    <value type="octet">u7s=</value>
  </add-attr>
</add>
```

Update Row (Query-Back)

The table below shows the contents of the event log table after the row in table `usr` has been updated.

Table 12-13 Event Log Table: Type 6

event_type	table	table_key	column_name	old_value	new_value
6	usr	idu=1	NULL	NULL	NULL

The Publisher channel generates the following XML. The values reflect the current contents of table `usr`, not the initial contents.

```
<modify class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <modify-attr attr-name="fname">
    <remove-all-values/>
    <add-value>
      <value type="string">John</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="lname">
    <remove-all-values/>
    <add-value>
      <value type="string">Doe</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="photo">
    <remove-all-values/>
    <add-value>
      <value type="octet">u7s=</value>
    </add-value>
  </modify-attr>
</modify>
```

Insert Field (Query-Back)

The following table shows the contents of the event log table after a new row is inserted into table `usr`. Old and new values are omitted because they are not used.

Table 12-14 Event Log Table: Type 7

event_type	table	table_key	column_name	old_value	new_value
7	usr	idu=1	fname	NULL	NULL
7	usr	idu=1	lname	NULL	NULL
7	usr	idu=1	photo	NULL	NULL

The Publisher channel generates the following XML. The values reflect the current contents of table `usr`, not the initial contents.

```
<add class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <add-attr attr-name="fname">
    <value type="string">John</value>
  </add-attr>
  <add-attr attr-name="lname">
    <value type="string">Doe</value>
  </add-attr>
  <add-attr attr-name="photo">
```

```

    <value type="octet">u7s=</value>
  </add-attr>
</add>

```

Update Field (Query-Back)

The following table shows the contents of the event log table after the row in table `usr` has been updated. Old and new values are omitted because they are not used.

Table 12-15 Event Log Table: Type 8

event_type	table	table_key	column_name	old_value	new_value
8	usr	idu=1	fname	NULL	NULL
8	usr	idu=1	lname	NULL	NULL
8	usr	idu=1	photo	NULL	NULL

The Publisher channel generates the following XML. The values reflect the current contents of table `usr`, not the initial contents.

```

<modify class-name="usr">
  <association>idu=1,table=usr,schema=indirect
</association>
  <modify-attr attr-name="fname">
    <remove-all-values/>
    <add-value>
      <value type="string">John</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="lname">
    <remove-all-values/>
    <add-value>
      <value type="string">Doe</value>
    </add-value>
  </modify-attr>
  <modify-attr attr-name="photo">
    <remove-all-values/>
    <add-value>
      <value type="octet">u7s=</value>
    </add-value>
  </modify-attr>
</modify>

```


Embedded SQL Statements in XDS Events

13

Embedded SQL allows you to embed SQL statements in XDS-formatted XML documents. You can use embedded SQL statements along with XDS events or use them standalone. When embedded SQL statements are used standalone, embedded SQL processing does not require that the driver know anything about tables/view in the target database. As such, the driver can run in schema-unaware mode. See [“Synchronization Filter” on page 64](#). When using embedded SQL standalone, you must establish associations manually. The driver won’t establish them for you.

When used in conjunction with XDS events, embedded SQL can act as a virtual database trigger. In the same way that you can install database triggers on a table and cause side effects in a database when certain SQL statements are executed, embedded SQL can cause side effects in a database in response to certain XDS events.

All examples in this section reference the following `indirect_usr` table.

```
CREATE TABLE indirect_usr
(
  idu    INTEGER NOT NULL,
  fname  VARCHAR2(64),
  lname  VARCHAR2(64),

  CONSTRAINT pk_usr_idu PRIMARY KEY(idu)
);
```

- ◆ [Section 13.1, “Common Uses of Embedded SQL,” on page 154](#)
- ◆ [Section 13.2, “Embedded SQL Basics,” on page 154](#)
- ◆ [Section 13.3, “Token Substitution,” on page 155](#)
- ◆ [Section 13.4, “Virtual Triggers,” on page 158](#)
- ◆ [Section 13.5, “Manual vs. Automatic Transactions,” on page 159](#)
- ◆ [Section 13.6, “Transaction Isolation Level,” on page 160](#)
- ◆ [Section 13.7, “Statement Type,” on page 161](#)
- ◆ [Section 13.8, “SQL Queries,” on page 162](#)
- ◆ [Section 13.9, “Data Definition Language \(DDL\) Statements,” on page 163](#)
- ◆ [Section 13.10, “Logical Operations,” on page 164](#)
- ◆ [Section 13.11, “Implementing Password Set with Embedded SQL,” on page 164](#)
- ◆ [Section 13.12, “Implementing Modify Password with Embedded SQL,” on page 165](#)
- ◆ [Section 13.13, “Implementing Check Object Password,” on page 165](#)
- ◆ [Section 13.14, “Calling Stored Procedures and Functions,” on page 166](#)
- ◆ [Section 13.15, “Best Practices,” on page 174](#)

13.1 Common Uses of Embedded SQL

You can accomplish the following by embedding SQL in XDS events:

- ♦ Create database users or roles.
- ♦ Manage user passwords
You can set, check or modify user passwords.
- ♦ Manage database user or role privileges.

For examples of each, consult the `User DDL Command Transformation` style sheet on the Subscriber channel in the example driver configuration.

13.2 Embedded SQL Basics

- ♦ [Section 13.2.1, “Elements,” on page 154](#)
- ♦ [Section 13.2.2, “Namespaces,” on page 154](#)
- ♦ [Section 13.2.3, “Embedded SQL Example,” on page 155](#)

13.2.1 Elements

SQL is embedded in XDS events through the `<jdbc:statement>` and `<jdbc:sql>` elements. The `<jdbc:statement>` element can contain one or more `<jdbc:sql>` elements.

13.2.2 Namespaces

The namespace prefix `jdbc` used throughout this section is implicitly bound to the namespace `urn:dirxml:jdbc` when referenced outside of an XML document.

You must use namespace-prefixed embedded SQL elements and attributes. Otherwise, the driver will not recognize them. In all examples in this section, the prefix used is `jdbc`. In practice, the prefix can be whatever you want it to be, as long as it is bound to the namespace value `urn:dirxml:jdbc`.

The following XML example illustrates how to use and properly namespace-prefix embedded SQL elements. In the following example, the namespace declaration and namespace prefixes are bolded:

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John'
  </jdbc:sql>
  </jdbc:statement>
</input>
```

13.2.3 Embedded SQL Example

The following XML example illustrates how to use the `<jdbc:statement>` and `<jdbc:sql>` elements and their interpretation. In the following example, embedded SQL elements are bolded:

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John'
  </jdbc:sql>
</jdbc:statement>
</input>
```

Because the Subscriber channel resolves `<add>` events to one or more INSERT statements, the XML shown above resolves to:

```
SET AUTOCOMMIT OFF
INSERT INTO indirect.usr(lname)VALUES('Doe');
COMMIT;--explicit commit
UPDATE indirect.usr SET fname = 'John';
COMMIT; --explicit commit
```

13.3 Token Substitution

Rather than require you to parse field values from an association, the Subscriber channel supports token substitution in embedded SQL statements. In the following examples, tokens and the values they reference are bolded:

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <modify class-name="usr">
    <association>idu=1,table=usr,schema=indirect</association>
    <modify-attr name="lname">
      <add-value>
        <value>DoeRaeMe</value>
      </add-value>
    </modify-attr>
  </modify>
  <jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
      idu = { $idu }</jdbc:sql>
  </jdbc:statement>
</input>
```

Token placeholders must adhere to the XSLT attribute value template syntax `{ $field-name }`. Also, the referenced association element must precede the `<jdbc:statement>` element in the XDS document, or must be present as a child of the `<jdbc:statement>` element. Alternatively, instead of copying the association element as child of the `<jdbc:statement>` element, you could copy the `src-entry-id` of the element containing the association element onto the `<jdbc:statement>` element. Both approaches are bolded in the following examples:

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <modify class-name="usr">
```

```

    <association>idu=1,table=usr,schema=indirect</association>
    <modify-attr name="lname">
      <add-value>
        <value>DoeRaeMe</value>
      </add-value>
    </modify-attr>
  </modify>
</jdbc:statement>
  <association>idu=1,table=usr,schema=indirect</association>
  <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
    idu = {$idu}</jdbc:sql>
</jdbc:statement>

</input>
<input xmlns:jdbc="urn:dirxml:jdbc">
  <modify class-name="usr" src-entry-id="0">
    <association>idu=1,table=usr,schema=indirect</association>
    <modify-attr name="lname">
      <add-value>
        <value>DoeRaeMe</value>
      </add-value>
    </modify-attr>
  </modify>
  <jdbc:statement src-entry-id="0">
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
      idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>

```

The `{$field-name}` token must refer to one of the naming RDN attribute names in the association value. The above examples have only one naming attribute: `idu`.

An `<add>` event is the only event where an association element is not required to precede embedded SQL statements with tokens because the association has not been created yet. Additionally, any embedded SQL statements using tokens must follow, not precede, the `<add>` event. For example:

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
      idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>

```

To prevent tracing of sensitive information, you can use the `{$$password}` token to refer to the contents of the immediately preceding `<password>` element within the same document. In the following example, the password token and the value it refers to are bolded:

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <password>some password</password>

```

```

    <add-attr name="fname">
      <value>John</value>
    </add-attr>
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <jdbc:sql>CREATE USER jdoe IDENTIFIED BY
      {$$password}</jdbc:sql>
  </jdbc:statement>
</input>

```

Furthermore, you can also refer to the driver's database authentication password specified by the `Application Password` parameter as `{$$$driver-password}`. See [“Application Password” on page 58](#). Named password substitution is not yet supported.

Just as with association elements, the referenced password element must precede the `<jdbc:statement>` element in the XDS document or must be present as a child of the `<jdbc:statement>` element. Alternatively, instead of copying the password element as child of the `<jdbc:statement>` element, you could copy the `src-entry-id` of the element containing the password element onto the `<jdbc:statement>` element. Both approaches are bolded in the following examples:

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <password>some password</password>
    <add-attr name="fname">
      <value>John</value>
    </add-attr>
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <password>some password</password>
    <jdbc:sql>CREATE USER jdoe IDENTIFIED BY
      {$$password}</jdbc:sql>
  </jdbc:statement>
</input>

```

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr" src-entry-id="0">
    <password>some password</password>
    <add-attr name="fname">
      <value>John</value>
    </add-attr>
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement src-entry-id="0">
    <jdbc:sql>CREATE USER jdoe IDENTIFIED BY
      {$$password}</jdbc:sql>

```

```

    </jdbc:statement>
</input>

```

13.4 Virtual Triggers

In the same way that database triggers can fire before or after a triggering statement, embedded SQL can be positioned before or after the triggering XDS event. The following examples show how you can embed SQL before or after an XDS event.

Virtual Before Trigger

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <jdbc:statement>
    <association>idu=1,table=usr,schema=indirect</association>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
      idu = {$idu}</jdbc:SQL>
  </jdbc:statement>
  <modify class-name="usr">
    <association>idu=1,table=usr,schema=indirect</association>
    <modify-attr name="lname">
      <remove-all-values/>
      <add-value>
        <value>Doe</value>
      </add-value>
    </modify-attr>
  </modify>
</input>

```

This XML resolves to:

```

SET AUTOCOMMIT OFF
UPDATE indirect.usr SET fname = 'John' WHERE idu = 1;
COMMIT; --explicit commit
UPDATE indirect.usr SET lname = 'Doe' WHERE idu = 1;
COMMIT; --explicit commit

```

Virtual After Trigger

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <modify class-name="usr">
    <association>idu=1,table=usr,schema=indirect</association>
    <modify-attr name="lname">
      <remove-all-values/>
      <add-value>
        <value>Doe</value>
      </add-value>
    </modify-attr>
  </modify>
  <jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
      idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>

```

This XML resolves to:

```
SET AUTOCOMMIT OFF
UPDATE indirect.usr SET lname = 'Doe' WHERE idu = 1;
COMMIT; --explicit commit
UPDATE indirect.usr SET fname = 'John' WHERE idu = 1;
COMMIT; --explicit commit
```

13.5 Manual vs. Automatic Transactions

You can manually group embedded SQL and XDS events by using two custom attributes:

- ♦ `jdbc:transaction-type`
- ♦ `jdbc:transaction-id`

`jdbc:transaction-type`

This attribute has two values: `manual` and `auto`. By default, most XDS events of interest (`<add>`, `<modify>`, and `<delete>`) are implicitly set to the manual transaction type. The manual setting enables XDS events to resolve to a transaction consisting of one or more SQL statement.

By default, embedded SQL events are set to auto transaction type because some SQL statements, such as DDL statements, cannot usually be included in a manual transaction. In the following example, the attribute is in bold text.

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr" jdbc:transaction-type="auto">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
      idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>
```

This XML resolves to:

```
SET AUTOCOMMIT ON
INSERT INTO indirect.usr(lname) VALUES('Doe');
-- implicit commit
UPDATE indirect.usr SET fname = 'John' WHERE idu = 1;
-- implicit commit
```

`jdbc:transaction-id`

The Subscriber channel ignores this attribute unless the element's `jdbc:transaction-type` attribute value defaults to or is explicitly set to `manual`. The following XML shows an example of a manual transaction. The attribute is in bold text.

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr" jdbc:transaction-id="0">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
</input>
```

```

</add>
  <jdbc:statement jdbc:transaction-type="manual"
                jdbc:transaction-id="0">
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John' WHERE
            idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>

```

This XML resolves to:

```

SET AUTOCOMMIT OFF
INSERT INTO indirect.usr(lname) VALUES('Doe');
UPDATE indirect.usr SET fname = 'John' WHERE idu = 1;
COMMIT; -- explicit commit

```

13.6 Transaction Isolation Level

In addition to grouping statements, you can use transactions to preserve the integrity of data in a database. Transactions can lock data to prevent concurrent access or modification. The isolation level of a transaction determines how locks are set. Usually, the default isolation level that the driver uses is sufficient and should not be altered.

The custom attribute `jdbc:isolation-level` allows you to adjust the isolation transaction level if necessary. The `java.sql.Connection` parameter defines five possible values in the interface. See [java.sql.Connection](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html>).

- ◆ none
- ◆ read uncommitted
- ◆ read committed
- ◆ repeatable read
- ◆ serializable

The driver's default transaction isolation level is `read committed` unless overridden by a descriptor file. In manual transactions, place the `jdbc:isolation-level` attribute on the first element in the transaction. This attribute is ignored on subsequent elements. In the following example, the attribute is in bold text.

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr" jdbc:transaction-id="0"
      jdbc:isolation-level="serializable">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement jdbc:transaction-type="manual"
                jdbc:transaction-id="0">
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John'
            WHERE idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>

```

This XML resolves to:

```

SET AUTOCOMMIT OFF
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE

```



```

INSERT INTO indirect.usr(lname) VALUES('Doe');
UPDATE indirect.usr SET fname = 'John' WHERE idu = 1;
COMMIT; -- explicit commit

```

13.7 Statement Type

The Subscriber channel executes embedded SQL statements, but it doesn't understand them. The JDBC 1 interface defines several methods for executing different types of SQL statements. The following table contains these methods:

Table 13-1 *Methods for Executing SQL Statements*

Statement Type	Method Executed
SELECT	java.sql.Statement.executeQuery(String query):java.sql.ResultSet
INSERT	java.sql.Statement.executeUpdate(String update):int
UPDATE	java.sql.Statement.executeUpdate(String update):int
DELETE	java.sql.Statement.executeUpdate(String update):int
CALL or EXECUTE SELECT INSERT UPDATE DELETE	java.sql.Statement.execute(String sql):boolean

The simplest solution is to map all SQL statements to the `java.sql.Statement.execute(String sql):boolean` method. By default, the Subscriber channel uses this method.

Some third-party drivers, particularly Oracle's JDBC drivers, incorrectly implement the methods used to determine the number of result sets that this method generates. Consequently, the driver can get caught in an infinite loop leading to high CPU utilization. To circumvent this problem, you can use the `jdbc:type` attribute on any `<jdbc:statement>` element to map the SQL statements contained in it to the following methods instead of the default method:

- ◆ `java.sql.Statement.executeQuery(String query):java.sql.ResultSet`
- ◆ `java.sql.Statement.executeUpdate(String update):int`

The `jdbc:type` attribute has two values: `update` and `query`. For `INSERT`, `UPDATE`, or `DELETE` statements, set the value to `update`. For `SELECT` statements, set the value to `query`. In the absence of this attribute, the driver maps all SQL statements to the default method. If placed on any element other than `<jdbc:statement>`, this attribute is ignored.

Recommendations:

- ◆ Place the `jdbc:type="query"` attribute value on all `SELECT` statements.
- ◆ Place the `jdbc:type="update"` attribute value on all `INSERT`, `UPDATE`, and `DELETE` statements.
- ◆ Place no attribute value on stored procedure/function calls.

The following XML shows an example of the `jdbc:type` attribute. The attribute is in bold text.

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement jdbc:type="update">
    <jdbc:sql>UPDATE indirect.usr SET fname = 'John'
      WHERE idu = {$idu}</jdbc:sql>
  </jdbc:statement>
</input>
```

13.8 SQL Queries

To fully support the query capabilities of a database and avoid the difficulty of translating native SQL queries into an XDS format, the driver supports native SQL query processing. You can embed select statements in XDS documents in exactly the same way as any other SQL statement.

For example, assume that the table `usr` has the following contents:

Table 13-2 Example Contents

idu	fname	lname
1	John	Doe

The XML document below results in an output document containing a single result set.

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <jdbc:statement jdbc:type="query">
    <jdbc:sql>SELECT * FROM indirect.usr</jdbc:sql>
  </jdbc:statement>
</input>
<output xmlns:jdbc="urn:dirxml:jdbc">
  <jdbc:result-set jdbc:number-of-rows="1">
    <jdbc:row jdbc:number="1">
      <jdbc:column jdbc:name="idu"
        jdbc:position="1"
        jdbc:type="java.sql.Types.BIGINT"
        <jdbc:value>1</jdbc:value>
      </jdbc:column>
      <jdbc:column jdbc:name="fname"
        jdbc:position="2"
        jdbc:type="java.sql.Types.VARCHAR"
        <jdbc:value>John</jdbc:value>
      </jdbc:column>
      <jdbc:column jdbc:name="lname"
        jdbc:position="3"
        jdbc:type="java.sql.Types.VARCHAR"
        <jdbc:value>Doe</jdbc:value>
      </jdbc:column>
    </jdbc:row>
  </jdbc:result-set>
</output>
```

```

    </jdbc:result-set>
    <status level="success"/>
</output>

```

SQL queries always produce a single `<jdbc:result-set>` element whether or not the result set contains any rows. If the result set is empty, the `jdbc:number-of-rows` attribute is set to zero.

You can embed more than one query in a document. SQL queries don't require that the referenced tables/views in the synchronization schema be visible to the driver. However, XDS queries do.

13.9 Data Definition Language (DDL) Statements

Generally, it is not possible to run a Data Definition Language (DDL) statement in a database trigger because most databases do not allow mixed DML and DDL transactions. Although virtual triggers do not overcome this transactional limitation, they do allow DDL statements to be executed as a side effect of an XDS event.

For example:

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr">
    <add-attr name="fname">
      <value>John</value>
    </add-attr>
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
  </add>
  <jdbc:statement>
    <jdbc:sql>CREATE USER jdoe IDENTIFIED BY novell</jdbc:sql>
  </jdbc:statement>
</input>

```

This XML resolves to:

```

SET AUTOCOMMIT OFF
INSERT INTO indirect.usr(fname, lname) VALUES('John', 'Doe');
COMMIT; -- explicit commit
SET AUTOCOMMIT ON
CREATE USER jdoe IDENTIFIED BY novell;
-- implicit commit

```

Using the `jdbc:transaction-id` and `jdbc:transaction-type` attributes to group DML and DDL statements into a single transaction causes the transaction to be rolled back on most databases. Because DDL statements are generally executed as separate transactions, it is possible that the `insert` statement in the above example might succeed and the `create user` statement might roll back.

It is not possible, however, that the `insert` statement fail and the `create user` statement succeed. The driver stops executing chained transactions at the point where the first transaction is rolled back.

13.10 Logical Operations

Because it is not generally possible to mix DML and DDL statements in a single transaction, a single event can consist of one or more transactions. You can use the `jdbc:op-id` and `jdbc:op-type` to group multiple transactions together into a single logical operation. When so grouped, all members of the operation are handled as a single unit with regard to status. If one member has an error, all members return the same status level. Similarly, all members share the same status type.

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr" jdbc:op-id="0"
      jdbc:op-type="password-set-operation">
    <add-attr name="fname">
      <value>John</value>
    </add-attr>
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
    <password>Doe{$idu}</password>
  </add>
  <jdbc:statement jdbc:op-id="0">
    <jdbc:sql>CREATE USER jdoe IDENTIFIED BY {$$password}</jdbc:sql>
  </jdbc:statement>
</input>
```

The `jdbc:op-type` attribute is ignored on all elements except the first element in a logical operation.

13.11 Implementing Password Set with Embedded SQL

Initially setting a password is usually accomplished by creating a database user account. Assuming that an `<add>` event is generated on the Subscriber channel, the following is an example of the output generated by XSLT style sheets that implement password set as a side effect of an XDS `<add>` event:

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <add class-name="usr" jdbc:op-id="0"
      jdbc:op-type="password-set-operation">
    <add-attr name="fname">
      <value>John</value>
    </add-attr>
    <add-attr name="lname">
      <value>Doe</value>
    </add-attr>
    <password>Doe{$idu}</password>
  </add>
  <jdbc:statement jdbc:op-id="0">
    <jdbc:sql>CREATE USER jdoe IDENTIFIED BY {$$password}</jdbc:sql>
  </jdbc:statement>
</input>
```

The `<add>` event is logically bound to the `CREATE USER` DDL statement by the `jdbc:op-id` and `jdbc:op-type` attributes.

The `User DDL Command Transformation` style sheet in the `example.xml` configuration file contains sample XSLT templates that bind user account creation DDL statements to `<add>` events for all databases that support them.

13.12 Implementing Modify Password with Embedded SQL

Modifying a password is usually accomplished by altering an existing database user account. Assuming that a `<modify-password>` event is generated on the `Subscriber` channel, the following is an example of the output generated by XSLT style sheets that implement `modify-password`:

NOTE: Some databases, such as Sybase Adaptive Server Enterprise and Microsoft SQL Server, differentiate between user account names and login account names. Therefore, you might need to supply the login name instead of the user name.

```
<input xmlns:jdbc="urn:dirxml:jdbc">
  <modify-password jdbc:op-id="0"
                  jdbc:op-type="password-set-operation">
    <password>new password</password>
  </modify-password>
  <jdbc:statement jdbc:op-id="0">
    <jdbc:sql>ALTER USER jdoe IDENTIFIED BY {$$password}</jdbc:sql>
  </jdbc:statement>
</input>
```

The `<modify-password>` event is logically bound to the `ALTER USER` DDL statement by the `jdbc:op-id` and `jdbc:op-type` attributes.

The `User DDL Command Transformation` style sheet in the `example.xml` configuration file contains sample XSLT templates that bind password maintenance DDL statements to `<modify-password>` events for all databases that support them.

13.13 Implementing Check Object Password

Unlike password set, check object password does not require embedded SQL statements or attributes. Only a user account name is required. This could be obtained from an association value (assuming that associations are being maintained manually), a directory attribute, or a database field. If stored in the directory or database, a query must be issued to retrieve the value.

The `example.xml` configuration file stores database user account names in database fields.

NOTE: Some databases, such as Sybase Adaptive Server Enterprise and Microsoft SQL Server, differentiate between user account names and login account names. Therefore, you might need to store two names, not just one.

To implement check object password, append a `dest-dn` attribute value to the `<check-object-password>` event. In the following example, the `dest-dn` attribute is bolded:

```

<input xmlns:jdbc="urn:dirxml:jdbc">
  <check-object-password dest-dn="jdoe">
    <password>whatever</password>
  </check-object-password>
</input>

```

13.14 Calling Stored Procedures and Functions

The Identity Manager Driver for JDBC 3.5.1 enables you to use stored procedures. The ability to use the `<jdbc:call-procedure>` and `<jdbc:call-function>` elements to call stored procedures from a policy has been tested against only Oracle and is supported only on that platform.

- [Section 13.14.1, “Using Embedded SQL to Call Stored Procedures or Functions,” on page 166](#)
- [Section 13.14.2, “Using the jdbc:call-procedure Element,” on page 167](#)
- [Section 13.14.3, “Using the jdbc:call-function Element,” on page 170](#)

13.14.1 Using Embedded SQL to Call Stored Procedures or Functions

You can call stored procedures or functions in one of two ways:

- Call the procedure or function by using a Statement object.
- Call the procedure by using a Callable Statement object.

Example 1: Calling a Stored Procedure by Using a Statement

```

<!-- call syntax is Oracle -->
<jdbc:statement>
  <jdbc:sql>CALL schema.procedure-name</jdbc:sql/>
</jdbc:statement>

```

Example 2: Calling a Stored Procedure as a CallableStatement

```

<!-- call syntax is vendor agnostic -->
<jdbc:statement>
  <jdbc:call-procedure jdbc:name="schema.procedure-name"/>
</jdbc:statement>

```

Example 3: Calling a Function by Using a Statement

```

<!-- call syntax is Informix -->
<jdbc:statement>
  <jdbc:sql>EXECUTE FUNCTION schema.function-name</jdbc:sql/>
</jdbc:statement>

```

Example 4: Calling a Function as a CallableStatement

```

<!-- call syntax is vendor agnostic -->
<jdbc:statement>
  <jdbc:call-function jdbc:name="schema.function-name"/>
</jdbc:statement>

```

The principle advantage of using the CallableStatement interface is that you do not need to know the proprietary call syntaxes of each database vendor or JDBC implementation. Other advantages include the following:

- ◆ It's much easier to build procedure or function calls in the Policy Builder.
- ◆ You can differentiate between Null and empty string parameter values.
- ◆ You can call functions on all database platforms.
Oracle, for instance, doesn't support calling functions by using a statement.
- ◆ You can retrieve Out parameter values from stored procedure calls.

13.14.2 Using the jdbc:call-procedure Element

Stored procedures do not necessarily require parameters. Only a name is required. If a database supports schemas, we recommend that you schema-qualify the name. If a schema qualifier isn't provided, how names are resolved depends upon your third-party JDBC implementation and might change, depending upon driver configuration settings.

NOTE: The jdbc:call-procedure element must be wrapped in a jdbc:statement element.

- ◆ [“Specifying a Procedure Name” on page 167](#)
- ◆ [“Passing In or In Out Parameter Values” on page 167](#)
- ◆ [“Handling Out or In Out Parameters” on page 168](#)
- ◆ [“Example Complex Stored Procedure Calls” on page 169](#)

Specifying a Procedure Name

```
<jdbc:call-procedure jdbc:name="schema.procedure-name"/>
```

Passing In or In Out Parameter Values

The number of jdbc:param elements specified must match the number of param elements declared in the procedure. Only jdbc:param elements corresponding to In or In Out procedure parameters can have values. Out parameters (those that can't be passed values) must be represented by an empty jdbc:param element.

Calling a Procedure with No Parameters

```
<jdbc:statement>  
  <jdbc:call-procedure jdbc:name="schema.procedure-name"/>  
</jdbc:statement>
```

Calling a Procedure with a Null Parameter

```
<jdbc:call-procedure jdbc:name="schema.procedure-name">  
  <!-- no value element = pass null -->  
  <jdbc:param/>  
</jdbc:call-procedure>
```

Calling a Procedure with an Empty String Parameter

```
<jdbc:call-procedure jdbc:name="schema.procedure-name">  
  <!-- empty value element = pass empty string -->
```

```

    <jdbc:param>
      <jdbc:value/>
    </jdbc:param>
  </jdbc:param>

```

NOTE: Literals can be passed only to procedure parameters declared as In or In Out. Passed literals must be type-compatible with declared procedure parameters.

Calling a Procedure with a Literal Value

```

<jdbc:call-procedure jdbc:name="schema.procedure-name">
  <!-- non-empty value element = pass literal -->
  <jdbc:param>
    <jdbc:value>literal</jdbc:value>
  </jdbc:param>
</jdbc:param>

```

Calling a Procedure with an Out Parameter

Assuming that a procedure has two parameters, the first Out and the second In, you invoke the procedure as follows:

```

<jdbc:call-procedure jdbc:name="schema.procedure-name">
  <!-- the OUT param place -->
  <jdbc:param/>
  <!-- the IN param -->
  <jdbc:param>
    <jdbc:value>literal</jdbc:value>
  </jdbc:param>
</jdbc:param>

```

Handling Out or In Out Parameters

Stored procedures with Out or In Out parameters can return values. These values are returned by the driver and are accessible to policies. Out or In Out parameters values are returned at the same position as their corresponding declared parameter.

Also, to facilitate correlation of procedure calls and output parameter values, Out parameters contain the same event-ID value as the procedure call that generated them. This is particularly useful when multiple calls are made in the same document.

Null or No Return Value

Assuming that a procedure has a single Out or In parameter, the following output is generated:

```

<output>
  <!-- no value element = OUT param returned null or IN param -->
  <jdbc:out-parameters event-id="0" jdbc:number-of-params="1">
    <jdbc:param/>
  </jdbc:out-parameters>
  <status event-id="0" level="success"/>
</output>

```

Empty String Return Value

Assuming that a procedure has a single Out or In Out parameter, the following output is generated:


```

<output>
  <!-- empty value element = returned empty string -->
  <jdbc:out-parameters event-id="0" jdbc:number-of-params="1">
    <jdbc:param>
      <jdbc:value/>
    </jdbc:param>
  </jdbc:out-parameters>
  <status event-id="0" level="success"/>
</output>

```

Literal Return Value

Assuming that a procedure has a single Out or In Out parameter, the following output is generated:

```

<output>
  <!-- no-empty value element = returned literal value -->
  <jdbc:out-parameters event-id="0" jdbc:number-of-params="2">
    <jdbc:param>
      <jdbc:value>literal<jdbc:value>
    </jdbc:param>
  </jdbc:out-parameters>
  <status event-id="0" level="success"/>
</output>

```

Example Complex Stored Procedure Calls

- ◆ [“Procedure Declaration” on page 169](#)
- ◆ [“Procedure Call from Policy” on page 169](#)
- ◆ [“Procedure Output to Policy” on page 170](#)

Procedure Declaration

NOTE: This procedure uses Oracle PSQL syntax.

```

CREATE PROCEDURE indirect.p1(i1 IN VARCHAR2, io2 IN OUT VARCHAR2, o3
OUT INTEGER, i4 IN VARCHAR2)
AS
BEGIN
  SELECT 'literal' INTO io2 FROM DUAL;
  SELECT 1 INTO o3 FROM DUAL;
END p1;

```

Procedure Call from Policy

```

<input>
  <jdbc:statement event-id="0">
    <jdbc:call-procedure jdbc:name="indirect.p1">
      <!-- i1 IN VARCHAR2 -->
      <jdbc:param>
        <!-- pass empty string -->
        <jdbc:value/>
      </jdbc:param>
      <!-- io2 IN OUT VARCHAR2 -->
      <jdbc:param>
        <!-- pass literal -->

```

```

        <jdbc:value>literal</jdbc:value>
    </jdbc:param>
    <!-- o3 OUT INTEGER -->
    <!-- param placeholder -->
    <jdbc:param/>
    <!-- o4 IN VARCHAR2 -->
    <!-- pass null -->
    <jdbc:param/>
    </jdbc:call-procedure>
</jdbc:statement>
</input>

```

Procedure Output to Policy

```

<output>
  <jdbc:out-parameters event-id="0" jdbc:number-of-params="2">
    <jdbc:param/>
    <jdbc:param jdbc:name="IO2"
      jdbc:param-type="INOUT"
      jdbc:position="2"
      jdbc:sql-type="java.sql.Types.VARCHAR">
      <jdbc:value>literal</jdbc:value>
    </jdbc:param>
    <jdbc:param jdbc:name="O3"
      jdbc:param-type="OUT"
      jdbc:position="3"
      jdbc:sql-type="java.sql.Types.DECIMAL">
      <jdbc:value>1</jdbc:value>
    </jdbc:param>
    <jdbc:param/>
  </jdbc:out-parameters>
  <status event-id="0" level="success"/>
</output>

```

13.14.3 Using the jdbc:call-function Element

Functions do not necessarily require parameters. Only a name is required. If a database supports schemas, we recommend that you schema-qualify the name. If a schema qualifier isn't provided, how names are resolved depends upon your third-party JDBC implementation and might change depending upon driver configuration settings.

NOTE: The `jdbc:call-function` element must be wrapped in a `jdbc:statement` element.

- ◆ [“Specifying a Function Name” on page 171](#)
- ◆ [“Passing In Parameter Values” on page 171](#)
- ◆ [“Calling a Function with No Parameter” on page 171](#)
- ◆ [“Calling a Function with a Null Parameter” on page 171](#)
- ◆ [“Calling a Function with an Empty String Parameter” on page 171](#)
- ◆ [“Calling a Function with a Literal Value” on page 171](#)
- ◆ [“Handling Function Results” on page 171](#)
- ◆ [“Example Complex Function Calls” on page 173](#)

Specifying a Function Name

```
<jdbc:call-function jdbc:name="schema.function-name"/>
```

Passing In Parameter Values

The number of jdbc:param elements specified must match the number of params declared in the function.

Calling a Function with No Parameter

```
<jdbc:call-function jdbc:name="schema.function-name"/>
```

Calling a Function with a Null Parameter

```
<jdbc:call-function jdbc:name="schema.function-name">  
  <!-- no value element = null -->  
  <jdbc:param/>  
</jdbc:call-procedure>
```

Calling a Function with an Empty String Parameter

```
<jdbc:call-function jdbc:name="schema.function-name">  
  <!-- empty value element = pass empty string -->  
  <jdbc:param>  
    <jdbc:value/>  
  </jdbc:param>  
<jdbc:param>
```

NOTE: Literals can be passed to function parameters declared as In. Passed literals must be type-compatible with declared function parameters.

Calling a Function with a Literal Value

```
<jdbc:call-function jdbc:name="schema.function-name">  
  <!-- non-empty value element = pass literal -->  
  <jdbc:param>  
    <jdbc:value>literal</jdbc:value>  
  </jdbc:param>  
<jdbc:param>
```

Handling Function Results

Unlike stored procedures, functions do not support Out or In Out parameters. They can, however, return a single, scalar value (such as an integer or string) or return a result set. Also, to facilitate correlation of function calls and results, results contain the same event-id value as the function call that generated them. This is particularly useful when multiple calls are made in the same document.

Scalar Return Value

Scalar return values are returned using the same syntax as stored procedure Out parameters. The scalar return value is always returned in the first parameter position.

```
<output>  
  <jdbc:out-parameters event-id="0" jdbc:number-of-params="1">  
    <jdbc:param jdbc:name="return value"
```

```

        jdbc:param-type="OUT"
        jdbc:position="1"
        jdbc:sql-type="java.sql.Types.VARCHAR">
        <jdbc:value>1</jdbc:value>
    </jdbc:param>
</jdbc:out-parameters>
<status event-id="0" level="success"/>
</output>

```

Empty Set

Assuming that a function returns no results set or an empty result set, the following output is generated:

```

<output>
    <jdbc:result-set event-id="0" jdbc:number-of-rows="0"/>
    <status event-id="0" level="success"/>
</output>

```

Non-Empty Results Set

Assuming a function returns a non-empty result set, output similar to the following is generated:

```

<output>
    <jdbc:result-set event-id="0" jdbc:number-of-rows="1">
        <jdbc:row jdbc:number="1">
            <jdbc:column jdbc:name="SYSDATE"
                jdbc:position="1"
                jdbc:type="java.sql.Types.TIMESTAMP">
                <jdbc:value>2007-01-03 14:52:20.0</jdbc:value>
            </jdbc:column>
        </jdbc:row>
    </jdbc:result-set>
    <status event-id="0" level="success"/>
</output>

```

Multiple Result Sets

Multiple result sets are returned in the order returned by the function. They all share a common event-id value.

```

<output>
    <jdbc:result-set event-id="0" jdbc:number-of-rows="0"/>
    <jdbc:result-set event-id="0" jdbc:number-of-rows="0"/>
    <status event-id="0" level="success"/>
</output>

```

Oracle Results Set

Oracle's JDBC implementation uses a proprietary mechanism to return a single result set from a function. To return a result set from an Oracle function, you need to explicitly set the **jdbc:return-type** value to `OracleTypes.CURSOR` on the `jdbc:call-function` element.

Returning Result Sets as Out Parameters

See the special attribute **“jdbc:return-format”** on page 173.

Special Attributes

jdbc:return-format

This attribute can be placed on the `jdbc:call-function` element to format the first row of a returned results set as stored procedure Out parameters of the result.

NOTE: This works only when the `jdbc:return-type` attribute isn't used.

```
<input>
  <jdbc:statement>
    <jdbc:call-function jdbc:name="schema.function-name"
      jdbc:return-format="return value">
    </jdbc:call-function>
  </jdbc:statement>
</input>
```

jdbc:return-type

This attribute can be placed on the `jdbc:call-function` element to allow Oracle functions to return a result set.

```
<input>
  <jdbc:statement>
    <jdbc:call-function jdbc:name="schema.function"
      jdbc:return-type="OracleTypes.CURSOR">
    </jdbc:call-function>
  </jdbc:statement>
</input>
```

Example Complex Function Calls

- ◆ [“Function Declaration” on page 173](#)
- ◆ [“Function Call from a Policy” on page 173](#)
- ◆ [“Function Results to a Policy” on page 174](#)

Function Declaration

NOTE: This declaration is for Oracle PSQL syntax.

```
CREATE OR REPLACE FUNCTION indirect.f1(i1 IN VARCHAR2, i2 IN INTEGER)
  RETURN VARCHAR2
AS
  o_idu VARCHAR2(32);
BEGIN
  SELECT 'literal' INTO o_idu FROM DUAL;
  RETURN o_idu;
END f1;
```

Function Call from a Policy

```
<input>
  <jdbc:statement>
    <jdbc:call-function jdbc:name="indirect.f1">
      <jdbc:param>
```

```

        <jdbc:value>literal</jdbc:value>
    </jdbc:param>
    <jdbc:param>
        <jdbc:value>1</jdbc:value>
    </jdbc:param>
</jdbc:call-function>
</jdbc:statement>
</input>

```

Function Results to a Policy

```

<output>
    <jdbc:out-parameters event-id="0" jdbc:number-of-params="1">
        <jdbc:param jdbc:name="return value"
            jdbc:param-type="OUT"
            jdbc:position="1"
            jdbc:sql-type="java.sql.Types.VARCHAR">
            <jdbc:value>literal</jdbc:value>
        </jdbc:param>
    </jdbc:out-parameters>
    <status event-id="0" level="success"/>
</output>

```

13.15 Best Practices

For performance reasons, it is better to call a single stored procedure/function that contains multiple SQL statements than to embed multiple statements in an XDS document.

In the following examples, the single stored procedure or function is preferred.

Single Stored Procedure

```

<input xmlns:jdbc="urn:dirxml:jdbc">
    <add class-name="usr">
        <add-attr name="fname">
            <value>John</value>
        </add-attr>
        <add-attr name="lname">
            <value>Doe</value>
        </add-attr>
    </add>
    <jdbc:statement>
        <jdbc:sql>CALL PROCEDURE set_name('John', 'Doe')</jdbc:sql>
    </jdbc:statement>
</input>

```

Multiple Embedded Statements

```

<input xmlns:jdbc="urn:dirxml:jdbc">
    <add class-name="usr">
        <add-attr name="lname">
            <value>Doe</value>
        </add-attr>
    </add>
    <jdbc:statement>
        <jdbc:sql>UPDATE indirect.usr SET fname = 'John'

```

```
        WHERE idu = {$idu}</jdbc:sql>
</jdbc:statement>
<jdbc:statement>
    <jdbc:sql>UPDATE indirect.usr SET lname = 'Doe'
        WHERE idu = {$idu}</jdbc:sql>
</jdbc:statement>
</input>
```

The syntax used to call stored procedures or functions varies by database. For additional information, see [“Syntaxes for Calling Stored Procedures and Functions”](#) on page 180.

Supported Databases

14

- ♦ [Section 14.1, “Database Interoperability,” on page 177](#)
- ♦ [Section 14.2, “Supported Databases,” on page 177](#)
- ♦ [Section 14.3, “Database Characteristics,” on page 178](#)

14.1 Database Interoperability

The Identity Manager Driver for JDBC is designed to interoperate with a specific set of JDBC driver implementations, instead of a specific set of databases. Consequently, the list of supported databases is primarily driven by the capabilities of supported third-party JDBC drivers. A secondary factor is testing resources.

14.2 Supported Databases

The following databases or database versions have been tested and are recommended for use with this product:

Table 14-1 *Supported Databases*

Database	Minor Version
IBM DB2 Universal Database (UDB) 7	7.2 or later
IBM DB2 Universal Database (UDB) 8	8.1 or later
Informix Dynamic Server (IDS) 9	9.40 or later
Informix Dynamic Server (IDS) 10	10.0 or later
Microsoft SQL Server 7	7.5, Service Pack 4 or later
Microsoft SQL Server 2000 (8)	Service Pack 3a or later
Microsoft SQL Server 2005 (9)	Service Pack 1 or later
MySQL 3	3.23.58 or later
MySQL 4	4.1 or later
MySQL 5	5.0.20 or later
Oracle 8i	Release 3 (8.1.7) or later
Oracle 9i	Release 2 (9.2.0.1) or later
Oracle 10g	Release 1 (10.0.2.1) or later
PostgreSQL 7	7.4.6 or later
Sybase Adaptive Server Enterprise (ASE) 12	12.5 or later

You can use the Driver for JDBC with other databases or database versions. However, Novell® does not support them. To interoperate with the Driver for JDBC, a database must meet the following requirements:

- ◆ Support the SQL-92 entry level grammar.
- ◆ Be JDBC-accessible.

14.3 Database Characteristics

- ◆ [Section 14.3.1, “Database Features,” on page 178](#)
- ◆ [Section 14.3.2, “Current Time Stamp Statements,” on page 179](#)
- ◆ [Section 14.3.3, “Syntaxes for Calling Stored Procedures and Functions,” on page 180](#)
- ◆ [Section 14.3.4, “Left Outer Join Operators,” on page 180](#)
- ◆ [Section 14.3.5, “Undelimited Identifier Case Sensitivity,” on page 181](#)
- ◆ [Section 14.3.6, “Supported Transaction Isolation Levels,” on page 181](#)
- ◆ [Section 14.3.7, “Commit Keywords,” on page 182](#)
- ◆ [Section 14.3.8, “IBM DB2 Universal Database \(UDB\),” on page 182](#)
- ◆ [Section 14.3.9, “Informix Dynamic Server \(IDS\),” on page 183](#)
- ◆ [Section 14.3.10, “Microsoft SQL Server,” on page 184](#)
- ◆ [Section 14.3.11, “MySQL,” on page 184](#)
- ◆ [Section 14.3.12, “Oracle,” on page 185](#)
- ◆ [Section 14.3.13, “PostgreSQL,” on page 186](#)
- ◆ [Section 14.3.14, “Sybase Adaptive Server Enterprise \(ASE\),” on page 186](#)

14.3.1 Database Features

Table 14-2 Database Features

Database	Schemas	Views	Identity Columns	Sequences	Stored Procedures	Functions	Triggers	Instead-Of-Triggers
IBM DB2 UDB 7	X	X	X	0	X ¹	X ¹	X	0
IBM DB2 UDB 8	X	X	X	0	X ¹	X ¹	X	X
Informix IDS 9	X	X	X ²	0	X ³	X	X	0
Informix IDS 10	X	X	X ²	0	X ³	X	X	X
MS SQL 7	X	X	X	0	X	0	X	0
MS SQL 2000, 2005	X	X	X	0	X	X	X	X

Database	Schemas	Views	Identity Columns	Sequences	Stored Procedures	Functions	Triggers	Instead-Of-Triggers
MySQL 4	0	0	X ⁴	0	0	0	0	0
MySQL 5	X	X	X ⁴	0	X	X	X	0
Oracle 8i, 9i, 10g	X	X	0	X	X	X	X	X
Postgres 7	X	X	X ⁵	X	X	X	X ⁶	X ⁶
Sybase ASE 12	X	X	X	0	X	0	X	0

¹ DB2 natively supports stored procedures or functions written in Java. To write procedures by using the native SQL procedural language, install a C compiler on the database server.

² The Informix identity column keyword is `SERIAL8`.

³ Informix stored procedures cannot return values through OUT parameters.

⁴ The MySQL identity column keyword is `AUTO_INCREMENT`.

⁵ You can use a Postgres sequence object to provide default values for primary key columns, effectively simulating an identity column.

Postgres has a native construct called rules. This construct can be used to effectively simulate triggers and instead-of-triggers. It also supports the use of triggers or instead-of-triggers written in a variety of procedural programming languages.

14.3.2 Current Time Stamp Statements

The following table lists SQL statements used to retrieve the current date and time by database:

Table 14-3 *Time Stamp Statements*

Database	Current Time Stamp Statement	ANSI-Compliant
IBM DB2 UDB	<code>SELECT (CURRENT_TIMESTAMP) FROM SYSIBM.SYSDUMMY1 FETCH FIRST 1 ROW ONLY</code>	No
Informix IDS	<code>SELECT FIRST 1 (CURRENT YEAR TO FRACTION(5)) FROM INFORMIX.SYSTABLES</code>	No
MSSQL	<code>SELECT (CURRENT_TIMESTAMP)</code>	Yes
MySQL	<code>SELECT (CURRENT_TIMESTAMP)</code>	Yes
Oracle	<code>SELECT (SYSDATE) FROM SYS.DUAL</code>	No
PostgreSQL	<code>SELECT (CURRENT_TIMESTAMP)</code>	Yes
Sybase ASE	<code>SELECT GETDATE()</code>	No

14.3.3 Syntaxes for Calling Stored Procedures and Functions

The following table lists the syntaxes for calling a stored procedure or function by database vendor. There's also a vendor-neutral JDBC escape syntax (see [JDBC Escape Syntax \(http://edocs.bea.com/wls/docs81/jdbc_drivers/sqlscape.html\)](http://edocs.bea.com/wls/docs81/jdbc_drivers/sqlscape.html)). Whenever possible, it is more secure to call a stored procedure or function by using the `jdbc:call-function` or `jdbc:call-procedure` syntax. (See [Section 13.14, "Calling Stored Procedures and Functions," on page 166.](#)) Other syntaxes should be used only when specifying procedure or function calls in driver parameters (for example, ["Post Polling Statements" on page 97](#) and ["Connection Initialization Statements" on page 69](#)).

Table 14-4 *Calling a Stored Procedure or Function*

Database	Stored Procedure/Function JDBC Call Syntax
IBM DB2 UDB	{call <i>schema-name.procedure-name</i> (<i>parameter-list</i>)}
Informix IDS	EXECUTE [PROCEDURE FUNCTION] <i>schema-name.routine-name</i> (<i>parameter-list</i>)
MSSQL	EXECUTE <i>schema-name.procedure-name</i> (<i>parameter-list</i>)
MySQL	[TODO]
Oracle ¹	CALL <i>schema-name.procedure-name</i> (<i>parameter-list</i>)
PostgreSQL	SELECT <i>schema-name.procedure-name</i> (<i>parameter-list</i>)
Sybase ASE	EXECUTE <i>schema-name.procedure-name</i> (<i>parameter-list</i>)

¹ Oracle's JDBC implementation does not support calling functions as a string.

14.3.4 Left Outer Join Operators

The following table lists outer join operators by database.

Table 14-5 *Outer Join Operators*

Database	Left Outer Join Operator	ANSI-Compliant
IBM DB2 UDB	LEFT OUTER JOIN	Yes
Informix IDS	LEFT OUTER JOIN	Yes
MSSQL 7.5, 2000	*=	No
MSSQL 2005	LEFT OUTER JOIN	Yes
MySQL	LEFT OUTER JOIN	Yes
Oracle	(+) [TODO]	No
PostgreSQL	LEFT OUTER JOIN	Yes
Sybase ASE	*=	No

NOTE: Oracle supports the ANSI-compliant left outer join operator LEFT OUTER JOIN as of version 10g.

14.3.5 Undelimited Identifier Case Sensitivity

Table 14-6 Case Sensitivity for Undelimited Identifiers

Database	Case-Sensitive?
IBM DB2 UDB	No
Informix IDS	No
MSSQL	No
MySQL	Yes
Oracle	No
PostgreSQL	No
Sybase ASE	Yes

14.3.6 Supported Transaction Isolation Levels

Table 14-7 Supported Transaction Isolation Levels

Database	No ne	Read Uncommit ted	Read Commit ted	Repeata ble Read	Serializa ble	URL
IBM DB2 UDB	0	X	X ¹	X	X	Setting JDBC Transaction Isolation Levels (http://publib.boulder.ibm.com/infocenter/db2help/index.jsp?topic=/com.ibm.db2.udb.doc/ad/tvjdiso.htm)
MySQL (InnoDB Table Type)	0	X	X	X ¹	X	InnoDB Transaction Isolation Levels (http://dev.mysql.com/doc/mysql/en/innodb-transaction-isolation.html)
Oracle	0	0	X ¹	0	X	JDBC Transaction Optimization (http://www.oracle.com/technology/oramag/oracle/02-jul/o42special_jdbc.html)
PostgreSQL	0	0 ²	X ¹	0 ²	X	Transaction Isolation (http://www.postgresql.org/docs/current/static/transaction-iso.html)

¹ This is the default isolation level for this database.

² Can be set, but it is aliased to a supported isolation level.

14.3.7 Commit Keywords

The following table identifies the commit keywords for supported databases:

Table 14-8 *Commit Keywords*

Database	Commit Keyword
IBM DB2 UDB	COMMIT
Informix IDS	COMMIT WORK ¹
MSSQL	GO
MySQL	COMMIT
Oracle	COMMIT
PostgreSQL	COMMIT
Sybase ASE	GO

¹ For logging and ANSI-compliant databases. Non-logging databases do not support transactions.

14.3.8 IBM DB2 Universal Database (UDB)

The following table lists properties for this database.

Table 14-9 *Properties for IBM DB2 UDB*

Property	Value
Current Timestamp Statement	SELECT (CURRENT TIMESTAMP) FROM SYSIBM.SYSDUMMY1 FETCH FIRST 1 ROW ONLY
Case-Sensitive?	No
Commit Keyword	COMMIT
Left Outer Join Operator	LEFT OUTER JOIN

Dynamic Defaults

The following table lists database compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly override these settings.

Table 14-10 *Dynamically Configured IBM DB2 Universal Database Settings*

Display Name	Tag Name	Value
Current Timestamp Statement:	current-timestamp-stmt	SELECT (CURRENT TIMESTAMP) FROM SYSIBM.SYSDUMMY1 FETCH FIRST 1 ROW ONLY
Timestamp Translator class:	time-translator-class	com.novell.nds.dirxml.driver.jdbc.db.DB2Timestamp

Known Issues

- ♦ The timestamp format is proprietary.

See [“Known Issues” on page 187](#).

14.3.9 Informix Dynamic Server (IDS)

The following table lists properties for this database.

Table 14-11 *Settings for Informix Dynamic Server*

Property	Value
Current Timestamp Statement	SELECT FIRST 1 (CURRENT YEAR TO FRACTION(5)) FROM INFORMIX.SYSTABLES
Case-Sensitive?	No
Commit Keyword	COMMIT WORK ¹
Left Outer Join Operator	LEFT OUTER JOIN

¹ For logging and ANSI-compliant databases. Nonlogging databases do not support transactions.

Dynamic Defaults

The following table lists database compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly overwrite these settings.

Table 14-12 *Dynamically Configured Informix Dynamic Server Settings*

Display Name	Tag Name	Value
Current Timestamp Statement:	current-timestamp-stmt	SELECT FIRST 1 (CURRENT YEAR TO FRACTION(5)) FROM INFORMIX.SYSTABLES

Known Issues

- ♦ NUMERIC or DECIMAL columns cannot be used as primary keys unless the scale (the number of digits to the right of the decimal point) is explicitly set to 0 when the table is created. By default, the scale is set to 255.

- ◆ DBAs cannot grant privileges to objects they don't own.

14.3.10 Microsoft SQL Server

The following table lists properties for this database:

Table 14-13 *Settings for Microsoft SQL Server*

Property	Value
Current Timestamp Statement	SELECT (CURRENT_TIMESTAMP)
Case-Sensitive?	No
Commit Keyword	GO
Left Outer Join Operator (7, 2000)	*=
Left Outer Join Operator (2005)	LEFT OUTER JOIN

Dynamic Defaults

The following table lists database compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly overwrite these settings.

Table 14-14 *Dynamically Configured Microsoft SQL Server Settings*

Display Name	Tag Name	Value
Add default values on insert?	add-default-values-on-view-insert	true
Left outer-join operator (7, 2000):	left-outer-join-operator	*=
Left outer-join operator (2005):	left-outer-join-operator	LEFT OUTER JOIN

14.3.11 MySQL

The following table lists properties for this database.

Table 14-15 *Settings for MySQL*

Property	Value
Current Timestamp Statement	SELECT (CURRENT_TIMESTAMP)
Case-Sensitive?	Yes
Commit Keyword	COMMIT
Left Outer Join Operator	LEFT OUTER JOIN

Dynamic Defaults

The following table lists database compatibility parameters that are dynamically configured at runtime for this database.

Table 14-16 *Dynamically Configured MySQL Settings*

Display Name	Tag Name	Value
Supports schemas in metadata retrieval?	supports-schemas-in-metadata-retrieval	false

Known Issues

- ◆ `TIMESTAMP` columns, when updated after being initially set to 0 or `NULL`, are always set to the current date and time. To compensate for this behavior, we recommend that you map Identity Vault Time and Timestamp syntaxes to `DATETIME` columns.

14.3.12 Oracle

The following table lists properties for this database:

Table 14-17 *Settings for Oracle*

Property	Value
Current Timestamp Statement	<code>SELECT (SYSDATE) FROM SYS.DUAL</code>
Case-Sensitive?	No
Commit Keyword	<code>COMMIT</code>
Left Outer Join Operator	<code>(+)</code>

Dynamic Defaults

The following table lists database compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly overwrite these settings.

Table 14-18 *Dynamically Configured Oracle Settings*

Display Name	Tag Name	Value
Left outer-join operator	left-outer-join-operator	<code>(+)</code>
Exclude filter expression	exclude-table-filter	<code>BIN\\$.{22}==\\$0</code>
Lock statement generator class	lock-generator-class	<code>com.novell.nds.dirxml.driver.jdbc.db.lock.OraLockGenerator</code>

NOTE: The default exclusion filter is intended to omit from the synchronization schema dropped tables that are visible in Oracle 10g.

Limitations

- ◆ LONG, LONG RAW and BLOB columns cannot be referenced in a trigger.

You can't reference columns of these types by using the :NEW qualifier in a trigger, including instead-of-triggers.

14.3.13 PostgreSQL

The following table lists properties for this database:

Table 14-19 Settings for PostgreSQL

Property	Value
Current Timestamp Statement	SELECT (CURRENT_TIMESTAMP)
Case-Sensitive?	No
Commit Keyword	COMMIT
Left Outer Join Operator	LEFT OUTER JOIN

Known Issues

- ◆ PostgreSQL does not support <check-object-password> events. You control authentication by manually inserting entries into the pg_hba.conf file.

14.3.14 Sybase Adaptive Server Enterprise (ASE)

The following table lists properties for this database:

Table 14-20 Settings for Sybase ASE

Property	Value
Current Timestamp Statement	SELECT GETDATE()
Case-Sensitive?	Yes
Commit Keyword	GO
Left Outer Join Operator	*=

Dynamic Defaults

The following table lists database compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly overwrite these settings.

Table 14-21 *Dynamically Configured Sybase ASE Settings*

Display Name	Tag Name	Value
Current timestamp statement	current-timestamp-stmt	SELECT GETDATE()
Left outer-join operator	left-outer-join-operator	*=
Timestamp Translator class	time-translator-class	com.novell.nds.dirxml.driver.jdbc.db.SybaseTimestamp

Known Issues

- ◆ Padding and truncation of binary values.

To ensure ANSI-compliant padding and truncation behavior for binary values, make sure that binary column types (other than `IMAGE`) meet the following criteria:

- ◆ They are exactly the size of the eDirectory™ attribute that maps to them.
- ◆ They are constrained `NOT NULL`.
- ◆ They are added to the Publisher and Subscriber Creation policies.

If they are constrained `NULL`, trailing zeros, which are significant to eDirectory, are truncated. If binary columns exceed the size of their respective eDirectory attributes, extra 0s are appended to the value.

The recommended solution is to use only the `IMAGE` data type when synchronizing binary values.

- ◆ `DATETIME` fractions of a second are rounded.

Sybase Timestamps are at best accurate to 1/300th of a second (approximately .003 seconds). The database server rounds to the nearest 1/300th of a second as opposed to the nearest 1/1000th of a second (.001 seconds or 1 millisecond).

- ◆ Timestamp formats are proprietary.

- ◆ [Section 15.1, “Third-Party JDBC Driver Interoperability,” on page 189](#)
- ◆ [Section 15.2, “JDBC Driver Types,” on page 189](#)
- ◆ [Section 15.3, “Third-Party Jar File Placement,” on page 190](#)
- ◆ [Section 15.4, “Supported Third-Party JDBC Drivers,” on page 191](#)
- ◆ [Section 15.5, “Supported-But-Discouraged Third-Party JDBC Drivers,” on page 202](#)
- ◆ [Section 15.6, “Deprecated Third-Party JDBC Drivers,” on page 207](#)
- ◆ [Section 15.7, “Other Third-Party JDBC Drivers,” on page 208](#)
- ◆ [Section 15.8, “Security Issues,” on page 210](#)

15.1 Third-Party JDBC Driver Interoperability

The Identity Manager Driver for JDBC is designed to interoperate with a specific set of third-party JDBC drivers, instead of a specific set of databases. In fact, the third-party JDBC driver, not the database, is the primary determinant of whether the Driver for JDBC works against any given database. As a general rule, if the Driver for JDBC interoperates well with a given third-party JDBC driver, it interoperates well with databases and database versions that the third-party driver supports.

We strongly recommend that you use the third-party JDBC drivers supplied by major enterprise database vendors whenever possible, such as those listed in [Section 15.4, “Supported Third-Party JDBC Drivers,” on page 191](#). They are usually free, mature, and known to interoperate well with the Driver for JDBC and the databases they target. You can use other third-party drivers, but Novell[®] does not support them.

In general, most third-party drivers are backward compatible. However, even if they are generally backward compatible, they are generally not forward compatible. Anytime a database server is upgraded, the third-party driver used with this product should probably be updated as well.

Also, as a general rule, we recommend that you use the latest version of a third-party driver, unless otherwise noted.

15.2 JDBC Driver Types

Type 1

A third-party JDBC driver that is partially Java and communicates indirectly with a database server through a native ODBC driver.

Type 1 drivers serve as a JDBC-ODBC bridge. Sun provides a JDBC-ODBC bridge driver for experimental use and for situations when no other type of third-party JDBC driver is available.

Type 2

A third-party JDBC driver that is part Java and communicates indirectly with a database server through its native client APIs.

Type 3

A third-party JDBC driver that is pure Java and communicates indirectly with a database server through a middleware server.

Type 4

A third-party JDBC driver that is pure Java and communicates directly with a database server.

15.2.1 Which Type To Use?

Type 3 and 4 drivers are generally more stable than type 1 and 2 drivers. Type 1 and 2 drivers are generally faster than type 3 and 4 drivers. Type 2 and 3 drivers are generally more secure than type 1 and 4 drivers.

Because Identity Manager uses a directory as its datastore, and because databases are usually significantly faster than directories, performance isn't a primary concern. Stability, however, is an issue. For this reason, we recommend that you use a type 3 or 4 third-party JDBC driver whenever possible.

IMPORTANT: If you choose to use a type 1 or type 2 driver (one containing native code) with the Driver for JDBC, use the Remote Loader to ensure the integrity of the directory process.

15.3 Third-Party Jar File Placement

The following tables identify the paths where third-party JDBC driver jar files should be placed on an Identity Manager or Remote Loader server assuming default installation paths.

15.3.1 Identity Manager File Paths

The following table identifies where to place third-party JDBC driver jar files on an Identity Management server, by platform.

Table 15-1 Locations for jar Files: Identity Manager Server

Platform	Directory Path
NetWare®	sys:\system\lib
Solaris, Linux, or AIX	/usr/lib/dirxml/classes (pre-eDirectory 8.8) /opt/novell/eDirectory/lib/dirxml/classes (eDirectory 8.8)
Windows NT/2000	novell\NDS\lib

15.3.2 Remote Loader File Paths

The following table identifies where to place third-party JDBC driver jar files on a Remote Loader server, by platform.

Table 15-2 Locations for jar Files: Remote Loader

Platform	Directory Path
Solaris, Linux, or AIX	/usr/lib/dirxml/classes (pre-eDirectory 8.8) /opt/novell/eDirectory/lib/dirxml/classes (eDirectory 8.8)
Windows NT/2000	novell\RemoteLoader\lib

15.4 Supported Third-Party JDBC Drivers

This section discusses supported drivers. Their use is encouraged.

- ◆ [Section 15.4.1, “Third-Party JDBC Driver Features,” on page 191](#)
- ◆ [Section 15.4.2, “JDBC URL Syntaxes,” on page 192](#)
- ◆ [Section 15.4.3, “JDBC Driver Class Names,” on page 192](#)
- ◆ [Section 15.4.4, “IBM DB2 Universal Database Type 4 JDBC Driver,” on page 193](#)
- ◆ [Section 15.4.5, “Informix JDBC Driver,” on page 194](#)
- ◆ [Section 15.4.6, “jTDS JDBC Driver,” on page 195](#)
- ◆ [Section 15.4.7, “MySQL Connector/J JDBC Driver,” on page 196](#)
- ◆ [Section 15.4.8, “Oracle Thin Client JDBC Driver,” on page 197](#)
- ◆ [Section 15.4.9, “Oracle OCI JDBC Driver,” on page 199](#)
- ◆ [Section 15.4.10, “PostgreSQL JDBC Driver,” on page 201](#)
- ◆ [Section 15.4.11, “Sybase Adaptive Server Enterprise JConnect JDBC Driver,” on page 201](#)

15.4.1 Third-Party JDBC Driver Features

The following table summarizes third-party JDBC driver features:

Table 15-3 Third-Party JDBC Driver Features

Driver	Supports Encrypted Transport?	Supports Retrieval of Auto-Generated Keys?
IBM DB2 UDB Type 4	No	No
Informix	No	No
Informix	No	No
MySQL Connector/J	Yes	Yes
jTDS	Yes	Yes
Oracle Thin Client	Yes	No
Oracle OCI	Yes	No
PostgreSQL	Yes*	No
Sybase jConnect	Yes	No

* For JDBC 3 (Java 1.4) versions and later.

15.4.2 JDBC URL Syntaxes

The following table lists URL syntaxes for supported third-party JDBC drivers:

Table 15-4 *URL Syntaxes*

Third-Party JDBC Driver	JDBC URL Syntax
IBM DB2 UDB Type 4, Universal	<code>jdbc:db2://ip-address:50000/database-name</code>
Informix	<code>jdbc:informix-sqli://ip-address:1526/database-name:informixserver=server-id</code>
jTDS	<code>jdbc:jtds:sqlserver://ip-address/database-name</code>
MySQL Connector/J	<code>jdbc:mysql://ip-address:3306/database-name</code>
Oracle OCI	<code>jdbc:oracle:oci8:@tns-name</code>
Oracle Thin Client	<code>jdbc:oracle:thin:@ip-address:1521:sid</code>
PostgreSQL	<code>jdbc:postgresql://ip-address:5432/database-name</code>
Sybase jConnect	<code>jdbc:sybase:Tds:ip-address:2048/database-name</code>

This information is used in conjunction with the Authentication Context parameter. For information on this parameter, see [“Authentication Context” on page 58](#).

15.4.3 JDBC Driver Class Names

The following table lists the fully-qualified Java class names of supported third-party JDBC drivers:

Table 15-5 *Class Names of Third-Party JDBC Drivers*

Third-party JDBC Driver	Class Name
IBM DB2 UDB Type 4, Universal	<code>com.ibm.db2.jcc.DB2Driver</code>
Informix	<code>com.informix.jdbc.IfxDriver</code>
jTDS	<code>net.sourceforge.jtds.jdbc.Driver</code>
MySQL Connector/J	<code>org.gjt.mm.mysql.Driver</code>
Oracle OCI	<code>oracle.jdbc.driver.OracleDriver</code>
Oracle Thin Client	<code>oracle.jdbc.driver.OracleDriver</code>
PostgreSQL	<code>org.postgresql.Driver</code>
Sybase jConnect 5.5	<code>com.sybase.jdbc2.jdbc.SybDriver</code>
Sybase jConnect 6.05	<code>com.sybase.jdbc3.jdbc.SybDriver</code>

This information is used in conjunction with the JDBC Driver Class Name parameter. For information on this parameter, see “[Third-Party JDBC Driver Class Name](#)” on page 60.

15.4.4 IBM DB2 Universal Database Type 4 JDBC Driver

Table 15-6 IBM DB2 Driver: Type 4

Supported Database Versions	8.x
Class Name	com.ibm.db2.jcc.DB2Driver
Type	4
URL Syntax	<code>jdbc:db2://ip-address:50000/database-name</code>
Download Instructions	Download as part of the latest FixPack (recommended). IBM Support & Downloads (http://www.ibm.com/support/us/) or Copy the file from the database server. <code>file:///database-installation-directory/java</code>
Filename	db2jcc.jar, db2jcc_license_cu.jar, db2jcc_javax.jar (optional)
Documentation URLs	DB2 Information Center (http://publib.boulder.ibm.com/infocenter/db2help) DB2 Universal JDBC Driver (http://publib.boulder.ibm.com/infocenter/db2help/index.jsp?topic=/com.ibm.db2.udb.doc/ad/t0010264.htm) Security under the DB2 Universal JDBC Driver (http://publib.boulder.ibm.com/infocenter/db2help/index.jsp?topic=/com.ibm.db2.udb.doc/ad/cjvjcsec.htm)

NOTE: Unlike the type 3 driver, the type 4 driver has only a minimal set of defined error codes. This absence inhibits the Driver for JDBC’s ability to distinguish between connectivity, retry, authentication, and fatal error conditions.

Compatibility

The IBM DB2 driver is backward compatible. However, it doesn’t work with database version 7. Database server updates are frequent. Driver updates are infrequent.

Security

The IBM DB2 driver supports a variety of authentication security mechanisms but does not support encrypted transport.

Known Issues

- ◆ It’s very difficult to diagnose and remedy Java-related errors on the database server.

Numerous error conditions and error codes can arise when you attempt to install and execute user-defined stored procedures and functions written in Java. Diagnosing these can prove time intensive and frustrating. A log file (`db2diag.log` on the database server) can often provide additional debugging information. In addition, all error codes are documented and available online.

15.4.5 Informix JDBC Driver

Table 15-7 *Informix JDBC Driver*

Supported Database Versions	Dynamic Server 7.x, 9.x, 10.x
Class Name	<code>com.informix.jdbc.IfxDriver</code>
Type	4
URL Syntax	<code>jdbc:informix-sqli://ip-address:1526/database-name:informixserver=server-id</code>
Download Instructions	Download URL (http://www-306.ibm.com/software/data/informix/tools/jdbc)
Filenames (7, 9)	<code>ifxjdbc.jar</code> , <code>ifxjdbcx.jar</code> (optional)
Filenames (10)	<code>jdbc.jar</code> , <code>jdbcx.jar</code> (optional)
Documentation URLs	Informix Information Center (http://publib.boulder.ibm.com/infocenter/ids9help/index.jsp) Informix JDBC Driver (http://www-306.ibm.com/software/data/informix/pubs/library/jdbc_2.html)

Compatibility

The Informix driver is backward compatible. Database server updates and driver updates are infrequent.

Security

The Informix driver does not support encrypted transport.

Required Parameter Settings for ANSI-Compliant Databases

The following table lists driver parameters that must be explicitly set for the Driver for JDBC to interoperate with the Informix driver against ANSI-compliant databases.

Table 15-8 *Driver Settings for ANSI-Compliant Databases*

Display Name	Tag Name	Value
Supports schemas in metadata retrieval?	<code>supports-schemas-in-metadata-retrieval</code>	false
	See “Supports Schemas in Metadata Retrieval?” on page 79.	

Display Name	Tag Name	Value
Force username case:	force-username-case	upper
	See “Force Username Case” on page 77 .	

Dynamic Parameter Defaults

The following table lists driver compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not override these settings.

Table 15-9 Informix JDBC Settings Not to Override

Display Name	Tag Name	Value
Function return method:	function-return-method	result set
	See “Function Return Method” on page 79 .	

Known Issues

- ◆ Schema names cannot be used to retrieve metadata against an ANSI-compliant database. Set the driver compatibility parameter [“Supports Schemas in Metadata Retrieval?” on page 79](#) to Boolean False.

The database objects available for metadata retrieval are those visible to the database user who authenticated to the database. Schema qualifiers cannot be used to identify database objects. Therefore, to avoid naming collisions (such as, owner1.table1, owner2.table1), give the database authentication user only SELECT privileges on objects being synchronized.

- ◆ When used against ANSI-compliant databases, usernames must be in uppercase. Set the driver compatibility parameter [“Force Username Case” on page 77](#) to upper.

15.4.6 jTDS JDBC Driver

Table 15-10 jTDS Driver Settings

Supported Database Versions:	Microsoft SQL Server 6.5, 7, 2000 (8), 2005 (9)
Class Name	net.sourceforge.jtds.jdbc.Driver
Type	4 (2 if NTLM or SSO authentication is enabled)
URL Syntax	<i>jdbc:jtds:sqlserver://ip-address/database-name</i>
Download Instructions	The jTDS Project (http://jtds.sourceforge.net/)
Filenames	jtds-<version>.jar

Limitations

Novell recommends that you use the Microsoft 2000 JDBC driver when Subscribing to views.

IMPORTANT: The jTDS JDBC driver does not support views or stored procedures.

Compatibility

The jTDS driver works with all versions of Microsoft SQL Server. It also supports all versions of Sybase ASE, but it hasn't been tested by Novell against that database server yet. Driver updates are infrequent.

Security

The jTDS driver supports encrypted transport.

URL Properties

Delimit URL properties by using a semicolon (;).

The following table lists values for the domain URL property for the jTDS driver.

Table 15-11 Values for the domain URL Property

Legal Value	Description
<any-domain-name>	When a domain name is specified, either NTLM or SSO authentication can be used. NTLM authentication is selected when a username and password are supplied. SSO authentication is selected when a username and password are not supplied.
<no-value>	JDBC authentication is used.

The following table lists values for the SSL URL property for the jTDS driver.

Table 15-12 Values for the ssl URL Property

Legal Value	Description
off	SSL is not used. This is the default.
request	SSL is used if the server supports it.
require	SSL is required. An exception is thrown if the server doesn't support it.
authenticate	SSL is required. An exception is thrown if the server doesn't support it. In addition, the server's certificate must be signed by a trusted certificate authority.

15.4.7 MySQL Connector/J JDBC Driver

Table 15-13 Settings for the MySQL Connector/J JDBC Driver

Supported Database Versions	3.x, 4.x
-----------------------------	----------

Class Name	org.gjt.mm.mysql.Driver
Type	4
URL Syntax	<code>jdbc:mysql://ip-address:3306/database-name</code>
Download Instructions	Download and extract. The <code>jar</code> file is located in the <code>extract-dir/mysql-connector-java-version</code> directory. MySQL Connector/J (http://www.mysql.com/products/connector/j/)
Filename	<code>mysql-connector-java-version-bin.jar</code>
Documentation URLs	MySQL Connector/J Documentation (http://dev.mysql.com/doc/refman/5.0/en/java-connector.html) Connecting Over SSL (http://dev.mysql.com/doc/refman/5.0/en/cj-using-ssl.html)

Also see “[Generation/Retrieval Method \(Table-Global\)](#)” on page 84.

Compatibility

The Connector/J driver is backward compatible. Database server updates are frequent. Driver updates are infrequent.

Security

The Connector/J driver supports JSSE (Java Secure Sockets Extension) SSL-encrypted transport.

Required Parameter Settings for MyISAM Tables

The following table lists driver parameters that you must set so that the Driver for JDBC can interoperate with the Connector/J driver against MyISAM tables.

Table 15-14 *Settings for MyISAM Tables*

Display Name	Tag Name	Value
Use manual transactions?	<code>use-manual-transactions</code>	<code>false</code>

15.4.8 Oracle Thin Client JDBC Driver

Table 15-15 *Oracle Thin Client Settings*

Supported Database Versions	8i, 9i, 10g
Class Name	<code>oracle.jdbc.driver.OracleDriver</code>
Type	4
URL Syntax	<code>jdbc:oracle:thin:@ip-address:1521:sid</code>

Download Instructions	Register for free and download. Oracle Technology Network (http://otn.oracle.com/software/tech/java/sqlj_jdbc/content.html)
Filenames	ojdbc14.jar, orai18n.jar (optional) Filenames for different JVM versions (http://www.oracle.com/technology/tech/java/sqlj_jdbc/htdocs/jdbc_faq.htm#02_07)
Documentation URLs	Oracle Advanced Security (http://www.stanford.edu/dept/itss/docs/oracle/10g/java.101/b10979/clientsec.htm) JDBC FAQ (http://www.oracle.com/technology/tech/java/sqlj_jdbc/htdocs/jdbc_faq.htm)

Compatibility

The Thin Client driver is backward compatible. Database server updates and driver updates are infrequent.

Oracle releases thin client drivers for various JVMs. Even though all of them work with this product, we recommend that you use the 1.4 version.

Security

The Thin Client driver supports Oracle Advanced Security encrypted transport.

Dynamic Parameter Defaults

The following table lists driver compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly override these settings.

Table 15-16 Oracle Thin Client Settings Not to Override

Display Name	Tag Name	Value
Number of returned result sets:	handle-stmt-results	single

Connection Properties

The following table lists important connection properties for this driver.

Table 15-17 Oracle Thin Client: Connection Properties

Property	Significance
includeSynonyms	If the value of this property is <code>true</code> , synonym column metadata is available.
ORACLE.NET.ENCRYPTION_CLIENT	Defines the level of security that the client wants to negotiate with the server.
ORACLE.NET.ENCRYPTION_TYPES_CLIENT	Defines the encryption algorithm to be used.

Property	Significance
ORACLE.NET.CRYPTO_CHECKSUM_CLIENT	Defines the level of security that it wants to negotiate with the server for data integrity.
ORACLE.NET.CRYPTO_CHEKSUM_TYPES_CLIENT	Defines the data integrity algorithm to be used.

Known Issues

- ♦ High CPU utilization triggered by execution of embedded SQL statements:

The most common problem experienced with this driver is high CPU utilization. As a result, this driver always indicates that more results are available from calls to method `java.sql.Statement.execute(String stmt)`, which can lead to an infinite loop condition. This condition occurs only if all the following happen:

- ♦ A value other than `single`, `no` or `one` in the driver compatibility parameter “Number of Returned Result Sets” on page 75 is being executed.
- ♦ An embedded SQL statement is being executed.
- ♦ The type of statement is not explicitly specified.

To avoid the conditions that produce high CPU utilization:

- ♦ Do not explicitly set this parameter. Defer to the dynamic default value.
- ♦ Always place a `jdbc:type` attribute on embedded `<jdbc:statement>` elements.

NOTE: The `jdbc` namespace prefix must map to `urn:dirxml:jdbc`.

- ♦ Can’t retrieve synonym column metadata:

The connection property `includeSynonyms` must be set to `true`.

- ♦ Can’t see synonym table primary key constraint:

The only known solution to this problem is to use a view.

15.4.9 Oracle OCI JDBC Driver

Table 15-18 Oracle OCI JDBC Driver Settings

Supported Database Versions	8i, 9i, 10g
Class Name	<code>oracle.jdbc.driver.OracleDriver</code>
Type	2
URL Syntax	<code>jdbc:oracle:oci8:@tns-name</code>

Download Instructions	<p>The SQLNet infrastructure is the main requirement for OCI. SQLNet can run on any platform that Oracle supports, not just Linux.</p> <p>For Linux, register for free and download the following:</p> <ul style="list-style-type: none"> ◆ The Oracle Instant Client (<code>instantclient-basic-linux32-10.2.0.1-20050713.zip</code>) from Instant Client Downloads (http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html). ◆ The Oracle SQL*Plus binary (<code>instantclient-sqlplus-linux32-10.2.0.1-20050713.zip</code>) from Instant Client Downloads (http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html).
Filenames	<p><code>ojdbc14.jar</code>, <code>orail8n.jar</code> (optional)</p> <p>Filenames for different JVM versions (http://www.oracle.com/technology/tech/java/sqlj_jdbc/htdocs/jdbc_faq.htm#02_07)</p>
Documentation URLs	<p>Oracle Call Interface (http://www.oracle.com/technology/tech/oci/index.html)</p> <p>OCI FAQ (http://www.oracle.com/technology/tech/oci/htdocs/oci_faq.html)</p> <p>Oracle Advanced Security (http://www.stanford.edu/dept/itss/docs/oracle/10g/java.101/b10979/clientsec.htm)</p> <p>Instant Client (http://www.oracle.com/technology/tech/oci/instantclient/index.html)</p> <p>Instant Client (http://download-west.oracle.com/docs/cd/B12037_01/java.101/b10979/instclient.htm#CHDGDIGG)</p>

You can install SQLNet by doing either of the following:

- ◆ Use the Instant Client (which bypasses unneeded components of the full version).
- ◆ Download the full package from Oracle.

If the database is running on the same server as Identity Manager, you don't need to install SQLNet because SQLNet comes as standard on the database server.

The Oracle OCI driver is essentially the same as the Thin Client driver. See [Section 15.4.8, “Oracle Thin Client JDBC Driver,” on page 197](#). The OCI client differs in the following ways:

- ◆ The OCI Client supports clustering, failover, and high availability.
- ◆ The OCI Client has additional security options.

For information on setting up the Oracle OCI Client, see [Appendix K, “Setting Up an OCI Client on Linux,” on page 251](#).

15.4.10 PostgreSQL JDBC Driver

Table 15-19 *PostgreSQL JDBC Driver Settings*

Supported Database Versions	6.x, 7.x, 8.x
Class Name	org.postgresql.Driver
Type	4
URL Syntax	<code>jdbc:postgresql://ip-address:5432/database-name</code>
Download Instructions	JDBC Driver Download (http://jdbc.postgresql.org/download.html)
Documentation URLs	JDBC Driver Documentation (http://jdbc.postgresql.org/documentation/docs.html) Using SSL (http://jdbc.postgresql.org/documentation/80/ssl.html)

NOTE: The filename of the PostgreSQL varies by database version.

Compatibility

The latest builds of the PostgreSQL driver are backward compatible through server version 7.2. Database server updates and driver updates are frequent.

Security

The PostgreSQL driver supports SSL-encrypted transport for JDBC 3 driver versions.

15.4.11 Sybase Adaptive Server Enterprise JConnect JDBC Driver

Table 15-20 *Settings for the Sybase Adaptive Server Enterprise Driver*

Supported Database Versions	Adaptive Server* Enterprise 11.x, 12.x
Class Name	com.sybase.jdbc2.jdbc.SybDriver (for jconn2.jar) com.sybase.jdbc3.jdbc.SybDriver (for jconn3.jar)
Type	4
URL Syntax	<code>jdbc:sybase:Tds:ip-address:2048/database-name</code>
Download Instructions	Sybase Downloads (http://www.sybase.com/downloads)
Filenames	<code>jconn2.jar</code> or <code>jconn3.jar</code>
Documentation URLs	JConnect Documentation (http://sybooks.sybase.com/onlinebooks/group-jc/jcg0600e/prjdbc)

* For JDBC 3 (Java 1.4) versions and later.

Compatibility

The Adaptive Server driver is backward compatible. Database server updates and driver updates are infrequent.

Security

The Adaptive Server driver supports SSL-encrypted transport. To enable SSL encryption, you must specify a custom socket implementation via the SYB SOCKET_FACTORY connection property. For additional information on how to set connection properties, see [“Connection Properties” on page 70](#).

Connection Properties

The SYB SOCKET_FACTORY property can be used to specify the class name of a custom socket implementation that supports encrypted transport.

15.5 Supported-But-Discouraged Third-Party JDBC Drivers

This section identifies drivers that are supported but whose use is discouraged.

- ◆ [Section 15.5.1, “Third-Party JDBC Driver Features,” on page 202](#)
- ◆ [Section 15.5.2, “JDBC URL Syntaxes,” on page 202](#)
- ◆ [Section 15.5.3, “JDBC Driver Class Names,” on page 203](#)
- ◆ [Section 15.5.4, “IBM DB2 Universal Database Type 3 JDBC Driver,” on page 203](#)
- ◆ [Section 15.5.5, “Microsoft SQL Server 2000 Driver for JDBC,” on page 204](#)
- ◆ [Section 15.5.6, “Microsoft SQL Server 2005 JDBC Driver,” on page 206](#)

15.5.1 Third-Party JDBC Driver Features

The following table summarizes third-party JDBC driver features:

Table 15-21 *Third-Party JDBC Driver Features*

Driver	Supports Encrypted Transport?	Supports Retrieval of Auto-Generated Keys?
IBM DB2 UDB Type 3	No	No
Microsoft 2000	No	No
Microsoft 2005	Yes	Yes

* For JDBC 3 (Java 1.4) versions and later.

15.5.2 JDBC URL Syntaxes

The following table lists URL syntaxes for supported third-party JDBC drivers:

Table 15-22 URL Syntaxes

Third-Party JDBC Driver	JDBC URL Syntax
IBM DB2 UDB Type 3	<code>jdbc:db2://ip-address:6789/database-name</code>
Microsoft SQL Server 7, 2000	<code>jdbc:microsoft:sqlserver://ip-address-or-dns-name:1433;DatabaseName=database-name</code>
Microsoft SQL Server 2005	<code>jdbc:sqlserver://ip-address-or-dns-name:1433;databaseName=database-name</code>

This information is used in conjunction with the Authentication Context parameter. For information on this parameter, see [“Authentication Context” on page 58](#).

15.5.3 JDBC Driver Class Names

The following table lists the fully-qualified Java class names of supported third-party JDBC drivers:

Table 15-23 Class Names of Third-Party JDBC Drivers

Third-party JDBC Driver	Class Name
IBM DB2 UDB Type 3	<code>COM.ibm.db2.jdbc.net.DB2Driver</code>
Microsoft 2000	<code>com.microsoft.jdbc.sqlserver.SQLServerDriver</code>
Microsoft 2005	<code>com.microsoft.sqlserver.jdbc.SQLServerDriver</code>

This information is used in conjunction with the JDBC Driver Class Name parameter. For information on this parameter, see [“Third-Party JDBC Driver Class Name” on page 60](#).

15.5.4 IBM DB2 Universal Database Type 3 JDBC Driver

Table 15-24 IBM DB2 Driver: Type 3

Supported Database Versions:	7.x
Class Name:	<code>COM.ibm.db2.jdbc.net.DB2Driver</code>
Type	3
URL Syntax:	<code>jdbc:db2://ip-address:6789/database-name</code>
Download Instructions:	Copy the file from the database server. <code>file:///database-installation-directory/java</code>
Filename:	<code>db2java.zip</code>

Documentation URLs: [DB2 Information Center \(http://publib.boulder.ibm.com/infocenter/db2v7luw\)](http://publib.boulder.ibm.com/infocenter/db2v7luw)

[JDBC Programming \(http://publib.boulder.ibm.com/infocenter/db2v7luw/index.jsp?topic=/com.ibm.db2v7.doc/db2a0/db2a0159.htm\)](http://publib.boulder.ibm.com/infocenter/db2v7luw/index.jsp?topic=/com.ibm.db2v7.doc/db2a0/db2a0159.htm)

IMPORTANT: The type 3 driver was deprecated as of DB2 UDB version 8.

Compatibility

The IBM DB2 driver can best be characterized as version-hypersensitive. It is not compatible across major or minor versions of DB2, including FixPacks. For this reason, we recommend that you use the file installed on the database server.

IMPORTANT: The IBM DB2 driver must be updated on the Identity Manager or Remote Loader server every time the target database is updated, even if only at the FixPack level.

Security

The IBM DB2 driver does not support encrypted transport.

Known Issues

- ♦ A version mismatch usually results in connectivity-related failures.

The most common problem experienced with the IBM DB2 driver is because of a driver/database version mismatch. The symptom of a version mismatch is connectivity-related failures such as "CLI0601E Invalid statement handle or statement is closed." To remedy the problem, overwrite the `db2java.zip` file on the Identity Manager or Remote Loader server with the version installed on the database server.

- ♦ It's very difficult to diagnose and remedy Java-related errors on the database server.

Numerous error conditions and error-codes can arise when you attempt to install and execute user-defined stored procedures and functions written in Java. Diagnosing them can prove time intensive and frustrating. A log file (`db2diag.log` on the database server) can often provide additional debugging information. In addition, all error codes are documented and available online.

15.5.5 Microsoft SQL Server 2000 Driver for JDBC

Table 15-25 *Microsoft SQL Server 2000 Driver Settings*

Supported Database Versions:	2000 (8)
Class Name	<code>com.microsoft.jdbc.sqlserver.SQLServerDriver</code>
Type	4
URL Syntax	<code>jdbc:sqlserver://ip-address-or-dns-name:1433;databaseName=database-name</code>

Download Instructions	Microsoft JDBC Downloads (http://www.microsoft.com/downloads/results.aspx?sortCriteria=date&OSID=&productID=&CategoryID=&featuretext=jdbc&DisplayLang=en&DisplayEnglishAlso=)
Filenames	msbase.jar, mssqlserver.jar, msutil.jar

NOTE: The filename, URL syntax and classname differ (often subtly) from those of the 2005 driver.

Compatibility

The SQL Server 2000 driver only works with SQL Server 2000. It is backward compatible. However, it doesn't work with database version 7. Database server and driver updates are infrequent.

Security

The SQL Server 2000 driver does not support encrypted transport.

URL Properties

Delimit URL properties by using a semicolon (;).

The following table lists values for the SelectMethod URL property for the SQL Server 2000 driver.

Table 15-26 Values for the SelectMethod URL Property

Legal Value	Description
direct	The default value. Doesn't allow for multiple active statements on a single connection
cursor	Allows for multiple active statements on a single connection

Dynamic Parameter Defaults

The following table lists driver compatibility parameters that the Driver for JDBC implicitly sets at runtime. Do not explicitly override these settings.

Table 15-27 SQL Server 2000 Settings Not to Override

Display Name	Tag Name	Value
Reuse Statements?	reuse-statements	false

Known Issues

- ♦ Can't start manual transaction because of cloned connections.

An implementation anomaly that doesn't allow concurrent statements to be active on the same connection causes the most common problem experienced with the SQL Server 2000 driver. Unlike other third-party implementations, the SQL Server 2000 driver can have only one

`java.sql.Statement` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html>) object active at a time on a given connection.

If you attempt to use more than one statement object, the following error is issued: “Can’t start manual transaction mode because there are cloned connections.” This error can occur only if the driver compatibility parameter “Reuse Statements?” on page 74 is set to Boolean True. As a best practice, never explicitly set this parameter. Instead, defer to the dynamic default value.

An alternative is to place the delimited property `; SelectMethod=cursor` at the end of the URL string. For additional information on this issue, consult the following support articles:

- ◆ Document 30096 (<http://knowledgebase.datadirect.com/kbase.nsf/SupportLink+Online/30096?OpenDocument>) by DataDirect Technologies*
- ◆ Article 313181 (<http://support.microsoft.com/default.aspx?scid=kb%3Ben-us%3B313181>) by Microsoft
- ◆ Association values that contain `UNIQUEIDENTIFIER` columns are inconsistent between driver versions.

Earlier versions of the SQL Server 2000 driver returned a non-standard `java.sql.Types` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>) value for native `UNIQUEIDENTIFIER` columns. To compensate, the Driver for JDBC mapped that non-standard type to the standard type `java.sql.Types.BINARY` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>) because it best mirrored the native database type, which is a 16-byte value. This mapping results in a Base64-encoded association value.

Later versions of the SQL Server 2000 driver return a standard type `java.sql.CHAR` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>). This mapping results in a non-Base64-encoded association value, effectively invalidating all associations generated by using earlier versions of the SQL Server 2000 driver. This change effectively breaks backward compatibility.

The best solution to this problem is to continue using the earlier version of the SQL Server 2000 driver. If you must upgrade, remove all invalidated associations and reassociate all previously-associated objects.

15.5.6 Microsoft SQL Server 2005 JDBC Driver

Table 15-28 *Microsoft SQL Server 2005 Driver Settings*

Supported Database Versions:	2005 (9)
Class Name	<code>com.microsoft.sqlserver.jdbc.SQLServerDriver</code>
Type	4 (2 if integrated security is enabled)
URL Syntax	<code>jdbc:sqlserver://ip-address-or-dns-name:1433;databaseName=database-name</code>
Download Instructions	Microsoft SQL Server 2005 JDBC Driver (http://msdn2.microsoft.com/en-us/data/aa937724.aspx)
Filenames	<code>sqljdbc.jar</code>

NOTE: The filename, URL syntax, and classname differ (often subtly) from those of the 2000 driver.

Compatibility

The SQL Server 2005 driver works with only SQL Server 2005. Database server and driver updates are infrequent.

Security

The SQL Server 2005 driver does not support encrypted transport.

URL Properties

Delimit URL properties by using a semicolon (;).

The following table lists values for the `integratedSecurity` URL property for the SQL Server 2005 driver.

Table 15-29 Values for the `integratedSecurity` URL Property

Legal Value	Description
false	The default value. JDBC authentication is used.
true	Windows process-level authentication is used.

15.6 Deprecated Third-Party JDBC Drivers

This section contains a list of third-party JDBC drivers that are no longer supported.

- ♦ [Section 15.6.1, “BEA Weblogic jDriver for Microsoft SQL Server,” on page 207](#)

15.6.1 BEA Weblogic jDriver for Microsoft SQL Server

Table 15-30 BEA Weblogic jDriver

Supported Database Version:	Microsoft SQL Server 6.5, 7.x, 8.x (2000)
Class Name	<code>weblogic.jdbc.mssqlserver4.Driver</code>
Type	4
URL Syntax	<code>jdbc:weblogic:mssqlserver4:database-name@ip-address:1433</code>
Download Instructions	Register for free and download the latest version of Weblogic server. Run the installer. The <code>weblogic.jar</code> file is installed in the <code>install-dir/server/lib</code> directory. BEA Download Center (http://commerce.bea.com/showallversions.jsp?family=WLS)
Filename	<code>weblogic.jar</code>

NOTE: The BEA Weblogic driver is included in the supported third-party driver listing to provide JDBC access to Microsoft SQL server 7. Microsoft's driver supports only version 8 (2000).

Compatibility

The BEA Weblogic driver is backward-compatible. Database server and driver updates are infrequent.

Security

The BEA Weblogic driver does not support encrypted transport.

Known Issues

- ♦ The BEA Weblogic driver is not free. It must be purchased and properly licensed.
- ♦ Association values that contain `UNIQUEIDENTIFIER` columns are inconsistent between driver versions.

Earlier versions of the BEA Weblogic driver returned a non-standard `java.sql.Types` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>) value for native `UNIQUEIDENTIFIER` columns. To compensate, the Driver for JDBC mapped that non-standard type to the standard type `java.sql.Types.BINARY` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>) because it best mirrored the native database type, which is a 16-byte value. This mapping results in a Base64-encoded association value.

Later versions of the BEA Weblogic driver return a standard type `java.sql.CHAR` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>). This mapping results in a non-Base64-encoded association value, effectively invalidating all associations generated by using earlier versions of the BEA Weblogic driver. This change effectively breaks backward compatibility.

The best solution to this problem is to continue using the earlier version of the BEA Weblogic driver. If you must upgrade, you must remove all invalidated associations and reassociate all previously-associated objects.

- ♦ The BEA Weblogic driver throws a `java.lang.IllegalMonitorStateException` (<http://java.sun.com/j2se/1.5.0/docs/api/java/lang/IllegalMonitorStateException.html>) when method `java.sql.Connection.getConnection(String url, String username, String password)` is called on AIX.

15.7 Other Third-Party JDBC Drivers

This section lists unsupported drivers of interest and discusses how to assess the feasibility of using unsupported third-party JDBC drivers with this product.

- ♦ [Section 15.7.1, “IBM Toolbox for Java/JTOpen,” on page 209](#)
- ♦ [Section 15.7.2, “Minimum Third-Party JDBC Driver Requirements,” on page 209](#)
- ♦ [Section 15.7.3, “Considerations When Using Other Third-Party JDBC Drivers,” on page 209](#)

15.7.1 IBM Toolbox for Java/JTOpen

Table 15-31 Settings for IBM Toolbox for Java/JTOpen

Database	IBM Toolbox for Java/JTOpen <ul style="list-style-type: none">◆ iSeries Toolbox for Java (alias)◆ AS/400 Toolbox for Java (alias)
Class Name	com.ibm.as400.access.AS400JDBCdriver
Type	4
URL Syntax	jdbc:as400://ip-address/database-name
Download Instructions	Download URLs for JTOpen <ul style="list-style-type: none">◆ JTOpen (http://jt400.sourceforge.net)◆ Toolbox for Java/JTOpen (http://www-03.ibm.com/servers/eserver/iseries/toolbox/downloads.html)
Filenames	jt400.jar
Documentation URLs	Toolbox for Java/JTOpen (http://www-03.ibm.com/servers/eserver/iseries/toolbox/)

If you use the IBM Toolbox for Java/JTOpen driver, you must manually enter values for the JDBC Driver Class Name and Authentication Context parameters. The settings are not automatically populated. See “[Third-Party JDBC Driver Class Name](#)” on page 60 and “[Authentication Context](#)” on page 58.

15.7.2 Minimum Third-Party JDBC Driver Requirements

The Driver for JDBC might not interoperate with all third-party JDBC drivers. If you use an unsupported third-party JDBC driver, it must meet the following requirements:

- ◆ Support required metadata methods
For a current list of the required and optional `java.sql.DatabaseMetaData` method calls that the Driver for JDBC makes, see [Appendix D, “java.sql.DatabaseMetaData Methods,”](#) on page 233.
- ◆ Support other required JDBC methods
For a list of required JDBC methods that the Driver for JDBC uses, refer to [Appendix E, “JDBC Interface Methods,”](#) on page 235. You can use this list in collaboration with third-party driver documentation to identify potential incompatibilities.

15.7.3 Considerations When Using Other Third-Party JDBC Drivers

- ◆ Because the Driver for JDBC is directly dependent upon third-party JDBC driver implementations, bugs in those implementations might cause this product to malfunction.

To assist you in debugging third-party JDBC drivers, the Driver for JDBC supports the following:

- ◆ Tracing at the JDBC API level (level 6)
- ◆ Third-party JDBC driver (level 7) tracing
- ◆ Stored procedure or function support is a likely point of failure.
- ◆ You'll probably need to write a custom driver descriptor file.

Specifically, you'll need to categorize error codes and SQL states for the third-party driver that you are using.

15.8 Security Issues

To ensure that a secure connection exists between the Identity Manager Driver for JDBC and a third-party driver, we recommend the following:

- ◆ Run the Driver for JDBC remotely on the database server.
- ◆ Use SSL to encrypt communications between the Identity Manager server and the database server.

If you cannot run the Driver for JDBC remotely, you might want to use a type 2 or type 3 JDBC driver. These driver types often facilitate a greater degree of security through middleware servers or client APIs unavailable to other JDBC driver types. Some type 4 drivers support encrypted transport, but encryption is the exception rather than the rule.

The Association Utility normalizes associations of objects associated under the 1.0 or later versions of the Driver for JDBC. It also provides several other features that simplify driver administration.

This version of the utility is compatible with the 1.0 and later versions of the Driver for JDBC, and supersedes all previous versions.

- ◆ [Section 16.1, “Independent Operations,” on page 211](#)
- ◆ [Section 16.2, “Before You Begin,” on page 212](#)
- ◆ [Section 16.3, “Using the Association Utility,” on page 213](#)
- ◆ [Section 16.4, “Editing Associations,” on page 213](#)

16.1 Independent Operations

The Association Utility supports seven independent operations:

Table 16-1 *Independent Operations*

Operation	Description	Read-Write Functionality
1	List objects associated with a driver (default).	Read-only
2	List objects that have multiple associations to a driver.	Read-only
3	List objects that have invalid associations to a driver. An association is invalid if: <ul style="list-style-type: none">◆ It is malformed. For example, the association is missing the schema RDN, missing the table RDN, or the schema keyword is misspelled.◆ It contains database identifiers that do not map to identifiers in the target database. For example, an association includes a mapping to a table that does not exist.◆ It maps to no row or multiple rows. An association is broken if it doesn't map to a row. Also, associations aren't unique if they map to more than one row.	Read-only
4	List objects that need to be normalized. A normalized association is valid, correctly ordered, and uses the correct case. Normal case is uppercase for case-insensitive databases and mixed case for case-sensitive databases.	Read-only
5	Normalize object associations listed during operation 4.	Write

Operation	Description	Read-Write Functionality
6	<p>List object associations to be modified.</p> <p>Allows for global replacement of schema, table, and column names based on search criteria.</p> <p>This operation requires two parameters (oldRDN and newRDN). See “Editing Associations” on page 213.</p>	Read-only
7	<p>Modify object associations listed during operation 6.</p> <p>This operation requires two parameters (oldRDN and newRDN). See “Editing Associations” on page 213.</p>	Write

16.2 Before You Begin

Modifying associations can potentially cause problems. If associations are corrupted, Identity Manager ceases to function. Therefore, use write operations only when necessary. To avoid unintentionally corrupting an association, the Association Utility creates an undo Idiff file for all write operations.

Review the following cautions before using the utility:

- ♦ The Association Utility, like the driver, assumes database identifiers are undelimited (unquoted and contain no special characters).
- ♦ Update all object associations related to a driver together.

IMPORTANT: It is extremely important that you update, at the same time, all object associations related to a driver.

To see all of the objects associated with a particular driver, run the Association Utility on the Identity Manager server associated with a particular driver instance.

The LDAP search base must contain all of the objects associated with a particular driver.

NOTE: To ensure complete containment, we recommend that you use your tree’s root container as the search base.

- ♦ Make sure that the JDBC URL of the target database supplied to this utility is the same as the URL that the driver uses. Pointing this utility at a case-insensitive database when the database is actually case-sensitive might result in associations being normalized to the wrong case.
 - ♦ Because the Association Utility runs locally, it uses an unsecured connection. Therefore, the Identity Vault LDAP server must be temporarily configured to accept clear text passwords. Depending upon the third-party JDBC driver you are using, the database connection established by this utility might be insecure.
-

NOTE: We recommend changing the driver’s authentication password on the database after you run this utility.

16.3 Using the Association Utility

Run the Association Utility once for each instance of the driver installed on an Identity Manager server. In the *install-dir\jdbc\util* directory, a batch file *association.bat* or shell script *association.sh* (depending upon your platform) starts the utility.

A properties file containing association utility parameters is provided for each supported database. These files are in the *install-dir\jdbc\util* directory.

Database	Properties Filename
IBM DB2 Universal Database	<i>properties_db2.txt</i>
Informix Dynamic Server	<i>properties_ifx_ansi.txt1</i> <i>properties_ifx_log.txt</i> <i>properties_ifx_no_log.txt</i>
Microsoft SQL Server	<i>properties_ms.txt</i>
MySQL	<i>properties_my.txt</i>
Oracle	<i>properties_ora.txt</i>
PostgreSQL	<i>properties_pg.txt</i>
Sybase Adaptive Server Enterprise	<i>properties_syb.txt</i>

¹This utility does not work with Informix ANSI-compliant databases.

NOTE: For more information on how to run the utility from the command line, refer to *run.bat* in the *install-dir\tools\util* directory.

- 1 Stop the driver.
- 2 Run the Association Utility to identify and remove extraneous associations (operations 2 and 3).

No object associated by this product should have multiple associations. Manually remove extraneous associations on a per object basis. Operation 3 might help you identify which of the multiple associations is actually valid. After you know this, you can probably discard the extraneous associations.

- 3 Run the Association Utility to identify and fix invalid associations (operation 3 and possibly operations 6 and 7).

As a general rule, if the problem is isolated, manually edit each invalid association. If the problem is repetitive and affects a large number of associations, consider using operations 6 and 7. This utility can replace bad identifiers on a global basis, but cannot insert or remove them where they do not already exist.

- 4 Run the Association Utility to normalize associations (operations 4 and 5).

16.4 Editing Associations

The Association Utility requires two parameters (oldRDN and newRDN) for operations 6 and 7, which search and replace.

The first value (for example, schema) in the parameter is the search criterion. The second value (for example, old) is the replacement value. Under certain scenarios, you can use the wildcard character * to generalize the search criterion or replacement value.

Three types of search and replace operations are possible:

Option	Description	Example
Replace the schema name	Replace schema <code>old</code> with schema <code>new</code> . Wildcards are supported on the right side only.	oldRDN: <code>schema=old</code> newRDN: <code>schema=new</code>
Replace the table name	Replace table <code>old</code> with table <code>new</code> . Wildcards are not supported.	oldRDN: <code>table=old</code> newRDN: <code>table=new</code>
Replace the column name	Replace column <code>old</code> with column <code>new</code> . Wildcards are required on the right side, but they aren't supported on the left side.	oldRDN: <code>old=*</code> newRDN: <code>new=*</code>

- ♦ [Section 17.1, “Recognizing Publication Events,” on page 215](#)
- ♦ [Section 17.2, “Executing Test Scripts,” on page 215](#)
- ♦ [Section 17.3, “Troubleshooting Driver Processes,” on page 215](#)

17.1 Recognizing Publication Events

Publication events might not be recognized by the Publisher channel unless you explicitly commit changes. For the commit keywords of supported databases, see [Section 14.3.7, “Commit Keywords,” on page 182](#).

17.2 Executing Test Scripts

The test scripts should be executed by a user other than the driver’s `idm` database user account. If you execute them as the `idm` user, events are ignored by the driver’s Publisher channel, unless publication loopback is allowed. For additional information on allowing or disallowing publication loopback, refer to [“Allow Loopback?” on page 94](#).

17.3 Troubleshooting Driver Processes

Viewing driver processes is necessary to analyze unexpected behavior. To view the driver processing events, use DSTrace. You should only use it during testing and troubleshooting the driver. Running DSTrace while the drivers are in production increases the utilization on the Identity Manager server and can cause events to process very slowly.

17.3.1 Viewing Driver Processes

To see the driver processes in DSTrace, values are added to the Driver Set object and the Driver object. You can do this in Designer or iManager.

- ♦ [“Adding Trace Levels in Designer” on page 215](#)
- ♦ [“Adding Trace Levels in iManager” on page 217](#)
- ♦ [“Capturing Driver Processes to a File” on page 218](#)


Adding Trace Levels in Designer

You can add trace levels to the Driver Set object or to each Driver object.

- ♦ [“Driver Set” on page 215](#)
- ♦ [“Driver” on page 216](#)

Driver Set

- 1 In an open project in Designer, select the Driver Set object in the *Outline* view.



- 2 Right-click, select *Properties*, then click 5. *Trace*.
- 3 Set the parameters for tracing, then click *OK*.

Parameter	Description
Driver trace level	As the Driver object trace level increases, the amount of information displayed in DSTrace increases. Trace level 1 shows errors, but not the cause of the errors. To see password synchronization information, set the trace level to 5.
XSL trace level	DSTrace displays XSL events. Set this trace level only when troubleshooting XSL style sheets. If you do not want to see XSL information, set the level to zero.
Java debug port	Allows developers to attach a Java* debugger.
Java trace file	When a value is set in this field, all Java information for the Driver Set object is written to a file. The value for this field is the path for that file. As long as the file is specified, Java information is written to this file. If you do not need to debug Java, leave this field blank.
Trace file size limit	Allows you to set a limit for the Java trace file. If you set the file size to <i>Unlimited</i> , the file grows in size until no disk space remains.

If you set the trace level on the Driver Set object, all drivers appear in the DSTrace logs.

Driver

- 1 In an open project in Designer, select the Driver object in the *Outline* view.
- 2 Right-click, select *Properties*, then click 8. *Trace*.
- 3 Set the parameters for tracing, then click *OK*.

Parameter	Description
Trace level	As the Driver object trace level increases, the amount of information displayed in DSTrace increases. Trace level 1 shows errors, but not the cause of the errors. To see password synchronization information, set the trace level to 5. if you select <i>Use setting from Driver Set</i> , the value is taken from the Driver Set object.

Parameter	Description
Trace file	Specify a filename and location for where the Identity Manager information is written for the selected driver. if you select <i>Use setting from Driver Set</i> , the value is taken from the Driver Set object.
Trace file size limit	Allows you to set a limit for the Java trace file. If you set the file size to <i>Unlimited</i> , the file grows in size until no disk space remains. If you select <i>Use setting from Driver Set</i> , the value is taken from the Driver Set object.
Trace name	The driver trace messages are prepended with the value entered instead of the driver name. Use this option if the driver name is very long.

If you set the parameters only on the Driver object, only information for that driver appears in the DSTrace log.

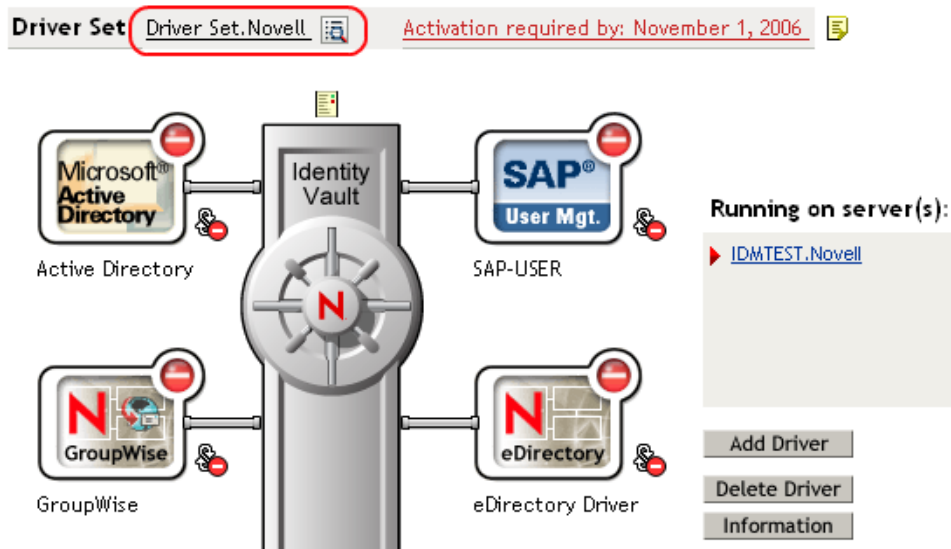
Adding Trace Levels in iManager

You can add trace levels to the Driver Set object or to each Driver object.

- ♦ “Driver Set” on page 217
- ♦ “Driver” on page 218

Driver Set

- 1 In iManager, select *Identity Manager > Identity Manager Overview*.
- 2 Browse to the Driver Set object, then click *Search*.
- 3 Click the driver set name.



- 4 Select the *Misc* tab for the Driver Set object.
- 5 Set the parameters for tracing, then click *OK*.

Driver

- 1 In iManager, select *Identity Manager > Identity Manager Overview*.
- 2 Browse to the Driver Set object where the Driver object resides, then click *Search*.
- 3 Click the upper right corner of the Driver object, then click *Edit properties*.
- 4 Select the *Misc* tab for the Driver object.
- 5 Set the parameters for tracing, then click *OK*.

NOTE: The option *Use setting from Driver Set* does not exist in iManager.

Capturing Driver Processes to a File

You can save driver processes to a file by using the parameter on the Driver object or by using DSTrace. The parameter on the Driver object is the *Trace file* parameter, under the *MISC* tab.

The driver processes that are captured through DSTrace are the processes that occur on the Identity Manager engine. If you use the Remote Loader, you need to capture a trace on the Remote Loader at the same time as you are capturing the trace on the Identity Manager engine.

The following methods help you capture and save Identity Manager processes through DSTrace on different platforms.

- ♦ [“NetWare” on page 218](#)
- ♦ [“Windows” on page 219](#)
- ♦ [“UNIX” on page 219](#)
- ♦ [“iMonitor” on page 219](#)
- ♦ [“Remote Loader” on page 220](#)

NetWare

Use `dstrace.nlm` to display trace messages on the system console or trace messages to a file (`sys:\system\dstrace.log`). Use `dstrace.nlm` to display the trace messages to a screen labeled DSTrace Console.

- 1 Enter `dstrace.nlm` at the server console to load `dstrace.nlm` into memory.
- 2 Enter `dstrace screen on` at the server console to allow trace messages to appear on the DSTrace Console screen.
- 3 Enter `dstrace file on` at the server console to capture trace messages sent to the DSTrace Console to the `dstrace.log` file.
- 4 (Optional) Enter `dstrace -all` at the server console to make it easier to read the trace log.
- 5 Enter `dstrace +dxml dstrace +dvrs` at the server console to display Identity Manager events.
- 6 Enter `dstrace +tags dstrace +time` at the server console to display message tags and time stamps.
- 7 Toggle to the DSTrace Console screen and watch for the event to pass.
- 8 Toggle back to the server console.
- 9 Enter `dstrace file off` at the server console.

This stops capturing trace messages to the log file. It also stops logging information into the file.

- 10 Open the `dstrace.log` in a text editor and search for the event or the object you modified.

Windows

- 1 Open the Control Panel, select *NDS Services* > `dstrace.dlm`, then click *Start* to display the NDS Server Trace utility window.
- 2 Click *Edit* > *Options*, then click *Clear All* to clear all of the default flags.
- 3 Select *DirXML* and *DirXML Drivers*.
- 4 Click OK.
- 5 Click *File* > *New*.
- 6 Specify the filename and location where you want the DSTrace information saved, then click *Open*.
- 7 Wait for the event to occur.
- 8 Click *File* > *Close*.

This stops the information from being written to the log file.

- 9 Open the file in a text editor and search for the event or the object you modified.

UNIX

- 1 Enter `ndstrace` to start the `ndstrace` utility.
- 2 Enter `set ndstrace=nodebug` to turn off all trace flags currently set.
- 3 Enter `set ndstrace on` to display trace messages to the console.
- 4 Enter `set ndstrace file on` to capture trace messages to the `ndstrace.log` file in the directory where eDirectory is installed. By default it is `/var/nds`.
- 5 Enter `set ndstrace=+dxml` to display the Identity Manager events.
- 6 Enter `set ndstrace=+dvrs` to display the Identity Manager driver events.
- 7 Wait for the event to occur.
- 8 Enter `set ndstrace file off` to stop logging information to the file.
- 9 Enter `exit` to quite the `ndstrace` utility.
- 10 Open the file in a text editor. Search for the event or the object that was modified.

iMonitor

iMonitor allows you to get DSTrace information from a Web browser. It does not matter where Identity Manager is running. The following files run iMonitor:

- ♦ `ndsimon.nlm` runs on NetWare®.
- ♦ `ndsimon.dlm` runs on Windows.
- ♦ `ndsimonitor` runs on UNIX*.

- 1 Access iMonitor from `http://server_ip:8008/nds`.

Port 8008 is the default.

- 2 Specify a username and password with administrative rights, then click *Login*.

- 3 Select *Trace Configuration* on the left side.
- 4 Click *Clear All*.
- 5 Select *DirXML* and *DirXML Drivers*.
- 6 Click *Trace On*.
- 7 Select *Trace History* on the left side.
- 8 Click the document with the *Modification Time* of *Current* to see a live trace.
- 9 Change the *Refresh Interval* if you want to see information more often.
- 10 Select *Trace Configuration* on the left side, then click *Trace Off* to turn the tracing off.
- 11 Select *Trace History* to view the trace history.

The files are distinguished by their time stamp.

If you need a copy of the HTML file, the default location is:

- ◆ NetWare: `sys:\system\ndsimon\dstrace*.htm`
- ◆ Windows: `Drive_letter:\novell\nds\ndsimon\dstrace*.htm`
- ◆ UNIX: `/var/nds/dstrace/*.htm`

Remote Loader

You can capture the events that occur on the machine running the Remote Loader service.

- 1 Launch the Remote Loader Console by clicking the icon.
- 2 Select the driver instance, then click *Edit*.
- 3 Set the *Trace Level* to 3 or above.
- 4 Specify a location and file for the trace file.
- 5 Specify the amount of disk space that the file is allowed.
- 6 Click *OK* twice to save the changes.

You can also enable tracing from the command line by using the following switches. For more information, see “[Configuring the Remote Loader](#)” in the *Novell Identity Manager 3.5.1 Administration Guide*.

Table 17-1 *Command Line Tracing Switches*

Option	Short Name	Parameter	Description
-trace	-t	integer	Specifies the trace level. This is used only when hosting an application shim. Trace levels correspond to those used on the Identity Manager server. Example: <code>-trace 3</code> or <code>-t3</code>
-tracefile	-tf	filename	Specify a file to write trace messages to. Trace messages are written to the file if the trace level is greater than zero. Trace messages are written to the file even if the trace window is not open. Example: <code>-tracefile c:\temp\trace.txt</code> or <code>-tf c:\temp\trace.txt</code>

Option	Short Name	Parameter	Description
-tracefilemax	-tfm	size	<p>Specifies the approximate maximum size that trace file data can occupy on disk. If you specify this option, Identity Manager creates a trace file with the name specified by using the tracefile option and up to 9 additional “roll-over” files. The roll-over files are named by using the base of the main trace filename plus “_n”, where n is 1 through 9.</p> <p>The size parameter is the number of bytes. Specify the size by using the suffixes K, M, or G for kilobytes, megabytes, or gigabytes.</p> <p>If the trace file data is larger than the specified maximum when the Remote Loader is started, the trace file data remains larger than the specified maximum until roll-over is completed through all 10 files.</p> <p>Example: <code>-tracefilemax 1000M</code> or <code>-tfm 1000M</code></p>

Backing Up the JDBC Driver

18

You can use Designer for Identity Manager or iManager to create an XML file of the driver. The file contains all of the information that you entered into the driver during configuration. If the driver becomes corrupted, you can restore the configuration information by importing the exported file.

IMPORTANT: If the driver has been deleted, all of the associations on the objects are purged. When you import the XML file, the migration process creates new associations.

Not all server-specific information stored on the driver is contained in the XML file. Make sure that this information is documented through the Document Generation process in Designer. See “[Documenting Projects](#)” in the *Designer 2.1 for Identity Manager 3.5.1* guide.

- ♦ [Section 18.1, “Exporting the Driver in Designer,” on page 223](#)
- ♦ [Section 18.2, “Exporting the Driver in iManager,” on page 223](#)

18.1 Exporting the Driver in Designer

- 1 Open a project in Designer, then right-click the Driver object.
- 2 Select *Export to Configuration File*.
- 3 Specify a unique name for the configuration file, browse to location where it should be saved, then click *Save*.
- 4 Click *OK* in the Export Configuration Results window.

18.2 Exporting the Driver in iManager

- 1 In iManager, select *Identity Manager > Identity Manager Overview*.
- 2 Browse to and select the Driver Set object, then click *Search*.
- 3 Click the driver icon.
- 4 Select *Export* in the Identity Manager Driver Overview window.
- 5 Browse to and select the Driver object that you want to export, then click *Next*.
- 6 Select *Export all policies, linked to the configuration or not* or select *Only export policies that are linked to the configuration*, depending upon the information you want to have stored in the XML file.
- 7 Click *Next*.
- 8 Click *Save As*, then click *Save*.
- 9 Browse and select a location to save the XML file, then click *Save*.
- 10 Click *Finish*.

Best Practices

A

The following section lists important best practices for using the Driver for JDBC. You can find additional information in [Chapter 6, “Configuring the JDBC Driver,” on page 55](#).

Security/Performance:

- ◆ For performance and security reasons, run the driver remotely on the database server whenever possible. Be sure to enable SSL encryption between the Identity Vault and the Remote Loader service.
- ◆ You should enable SSL encryption for third-party drivers whenever the Driver for JDBC is not running remotely on the database server. For information on the security capabilities of supported third-party drivers, see [“Third-Party JDBC Drivers” on page 189](#).
- ◆ In a production environment, turn off tracing.

Other:

- ◆ For direct synchronization, prefix one or more view column names with “pk_” (case-insensitive).
- ◆ For both direct and indirect synchronization, use different primary key column names between logical database classes.
- ◆ Delimit (double-quote) primary key values placed in the event log `table_key` field if they contain the following characters:

`, ; ' + = \ " < >`

This caution is usually an issue only if the primary key column is a binary type.

- ◆ When an Identity Vault is the authoritative source of primary key values, GUID rather than CN is recommended for use as a primary key. Unlike CN, GUID is single-valued and does not change.
- ◆ Omit from publication triggers foreign key columns that link child and parent tables.
- ◆ If primary key columns are static (they do not change), do not include them in publication triggers.
- ◆ Place the `jdbc:type="query"` attribute value on all embedded SELECT statements. Place the `jdbc:type="update"` attribute value on all embedded INSERT, UPDATE and DELETE statements.

- ◆ [Section B.1, “Can’t See Tables or Views,” on page 227](#)
- ◆ [Section B.2, “Synchronizing with Tables,” on page 227](#)
- ◆ [Section B.3, “Processing Rows in the Event Log Table,” on page 228](#)
- ◆ [Section B.4, “Managing Database User Accounts,” on page 228](#)
- ◆ [Section B.5, “Synchronizing Large Data Types,” on page 228](#)
- ◆ [Section B.6, “Slow Publication,” on page 228](#)
- ◆ [Section B.7, “Synchronizing Multiple Classes,” on page 229](#)
- ◆ [Section B.8, “Encrypted Transport,” on page 229](#)
- ◆ [Section B.9, “Mapping Multivalued Attributes,” on page 229](#)
- ◆ [Section B.10, “Synchronizing Garbage Strings,” on page 229](#)
- ◆ [Section B.11, “Running Multiple Driver for JDBC Instances,” on page 229](#)

B.1 Can’t See Tables or Views

Question: Why can’t the driver see my tables or views?

Answer: The driver is capable of synchronizing only tables that have explicit primary key constraints and views that contain one or more columns prefixed with “pk_” (case-insensitive). The driver uses these constraints to determine which fields to use when constructing associations. As such, the driver ignores any unconstrained tables.

If you are trying to synchronize with tables or views that lack the necessary constraints, either add them or synchronize to intermediate tables with the required constraints.

Another possibility is that the driver lacks the necessary database privileges to see the tables. Usually, visibility is determined by the presence or absence of the `SELECT` privilege.

B.2 Synchronizing with Tables

Question: How do I synchronize with tables located in multiple schemas?

Answer: Do one of the following:

- ◆ Alias the tables into the synchronization schema.
- ◆ Synchronize to intermediate tables in the synchronization schema and move the data across schema boundaries.
- ◆ Use a view.
- ◆ Create a virtual schema by using the Table/View Names parameter.
See [“Table/View Names” on page 68](#).

B.3 Processing Rows in the Event Log Table

Question: Why isn't the driver processing rows in the Event Log Table?

Answer: Do the following:

- 1 Check the `perpetrator` field of the rows in question and make sure that the value is set to something other than the driver's database username.

The Publisher channel checks the `perpetrator` field to detect loopback events if the Publisher channel Allow Loopback parameter is set to Boolean False (the default). See [“Allow Loopback?” on page 94](#).

When the Allow Loopback parameter is set to Boolean False, the Publisher channel ignores all records where the `perpetrator` field value is equal to the driver's database username. The driver's database username is specified by using the Authentication ID parameter. See [“Authentication ID” on page 58](#).

- 2 Ensure that the record's `status` field is set to N (new).
Records with `status` fields set to something other than N will not be processed.
- 3 Make sure to explicitly commit changes.
Changes are often tentative until explicitly committed.

B.4 Managing Database User Accounts

Question: Can the driver manage database user accounts?

Answer: Yes. You can manage database accounts by using embedded SQL. For more information, see [Chapter 13, “Embedded SQL Statements in XDS Events,” on page 153](#).

B.5 Synchronizing Large Data Types

Question: Can the driver synchronize large binary and string data types?

Answer: Yes. Large binary and string data types can be subscribed and published. Publish large binary and string data types by using query-back event types. For additional information, see [Section 12.2, “Event Types,” on page 145](#).

B.6 Slow Publication

Question: Why is publication slow?

Answer: If the event log table contains a large number of rows, index the table. Example indexes are provided in all database installation scripts. By using trace level 3, you can view the statements that the driver uses to maintain the event log.

You can further refine indexes in the installation scripts to enhance publication performance. Placing indexes in a different tablespace or physical disk than the event log table also enhances publication performance.

Furthermore, in a production environment, set the Delete Processed Rows parameter to Boolean False, unless processed rows are being periodically moved to another table. See [“Delete Processed Rows?” on page 93](#).

B.7 Synchronizing Multiple Classes

Question: Can the driver synchronize multiple classes?

Answer: Yes. However, primary key column names must be unique between logical database classes. For example, if *class1* is mapped to *table1* with primary key column name *key1*, and *class2* is mapped to *table2* with primary key column name *key2*, then the name of *key1* cannot equal *key2*.

This requirement can always be satisfied, no matter which synchronization model is employed.

B.8 Encrypted Transport

Question: Does the driver support encrypted transport?

Answer: No. How the driver communicates with a given database depends upon the third-party driver being used. Some third-party drivers support encrypted transport, while others do not. Even if encrypted transport is supported, no standardized way exists to enable encryption between third-party JDBC drivers.

The general solution for this problem is to remotely run the Driver for JDBC and your third-party driver. This method allows both the Driver for JDBC and the third-party driver to run locally on the database server. Then all data traveling across the network between the Metadirectory engine and the Driver for JDBC are SSL encrypted.

Another possibility is to use a type 3 or type 2 third-party JDBC driver. Database middleware and client APIs usually provide encrypted transport mechanisms.

B.9 Mapping Multivalue Attributes

Question: How do I map multivalue attributes to single-value database fields?

Answer: See [Section 10.7, “Mapping Multivalue Attributes to Single-Value Database Fields,”](#) on [page 139](#).

B.10 Synchronizing Garbage Strings

Question: Why is the driver synchronizing garbage strings?

Answer: The database and the third-party driver are probably using incompatible character encoding. Adjust the character encoding that your third-party driver uses.

For more information, refer to [Character Encoding Values \(http://java.sun.com/j2se/1.5.0/docs/guide/intl/encoding.doc.html\)](http://java.sun.com/j2se/1.5.0/docs/guide/intl/encoding.doc.html), defined by Sun.

B.11 Running Multiple Driver for JDBC Instances

Question: How do I run multiple Driver for JDBC instances in the same driver set? The instances require different versions of the same third-party JDBC driver (for example, the Oracle JDBC driver or the IBM DB2 Type 3 JDBC driver).

Answer: Use the Remote Loader to load each Driver for JDBC instance in a separate Java Virtual Machine (JVM). When run locally in the same JVM, different versions of the same third-party classes collide.

Supported Data Types

C

The driver for JDBC can synchronize all JDBC 1 data types and a small subset of JDBC 2 data types. How JDBC data types map to a database's native data types depends on the third-party driver.

The following list includes the supported JDBC 1 `java.sql.Types` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>).

- ◆ Numeric Types:
 - ◆ `java.sql.Types.BIGINT`
 - ◆ `java.sql.Types.BIT`
 - ◆ `java.sql.Types.DECIMAL`
 - ◆ `java.sql.Types.DOUBLE`
 - ◆ `java.sql.Types.NUMERIC`
 - ◆ `java.sql.Types.REAL`
 - ◆ `java.sql.Types.FLOAT`
 - ◆ `java.sql.Types.INTEGER`
 - ◆ `java.sql.Types.SMALLINT`
 - ◆ `java.sql.Types.TINYINT`
- ◆ String Types:
 - ◆ `java.sql.Types.CHAR`
 - ◆ `java.sql.Types.LONGCHAR`
 - ◆ `java.sql.Types.VARCHAR`
- ◆ Time Types:
 - ◆ `java.sql.Types.DATE`
 - ◆ `java.sql.Types.TIME`
 - ◆ `java.sql.Types.TIMESTAMP`
- ◆ Binary Types:
 - ◆ `java.sql.Types.BINARY`
 - ◆ `java.sql.Types.VARBINARY`
 - ◆ `java.sql.Types.LONGVARBINARY`

The following list includes the supported JDBC 2 `java.sql.Types` (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Types.html>).

- ◆ Large Object (LOB) Types:
 - ◆ `java.sql.Types.CLOB`
 - ◆ `java.sql.Types.BLOB`

java.sql.DatabaseMetaData Methods

D

This section lists the required and optional [java.sql.DatabaseMetaData](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DatabaseMetaData.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DatabaseMetaData.html>) methods.

The following JDBC 1 methods are required only if the **Synchronization Filter** parameter is set to something other than *Exclude all tables/views*:

- ◆ `getColumns(java.lang.String catalog, java.lang.String schemaPattern, java.lang.String tableNamePattern, java.lang.String columnNamePattern):java.sql.ResultSet`
- ◆ `getPrimaryKeys(java.lang.String catalog, java.lang.String schema, java.lang.String table):java.sql.ResultSet`
- ◆ `getTables(java.lang.String catalog, java.lang.String schemaPattern, java.lang.String tableNamePattern, java.lang.String[] types):java.sql.ResultSet`
- ◆ `storesLowerCaseIdentifiers():boolean`
- ◆ `storesMixedCaseIdentifiers():boolean`
- ◆ `storesUpperCaseIdentifiers():boolean`

Optional JDBC 1 methods:

- ◆ `dataDefinitionCausesTransactionCommit():boolean`
- ◆ `dataDefinitionIgnoredInTransactions():boolean`
- ◆ `getColumnPrivileges(String catalog, String schema, String table, String columnNamePattern):java.sql.ResultSet`
- ◆ `getDatabaseProductName():java.lang.String`
- ◆ `getDatabaseProductVersion():java.lang.String`
- ◆ `getDriverMajorVersion():int`
- ◆ `getDriverMinorVersion():int`
- ◆ `getDriverName():java.lang.String`
- ◆ `getDriverVersion():java.lang.String`
- ◆ `getExportedKeys(java.lang.String catalog, java.lang.String schema, java.lang.String table):java.sql.ResultSet`
- ◆ `getMaxStatements():int`
- ◆ `getMaxConnections():int`
- ◆ `getMaxColumnsInSelect():int`
- ◆ `getProcedureColumns(String catalog, String schemaPattern, String procedureNamePattern, String columnNamePattern):java.sql.ResultSet`
- ◆ `getSchemas():java.sql.ResultSet`
- ◆ `getTableTypes():java.sql.ResultSet`
- ◆ `getUserName():java.lang.String`

- ◆ supportsColumnAliasing():boolean
- ◆ supportsDataDefinitionAndDataManipulationTransactions():boolean
- ◆ supportsDataManipulationTransactionsOnly():boolean
- ◆ supportsLimitedOuterJoins():boolean
- ◆ supportsMultipleTransactions():boolean
- ◆ supportsSchemasInDataManipulation():boolean
- ◆ supportsSchemasInProcedureCalls():boolean
- ◆ supportsTransactionIsolationLevel(int level):boolean
- ◆ supportsTransactions():boolean

Optional JDBC 2 methods:

- ◆ supportsBatchUpdates():boolean

Optional JDBC 3 methods:

- ◆ supportsGetGeneratedKeys():boolean

JDBC Interface Methods

E

This section lists the JDBC interface methods (other than [java.sql.DatabaseMetaData](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DatabaseMetaData.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DatabaseMetaData.html>) methods) that the driver for JDBC uses. Methods are organized by class.

Often, third-party JDBC driver vendors list defects or known issues by method. You can use the following methods in collaboration with third-party JDBC driver documentation to troubleshoot or anticipate potential interoperability problems.

- ♦ [java.sql.DriverManager](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DriverManager.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DriverManager.html>)
- ♦ [java.sql.CallableStatement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/CallableStatement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/CallableStatement.html>)
- ♦ [java.sql.Connection](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html>)
- ♦ [java.sql.PreparedStatement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/PreparedStatement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/PreparedStatement.html>)
- ♦ [java.sql.ResultSet](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSet.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSet.html>)
- ♦ [java.sql.ResultSetMetaData](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSetMetaData.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSetMetaData.html>)
- ♦ [java.sql.Statement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html>)
- ♦ [java.sql.Timestamp](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html>)

The following table lists [java.sql.DriverManager](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DriverManager.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/DriverManager.html>) methods that the Driver for JDBC uses:

Table E-1 *java.sql.DriverManager Methods*

Method Signature	JDBC Version	Required?
<code>getConnection(String url, java.util.Properties info):java.sql.Connection</code>	1	yes ¹
<code>getConnection(String url, java.util.Properties info):java.sql.Connection</code>	1	yes ¹
<code>setLogStream(java.io.PrintStream out):void</code>	1	no

¹One method or the other.

The following table lists [java.sql.CallableStatement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/CallableStatement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/CallableStatement.html>) methods that the Driver for JDBC uses:

Table E-2 *java.sql.CallableStatement Methods*

Method Signature	JDBC Version	Required?
<code>getBigDecimal(int parameterIndex, int scale):java.math.BigDecimal</code>	1	yes

Method Signature	JDBC Version	Required?
getBoolean(int parameterIndex):boolean	1	yes
getBoolean(String parameterName):boolean	3	no
getBytes(int parameterIndex):byte	1	yes
getBytes(String parameterName):byte	3	no
getBytes(int parameterIndex):byte[]	1	yes
getBytes(String parameterName):byte[]	3	no
getDate(int parameterIndex):java.sql.Date	1	yes
getDate(String parameterName):java.sql.Date	3	no
getDouble(int parameterIndex):double	1	yes
getDouble(String parameterName):double	3	no
getFloat(int parameterIndex):float	1	yes
getFloat(String parameterName):float	3	no
getInt(int parameterIndex):int	1	yes
int getInt(String parameterName)	3	no
getLong(int parameterIndex):long	1	yes
getLong(String parameterName):long	3	no
getShort(int parameterIndex):short	1	yes
getShort(String parameterName):short	3	no
getString(int parameterIndex):String	1	yes
getString(String parameterName):String	3	no
getTime(int parameterIndex):java.sql.Time	1	yes
getTime(String parameterName):java.sql.Time	3	no
getTimestamp(int parameterIndex):java.sql.Timestamp	1	yes
getTimestamp(String parameterName):java.sql.Timestamp	3	no
registerOutParameter(int parameterIndex, int sqlType):void	1	yes
wasNull():boolean	1	yes

The following table lists [java.sql.Connection](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Connection.html>) methods that the driver for JDBC uses:

Table E-3 *java.sql.Connection Methods*

Method Signature	JDBC Version	Required?
close():void	1	yes
commit():void	1	no
createStatement():java.sql.Statement	1	yes
getAutoCommit():boolean	1	no
getMetaData():java.sql.DatabaseMetaData	1	yes
getTransactionIsolation():int	1	no
getWarnings():java.sql.SQLWarning	1	no
isClosed():boolean	1	no
prepareCall(String sql):java.sql.CallableStatement	1	no
prepareStatement(String sql):java.sql.PreparedStatement	1	yes
rollback():void	1	no
setAutoCommit(boolean autoCommit):void	1	no
setTransactionIsolation(int level):void	1	no

The following table lists [java.sql.PreparedStatement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/PreparedStatement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/PreparedStatement.html>) methods that the Driver for JDBC uses:

Table E-4 *java.sql.PreparedStatement Methods*

Method Signature	JDBC Version	Required?
clearParameters() :void	1	no
execute():boolean	1	yes
executeQuery():java.sql.ResultSet	1	yes
executeUpdate():int	1	yes
setBigDecimal(int parameterIndex, java.math.BigDecimal x):void	1	yes
setBoolean(int parameterIndex, boolean x):void	1	yes
setByte(int parameterIndex, byte x):void	1	yes
setBytes(int parameterIndex, byte x[]):void	1	yes
setDate(int parameterIndex, java.sql.Date x):void	1	yes
setDouble(int parameterIndex, double x):void	1	yes
setFloat(int parameterIndex, float x):void	1	yes
setInt(int parameterIndex, int x):void	1	yes
setLong(int parameterIndex, long x):void	1	yes

Method Signature	JDBC Version	Required?
setNull(int parameterIndex, int sqlType):void	1	yes
setShort(int parameterIndex, short x):void	1	yes
setString(int parameterIndex, String x):void	1	yes
setTime(int parameterIndex, java.sql.Time x):void	1	yes
setTimestamp(int parameterIndex, java.sql.Timestamp x):void	1	yes

The following table lists [java.sql.ResultSet](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSet.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSet.html>) methods that the Driver for JDBC uses:

Table E-5 *java.sql.ResultSet Methods*

Method Signature	JDBC Version	Required?
close():void	1	yes
getBigDecimal(int columnIndex, int scale):java.math.BigDecimal	1	yes
getBigDecimal(String columnName, int scale):java.math.BigDecimal	1	yes
getBinaryStream(int columnIndex):java.io.InputStream	1	yes
getBinaryStream(String columnName)java.io.InputStream	1	yes
getBoolean(int columnIndex):boolean	1	yes
getBoolean(String columnName):boolean	1	yes
getByte(int columnIndex):byte	1	yes
getByte(String columnName):byte	1	yes
getBytes(int columnIndex):byte[]	1	yes
getBytes(String columnName):byte[]	1	yes
getDate(int columnIndex):java.sql.Date	1	yes
getDate(String columnName)java.sql.Date	1	yes
getFloat(int columnIndex):float	1	yes
getFloat(String columnName):float	1	yes
getInt(int columnIndex):int	1	yes
getInt(String columnName):int	1	yes
getLong(int columnIndex):long	1	yes
getLong(String columnName):long	1	yes
getMetaData():java.sql.ResultSetMetaData	1	no
getShort(int columnIndex):short	1	yes
getShort(String columnName):short	1	yes

Method Signature	JDBC Version	Required?
getString(int columnIndex):String	1	yes
getString(String columnName):String	1	yes
getTime(int columnIndex):java.sql.Time	1	yes
getTime(String columnName):java.sql.Time	1	yes
getTimestamp(int columnIndex):java.sql.Timestamp	1	yes
getTimestamp(String columnName):java.sql.Timestamp	1	yes
getWarnings():java.sql.SQLWarning	1	no

The following table lists [java.sql.ResultSetMetaData](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSetMetaData.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/ResultSetMetaData.html>) methods that the Driver for JDBC uses:

Table E-6 *java.sql.ResultSetMetaData Methods*

Method Signature	JDBC Version	Required?
getColumnCount():int	1	yes
columnName(int column):String	1	no
getColumnType(int column):int	1	no

The following table lists [java.sql.Statement](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Statement.html>) methods that the Driver for JDBC uses:

Table E-7 *java.sql.Statement Methods*

Method Signature	JDBC Version	Required?
addBatch(java.lang.String sql):void	2	no
clearBatch():void	2	no
clearWarnings():void	1	no
close():void	1	yes
execute(java.lang.String sql):boolean	1	yes
executeBatch():int[]	2	no
executeUpdate(String sql):int	1	yes
executeQuery(String sql):java.sql.ResultSet	1	yes
getGeneratedKeys():java.sql.ResultSet	3	no
getMoreResults():boolean	1	no
getResultSet():java.sql.ResultSet	1	yes
getUpdateCount():int	1	no

Method Signature	JDBC Version	Required?
getWarnings():java.sql.SQLWarning	1	no

The following table lists [java.sql.Timestamp](http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html) (<http://java.sun.com/j2se/1.5.0/docs/api/java/sql/Timestamp.html>) methods that the Driver for JDBC uses:

Table E-8 *java.sql.Timestamp Methods*

Method Signature	JDBC Version	Required?
getNanos():int	1	yes
getTime():long	1	yes
setNanos(int n):void	1	yes
setTime(long time):void	1	yes
toString ():String	1	yes

Third-Party JDBC Driver Descriptor DTD

F

This section contains the DTD for third-party JDBC descriptor files.

```
<?xml version="1.0" encoding="UTF-8"?>
<!ELEMENT actions (exec-sql | check-for-closed-connection | fetch
metadata | rollback)*>
<!ELEMENT add-default-values-on-view-insert (#PCDATA)>
<!ELEMENT authentication (regular-expression | sql-state | error-code
| sql-state-class | error-code-range | actions)*>
<!ELEMENT check-for-closed-connection EMPTY>
<!ELEMENT column-position-comparator (#PCDATA)>
<!ELEMENT connection-properties (property)*>
<!ELEMENT connectivity (regular-expression | sql-state | error-code |
sql-state-class | error-code-range | actions)*>
<!ELEMENT current-timestamp-stmt (#PCDATA)>
<!ELEMENT error-code (value)>
<!ATTLIST error-code
  description CDATA #IMPLIED
>
<!ELEMENT error-code-range (from, to)>
<!ATTLIST error-code-range
  description CDATA #IMPLIED
>
<!ELEMENT errors (connectivity | authentication | retry | fatal)*>
<!ELEMENT exclude-table-filter (#PCDATA)>
<!ELEMENT exec-sql (#PCDATA)>
<!ELEMENT fatal (regular-expression | sql-state | error-code | sql
state-class | error-code-range | actions)*>
<!ELEMENT fetch-metadata EMPTY>
<!ELEMENT from (#PCDATA)>
<!ELEMENT function-return-method (#PCDATA)>
<!ELEMENT handle-stmt-results (#PCDATA)>
<!ELEMENT identity (name?, target-database?, jdbc-type?, jdbc-class?)>
<!ELEMENT import (#PCDATA)>
<!ELEMENT imports (import)*>
<!ELEMENT include-table-filter (#PCDATA)>
<!ELEMENT jdbc-class (#PCDATA)>
<!ELEMENT jdbc-driver (imports?, identity, (metadata-override |
connection-properties | sql-type-map | options | errors)*>
<!ELEMENT jdbc-type (#PCDATA)>
<!ELEMENT key (#PCDATA)>
<!ELEMENT left-outer-join-operator (#PCDATA)>
<!ELEMENT lock-generator-class (#PCDATA)>
<!ELEMENT metadata-override (supports-schemas-in-procedure-calls?)>
<!ELEMENT minimal-metadata (#PCDATA)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT options (lock-generator-class | supports-schemas-in
```

```

metadata-retrieval | time-translator-class | column-position
comparator | use-manual-transactions | minimal-metadata | transaction
isolation-level | use-single-connection | exclude-table-filter |
include-table-filter | left-outer-join-operator | current-timestamp
stmt | add-default-values-on-view-insert | reuse-statements |
function-return-method | handle-stmt-results)*>
<!ELEMENT property (key, value)>
<!ELEMENT regular-expression (value)>
<!ELEMENT retry (regular-expression | sql-state | error-code | sql
state-class | error-code-range | actions)*>
<!ELEMENT reuse-statements (#PCDATA)>
<!ELEMENT rollback EMPTY>
<!ELEMENT sql-state (value)>
<!ATTLIST sql-state
  description CDATA #IMPLIED
>
<!ELEMENT sql-state-class (value)>
<!ATTLIST sql-state-class
  description CDATA #IMPLIED
>
<!ELEMENT sql-type-map (type*)>
<!ELEMENT supports-schemas-in-metadata-retrieval (#PCDATA)>
<!ELEMENT supports-schemas-in-procedure-calls (#PCDATA)>
<!ELEMENT target-database (#PCDATA)>
<!ELEMENT time-translator-class (#PCDATA)>
<!ELEMENT to (#PCDATA)>
<!ELEMENT transaction-isolation-level (#PCDATA)>
<!ELEMENT type (from, to)>
<!ELEMENT use-manual-transactions (#PCDATA)>
<!ELEMENT use-single-connection (#PCDATA)>
<!ELEMENT value (#PCDATA)>

```

Third-Party JDBC Driver Descriptor Import DTD



This section contains the DTD for third-party JDBC descriptor import files.

```
<?xml version="1.0" encoding="UTF-8"?>
<!ELEMENT actions (exec-sql | check-for-closed-connection | fetch-
metadata | rollback)*>
<!ELEMENT add-default-values-on-view-insert (#PCDATA)>
<!ELEMENT authentication (regular-expression | sql-state | error-code
| sql-state-class | error-code-range | actions)*>
<!ELEMENT check-for-closed-connection EMPTY>
<!ELEMENT column-position-comparator (#PCDATA)>
<!ELEMENT connection-properties (property)*>
<!ELEMENT connectivity (regular-expression | sql-state | error-code |
sql-state-class | error-code-range | actions)*>
<!ELEMENT current-timestamp-stmt (#PCDATA)>
<!ELEMENT error-code (value)>
<!ATTLIST error-code
  description CDATA #IMPLIED
>
<!ELEMENT error-code-range (from, to)>
<!ATTLIST error-code-range
  description CDATA #IMPLIED
>
<!ELEMENT errors (connectivity | authentication | retry | fatal)*>
<!ELEMENT exclude-table-filter (#PCDATA)>
<!ELEMENT exec-sql (#PCDATA)>
<!ELEMENT fatal (regular-expression | sql-state | error-code | sql-
state-class | error-code-range | actions)*>
<!ELEMENT fetch-metadata EMPTY>
<!ELEMENT from (#PCDATA)>
<!ELEMENT function-return-method (#PCDATA)>
<!ELEMENT handle-stmt-results (#PCDATA)>
<!ELEMENT include-table-filter (#PCDATA)>
<!ELEMENT jdbc-driver (metadata-override | connection-properties |
sql-type-map | options | errors)*>
<!ELEMENT key (#PCDATA)>
<!ELEMENT left-outer-join-operator (#PCDATA)>
<!ELEMENT lock-generator-class (#PCDATA)>
<!ELEMENT metadata-override (supports-schemas-in-procedure-calls?)>
<!ELEMENT minimal-metadata (#PCDATA)>
<!ELEMENT options (lock-generator-class | supports-schemas-in-
metadata-retrieval | time-translator-class | column-position-
comparator | use-manual-transactions | minimal-metadata | transaction-
isolation-level | use-single-connection | exclude-table-filter |
include-table-filter | left-outer-join-operator | current-timestamp-
stmt | add-default-values-on-view-insert | reuse-statements |
function-return-method | handle-stmt-results)*>
<!ELEMENT property (key, value)>
```

```

<!ELEMENT regular-expression (value)>
<!ELEMENT retry (regular-expression | sql-state | error-code | sql-
state-class | error-code-range | actions)*>
<!ELEMENT reuse-statements (#PCDATA)>
<!ELEMENT rollback EMPTY>
<!ELEMENT sql-state (value)>
<!ATTLIST sql-state
  description CDATA #IMPLIED
>
<!ELEMENT sql-state-class (value)>
<!ATTLIST sql-state-class
  description CDATA #IMPLIED
>
<!ELEMENT sql-type-map (type*)>
<!ELEMENT supports-schemas-in-metadata-retrieval (#PCDATA)>
<!ELEMENT supports-schemas-in-procedure-calls (#PCDATA)>
<!ELEMENT time-translator-class (#PCDATA)>
<!ELEMENT to (#PCDATA)>
<!ELEMENT transaction-isolation-level (#PCDATA)>
<!ELEMENT type (from, to)>
<!ELEMENT use-manual-transactions (#PCDATA)>
<!ELEMENT use-single-connection (#PCDATA)>
<!ELEMENT value (#PCDATA)>

```

Database Descriptor DTD



This section contains the DTD for database descriptor files.

```
<?xml version="1.0" encoding="UTF-8"?>
<!ELEMENT add-default-values-on-view-insert (#PCDATA)>
<!ELEMENT column-position-comparator (#PCDATA)>
<!ELEMENT current-timestamp-stmt (#PCDATA)>
<!ELEMENT database (imports?, identity, options?)>
<!ELEMENT exclude-table-filter (#PCDATA)>
<!ELEMENT function-return-method (#PCDATA)>
<!ELEMENT handle-stmt-results (#PCDATA)>
<!ELEMENT include-table-filter (#PCDATA)>
<!ELEMENT identity (name?, regex-name?, regex-version?)>
<!ELEMENT import (#PCDATA)>
<!ELEMENT imports (import*)>
<!ELEMENT left-outer-join-operator (#PCDATA)>
<!ELEMENT lock-generator-class (#PCDATA)>
<!ELEMENT minimal-metadata (#PCDATA)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT options (lock-generator-class | supports-schemas-in-
metadata-retrieval | time-translator-class | column-position-
comparator | use-manual-transactions | minimal-metadata | transaction-
isolation-level | use-single-connection | exclude-table-filter |
include-table-filter | left-outer-join-operator | current-timestamp-
stmt | add-default-values-on-view-insert | reuse-statements |
function-return-method | handle-stmt-results)*>
<!ELEMENT regex-name (#PCDATA)>
<!ELEMENT regex-version (#PCDATA)>
<!ELEMENT reuse-statements (#PCDATA)>
<!ELEMENT supports-schemas-in-metadata-retrieval (#PCDATA)>
<!ELEMENT time-translator-class (#PCDATA)>
<!ELEMENT transaction-isolation-level (#PCDATA)>
<!ELEMENT use-manual-transactions (#PCDATA)>
<!ELEMENT use-single-connection (#PCDATA)>
```


Database Descriptor Import DTD

This section contains the DTD for database descriptor import files.

```
<?xml version="1.0" encoding="UTF-8"?>
<!ELEMENT add-default-values-on-view-insert (#PCDATA)>
<!ELEMENT column-position-comparator (#PCDATA)>
<!ELEMENT current-timestamp-stmt (#PCDATA)>
<!ELEMENT exclude-table-filter (#PCDATA)>
<!ELEMENT function-return-method (#PCDATA)>
<!ELEMENT handle-stmt-results (#PCDATA)>
<!ELEMENT include-table-filter (#PCDATA)>
<!ELEMENT database (options?)>
<!ELEMENT left-outer-join-operator (#PCDATA)>
<!ELEMENT lock-generator-class (#PCDATA)>
<!ELEMENT minimal-metadata (#PCDATA)>
<!ELEMENT options (lock-generator-class | supports-schemas-in-
metadata-retrieval | time-translator-class | column-position-
comparator | use-manual-transactions | minimal-metadata | transaction-
isolation-level | use-single-connection | exclude-table-filter |
include-table-filter | left-outer-join-operator | current-timestamp-
stmt | add-default-values-on-view-insert | reuse-statements |
function-return-method | handle-stmt-results)*>
<!ELEMENT reuse-statements (#PCDATA)>
<!ELEMENT supports-schemas-in-metadata-retrieval (#PCDATA)>
<!ELEMENT time-translator-class (#PCDATA)>
<!ELEMENT transaction-isolation-level (#PCDATA)>
<!ELEMENT use-manual-transactions (#PCDATA)>
<!ELEMENT use-single-connection (#PCDATA)>
```


Policy Example: Triggerless Future Event Processing



The following example assumes that a “commence” attribute exists and does the following:

- ◆ Holds the timestamp value of when an event should be processed
- ◆ Contains an integer or java string timestamp value. See [“Time Syntax” on page 61](#).

```
<policy xmlns:Timestamp="http://www.novell.com/nxsl/java/java.sql.Timestamp"
        xmlns:TimestampUtil="http://www.novell.com/nxsl/java/com.novell.nds.dirxml.driver.jdbc.db.TimestampUtil"
        xmlns:jdbc="urn:dirxml:jdbc">
<rule>
<description>Get commencement date from datasource.</description>
  <conditions>
    <and>
      <if-xpath op="true">.</if-xpath>
    </and>
  </conditions>
  <actions>
    <do-set-local-variable name="commence">
      <arg-string>
        <token-src-attr class-name="User"
name="commence"/>
      </arg-string>
    </do-set-local-variable>
  </actions>
</rule>
<rule>
  <description>Break if commencement date unavailable.</
description>
  <conditions>
    <and>
      <if-local-variable name="commence" op="equal"/>
    </and>
  </conditions>
  <actions>
    <do-break/>
  </actions>
</rule>
<rule>
<description>Parse times.</description>
  <conditions>
    <and>
      <if-xpath op="true">.</if-xpath>
    </and>
  </conditions>
  <actions>
    <do-set-local-variable name="dbTime">
```

```

        <arg-object>
            <token-xpath
expression="Timestamp:valueOf (@jdbc:database-local-time)"/>
            </arg-object>
        </do-set-local-variable>
        <do-set-local-variable name="eventTime">
            <arg-object>
                <token-xpath
expression="Timestamp:valueOf ($commence)"/>
            </arg-object>
        </do-set-local-variable>
    </actions>
</rule>
<rule>
    <description>Is commencement date after database time?</
description>
    <conditions>
        <and>
            <if-xpath op="true">.</if-xpath>
        </and>
    </conditions>
    <actions>
        <do-set-local-variable name="after">
            <arg-string>
                <token-xpath
expression="TimestampUtil:after ($eventTime, $dbTime)"/>
            </arg-string>
        </do-set-local-variable>
    </actions>
</rule>
<rule>
<description>Retry if future event.</description>
    <conditions>
        <and>
            <if-local-variable name="after" op="equal">true</if-
local-variable>
        </and>
    </conditions>
    <actions>
        <do-status level="retry">
            <arg-string>
                <token-text xml:space="preserve">Future event
detected.</token-text>
            </arg-string>
        </do-status>
    </actions>
</rule>
</policy>

```

Setting Up an OCI Client on Linux

K

- ♦ [Section K.1, “Downloading the Instant Client,” on page 251](#)
- ♦ [Section K.2, “Setting Up the OCI Client,” on page 251](#)
- ♦ [Section K.3, “Configuring the OCI Driver,” on page 252](#)

K.1 Downloading the Instant Client

- 1 Download the Oracle Instant Client (`instantclient-basic-linux32-10.2.0.1-20050713.zip`).

The file is available from [Instant Client Downloads \(http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html\)](http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html).

- 2 Download the Oracle SQL*Plus binary (`instantclient-sqlplus-linux32-10.2.0.1-20050713.zip`).

The file is available from [Instant Client Downloads \(http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html\)](http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/linuxsoft.html).

K.2 Setting Up the OCI Client

Set up the Oracle Instant Client on the machine where the JDBC driver is running (not on the machine where Oracle is running).

- 1 Log into Linux as `root`, and create the following structure:

```
/oracle
/oracle/client
/oracle/client/bin
/oracle/client/lib
/oracle/client/network/admin
```

- 2 Unzip all files from `instantclient-basic-linux32-10.2.0.1-20050713.zip` to `/oracle/client/lib`.

- 3 Unzip all files from `instantclient-sqlplus-linux32-10.2.0.1-20050713.zip` to `/oracle/client/bin`.

- 4 Copy `libsqlplus.so` from `/oracle/client/bin` to `/oracle/client/lib`.

- 5 Copy `libsqlplusic.so` from `/oracle/client/bin` to `/oracle/client/lib`.

- 6 Using `chmod`, ensure that the file `sqlplus` in `/oracle/client/bin` is executable.

- 7 Copy a valid `tnsnames.ora` into `/oracle/client/network/admin`.

If you don't have a `tnsnames.ora` file, use the Oracle configuration tool to create one.

Make sure that the `tnsnames.ora` filename is in lowercase.

- 8 Modify the `profile.local` file by adding the following lines:

```
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/oracle/client/lib
export TNS_ADMIN=/oracle/client/network/admin
export PATH=$PATH:/oracle/client/lib
```

The `profile.local` file is in the `/etc` folder. If the file doesn't exist, create one. The file can consist of only the three export lines.

The `profile.local` file extends the `LD_LIBRARY_PATH`, sets `TNS_ADMIN`, and extends the `PATH`. This file is read when the server boots.

- 9 Ensure that the exports in the `profile.local` file are always valid.
- 10 Copy the `classes12.jar` and `ojdbc14.jar` to the IDM classes directory.
These `.jar` files are supplied with the Instant Client.
The IDM classes directory is the directory where your driver is located.
- 11 Start SQL*Plus with the following example command (assuming that the directory is `/oracle/client/bin`):

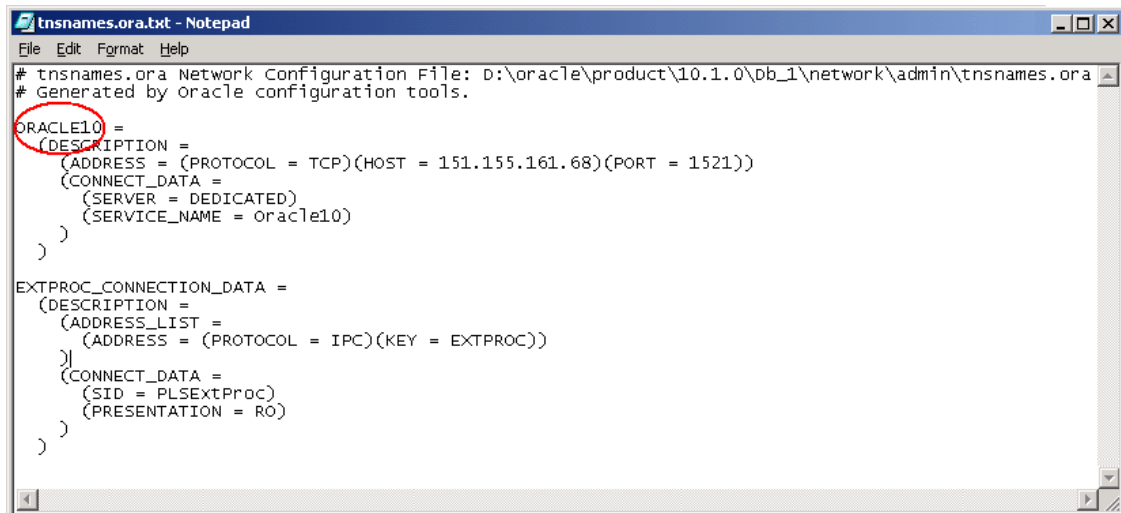
```
./sqlplus username/password@sid
```

K.3 Configuring the OCI Driver

To configure the driver, customize the driver's URL syntax. See [Table 15-18 on page 199](#).

An example URL syntax is `jdbc:oracle:oci8:@ORACLE10`. In this example, `ORACLE10` is the connection string in the `tnsnames.ora` file.

Figure K-1 Example `tnsnames.ora` File



```
tnsnames.ora.txt - Notepad
File Edit Format Help
# tnsnames.ora Network Configuration File: D:\oracle\product\10.1.0\db_1\network\admin\tnsnames.ora
# Generated by Oracle configuration tools.

ORACLE10 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = 151.155.161.68)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = Oracle10)
    )
  )

EXTPROC_CONNECTION_DATA =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC))
    )
    (CONNECT_DATA =
      (SID = PLSEXTProc)
      (PRESENTATION = RO)
    )
  )
```

Sybase Chain Modes and the Identity Manager Driver for JDBC



Sybase can execute stored procedures in two distinct modes: chained and unchained. Depending upon the configuration of the Identity Manager Driver for JDBC and stored procedures in a database, various problems can arise. This section can help you understand and resolve those problems.

- ♦ [Section L.1, “Error Codes,” on page 253](#)
- ♦ [Section L.2, “Procedures and Modes,” on page 254](#)

L.1 Error Codes

- ♦ [“Error 226: SET CHAINED command not allowed within multi-statement transaction” on page 253](#)
- ♦ [“Error 7112: Stored procedure 'x' may be run only in chained transaction mode” on page 253](#)
- ♦ [“Error 7113: Stored procedure 'x' may be run only in unchained transaction mode” on page 254](#)

Error 226: SET CHAINED command not allowed within multi-statement transaction

Effect: Throws the exception of `com.sybase.jdbc2.jdbc.SybSQLException` with error code 226 and an SQL state of `ZZZZZ`.

Cause: This exception is usually caused by a defect in older versions of `jConnect*`.

Solution: Download and upgrade to the latest version. Downloads are available at the [jConnect for JDBC Web page \(http://www.sybase.com/products/informationmanagement/softwaredeveloperkit/jconnect\)](http://www.sybase.com/products/informationmanagement/softwaredeveloperkit/jconnect).

Error 7112: Stored procedure 'x' may be run only in chained transaction mode

Effect: Throws the exception of `com.sybase.jdbc2.jdbc.SybSQLException` with error code 7112 and an SQL state of `ZZZZZ`.

Cause: The stored procedure was created in chained mode, or later altered to run in chained mode, but the driver is currently running in unchained mode. The probable cause is that the [Use Manual Transactions?](#) parameter is set to `False`. Another possibility is that the transaction type has been overridden to `auto` in a policy.

Solution: Do one of the following:

- ♦ Use stored procedure `sp_procxmode` to change the stored procedure's mode to `unchained` or `anymode` (preferred).
- ♦ Change the driver's [Use Manual Transactions?](#) parameter to `True`, or change the policy transaction type to `manual`.

Error 7113: Stored procedure 'x' may be run only in unchained transaction mode

Effect: Throws the exception `com.sybase.jdbc2.jdbc.SybSQLException` with error code 7713 and an SQL state of `ZZZZZ`.

Cause: The stored procedure was created in unchained mode, or later altered to run in unchained mode, but the driver is currently running in chained mode. The probable cause is that the **Use Manual Transactions?** parameter is set to *True*. Another possibility is that the transaction type has been overridden to *manual* in policy.

Solution: Do one of the following:

- ◆ Use stored procedure `sp_proxmode` to change the stored procedure's mode to *chained* or *anymode* (preferred).
- ◆ Change the driver's **“Use Manual Transactions?”** on page 72 parameter to *False*, or change the policy transaction type to *auto*.

NOTE: If you set `use-manual-transactions` to *False*, all transactions will consist of a maximum of one statement.

L.2 Procedures and Modes

- ◆ [Section L.2.1, “Using Stored Procedure `sp_proxmode`,” on page 254](#)
- ◆ [Section L.2.2, “Chained and Unchained Modes,” on page 254](#)
- ◆ [Section L.2.3, “Managing Transactions in a Policy,” on page 255](#)
- ◆ [Section L.2.4, “Useful Links,” on page 255](#)

L.2.1 Using Stored Procedure `sp_proxmode`

The preferred way to avoid errors 7112 and 7113 is to alter all stored procedures invoked directly or indirectly by the driver (via triggers, for example) to run in both chained and unchained mode. To alter a procedure, invoke the `sp_proxmode` procedure with two arguments:

- ◆ The procedure name
- ◆ The mode

The following example illustrates how to invoke the `sp_proxmode` procedure from the `isql` command line:

```
client:sp_proxmode my_procedure, anymode
go
```

Of course, not all customers are willing to alter stored procedure modes. Altering a procedure's mode might alter its runtime behavior, which could alter the behavior of other applications that invoke the procedure.

L.2.2 Chained and Unchained Modes

Unchained mode is Sybase's native way of executing SQL. A second mode, chained mode, was later added to make the database compatible with SQL standards.

Table L-1 Modes and Compatibility

Mode	Compatibility
Chained	SQL-compatible mode
Unchained	Sybase native mode

Sybase provides a third-party JDBC driver called jConnect. The default mode of jConnect is unchained. Whenever the method `Connection.setAutoCommit(boolean autoCommit):void` is invoked, jConnect switches modes. See [java.sql Interface Connection \(http://java.sun.com/j2se/1.4.2/docs/api/java/sql/Connection.html\)](http://java.sun.com/j2se/1.4.2/docs/api/java/sql/Connection.html).

Table L-2 Methods and Switches

Method	Effect
<code>Connection.setAutoCommit(true)</code>	Switches to unchained mode
<code>Connection.setAutoCommit(false)</code>	Switches to chained mode

If the **Use Manual Transactions?** parameter is set to *False*, the driver invokes `Connection.setAutoCommit(true)`. That is, the driver enters unchained mode. This is the normal processing mode for SELECT statements and SQL embedded in a policy where the transaction type is set to `auto`. See [Section 13.5, “Manual vs. Automatic Transactions,” on page 159](#). When the driver is in this state, any chained stored procedures invoked directly or indirectly by the driver yield the 7112 error.

If the **Use Manual Transactions?** parameter is set to *True*, the driver invokes `Connection.setAutoCommit(false)`. That is, the driver enters chained mode. This is the normal processing mode for all statements except SELECT statements and SQL embedded in a policy where the transaction type is set to `manual`. See [Manual vs. Automatic Transactions \(http://www.novell.com/documentation/idmdrivers/index.html?page=/documentation/idmdrivers/jdbc/data/af899ky.html#af8bdjt\)](http://www.novell.com/documentation/idmdrivers/index.html?page=/documentation/idmdrivers/jdbc/data/af899ky.html#af8bdjt). When the driver is in this state, any unchained stored procedures invoked directly or indirectly by the driver yield the 7113 error.

L.2.3 Managing Transactions in a Policy

For information on managing transactions in a policy, see [Manual vs. Automatic Transactions \(http://www.novell.com/documentation/idmdrivers/index.html?page=/documentation/idmdrivers/jdbc/data/af899ky.html#af8bdjt\)](http://www.novell.com/documentation/idmdrivers/index.html?page=/documentation/idmdrivers/jdbc/data/af899ky.html#af8bdjt).

L.2.4 Useful Links

- ◆ [Transaction modes and stored procedures \(http://manuals.sybase.com/onlinebooks/group-as/asg1250e/sqlug/@Generic__BookTextView/55096;hf=0;pt=55096#X\)](http://manuals.sybase.com/onlinebooks/group-as/asg1250e/sqlug/@Generic__BookTextView/55096;hf=0;pt=55096#X) in the *Transact-SQL User's Guide*
- ◆ [Selecting the transaction mode and isolation level \(http://manuals.sybase.com/onlinebooks/group-as/asg1250e/sqlug/@Generic__BookTextView/53713;pt=53001\)](http://manuals.sybase.com/onlinebooks/group-as/asg1250e/sqlug/@Generic__BookTextView/53713;pt=53001) in the *Transact-SQL User's Guide*

The DirXML Command Line Utility

M

The DirXML[®] Command Line utility allows you to use a command line interface to manage the driver. You can create scripts that have the commands to manage the driver.

The utility and scripts are installed on all platforms during the Identity Manager installation. The utility is installed to the following locations:

- ♦ Windows: \Novell\Nds\dxcmd.bat
- ♦ NetWare[®]: sys:\system\dxcmd.ncf
- ♦ UNIX: /usr/bin/dxcmd

Either of the following methods enable you to use the DirXML Command Line utility:

- ♦ [Section M.1, “Interactive Mode,” on page 257](#)
- ♦ [Section M.2, “Command Line Mode,” on page 266](#)

M.1 Interactive Mode

The interactive mode provides a text interface to control and use the DirXML Command Line utility.

- 1 At the console, enter dxcmd.
- 2 Enter the name of a user with sufficient rights to the Identity Manager objects, such as admin.novell.
- 3 Enter the user’s password.

```
DirXML commands
1: Start driver
2: Stop driver
3: Driver operations...
4: Driver set operations...
5: Log events operations...
6: Get DirXML version
7: Job operations...
99: Quit
Enter choice:
```

- 4 Enter the number of the command that you want to perform.
[Table M-1 on page 258](#) contains the list of options and what functionality is available.
- 5 To quit the utility, enter 99.

NOTE: If you are running eDirectory[™] 8.8 on UNIX or Linux*, you must specify the -host and -port parameters. For example, dxcmd -host 10.0.0.1 -port 524. If the parameters are not specified, a jclient error occurs.

```
novell.jclient.JCException: connect (to address) 111 UNKNOWN ERROR
```

By default, eDirectory 8.8 is not listening to localhost. The DirXML Command Line utility needs to resolve the server IP address or hostname and the port to be able to authenticate.

Table M-1 *Interactive Mode Options*

Option	Description
1: <i>Start Driver</i>	Starts the driver. If more than one driver exists, each driver is listed with a number. Enter the number of the driver to start the driver.
2: <i>Stop Driver</i>	Stops the driver. If more than one driver exists, each driver is listed with a number. Enter the number of the driver to stop the driver.
3: <i>Driver operations</i>	Lists the operations available for the driver. If more than one driver exists, each driver is listed with a number. Enter the number of the driver to see the operations available. See Table M-2 on page 259 for a list of operations.
4: <i>Driver set operations</i>	Lists the operations available for the driver set. <ul style="list-style-type: none">◆ 1: Associate driver set with server◆ 2: Disassociate driver set from server◆ 99: Exit
5: <i>Log events operations</i>	Lists the operations available for logging events through Novell® Audit. See Table M-5 on page 263 for a description of these options.
6: <i>Get DirXML version</i>	Lists the installed version of Identity Manager.
7: <i>Job operations</i>	Manages jobs created for Identity Manager.
99: <i>Quit</i>	Exits the DirXML Command Line utility

Figure M-1 *Driver Options*

```
Select a driver operation for:
Active Directory.Driver Set.Novell.IDMDESIGNTREE.

1: Start driver
2: Stop driver
3: Get driver state
4: Get driver start option
5: Set driver start option
6: Resync driver
7: Migrate from application into DirXML
8: Submit XDS command document to driver
9: Submit XDS event document to driver
10: Queue event for driver
11: Check object password
12: Initialize new driver object
13: Passwords operations
14: Cache operations
99: Exit

Enter choice:
```

Table M-2 *Driver Options*

Options	Description
1: <i>Start driver</i>	Starts the driver.
2: <i>Stop driver</i>	Stops the driver.
3: <i>Get driver state</i>	Lists the state of the driver. <ul style="list-style-type: none">◆ 0 - Driver is stopped◆ 1 - Driver is starting◆ 2 - Driver is running◆ 3 - Driver is stopping
4: <i>Get driver start option</i>	Lists the current driver start option. <ul style="list-style-type: none">◆ 1 - Disabled◆ 2 - Manual◆ 3 - Auto
5: <i>Set driver start option</i>	Changes the start option of the driver. <ul style="list-style-type: none">◆ 1 - Disabled◆ 2 - Manual◆ 3 - Auto◆ 99 - Exit
6: <i>Resync driver</i>	<p>Forces a resynchronization of the driver. It prompts for a time delay: <i>Do you want to specify a minimum time for resync? (yes/no)</i>.</p> <p>If you enter Yes, specify the date and time you want the resynchronization to occur: <i>Enter a date/time (format 9/27/05 3:27 PM)</i>.</p> <p>If you enter No, the resynchronization occurs immediately.</p>
7: <i>Migrate from application into DirXML</i>	<p>Processes an XML document that contains a query command: <i>Enter filename of XDS query document:</i></p> <p>Create the XML document that contains a query command by using the Novell nds.dtd (http://developer.novell.com/ndk/doc/dirxml/dirxmlbk/ref/ndsstdtd/query.html).</p> <p>Examples:</p> <p>NetWare: <code>sys:\files\query.xml</code></p> <p>Windows: <code>c:\files\query.xml</code></p> <p>Linux: <code>/files/query.xml</code></p>

Options	Description
8: <i>Submit XDS command document to driver</i>	<p>Processes an XDS command document:</p> <p><i>Enter filename of XDS command document:</i></p> <p>Examples:</p> <p>NetWare: <code>sys:\files\user.xml</code></p> <p>Windows: <code>c:\files\user.xml</code></p> <p>Linux: <code>/files/user.xml</code></p> <p><i>Enter name of file for response:</i></p> <p>Examples:</p> <p>NetWare: <code>sys:\files\user.log</code></p> <p>Windows: <code>c:\files\user.log</code></p> <p>Linux: <code>/files/user.log</code></p>
9: <i>Submit XDS event document to driver</i>	<p>Processes an XDS event document:</p> <p><i>Enter filename of XDS event document:</i></p> <p>Examples:</p> <p>NetWare: <code>sys:\files\add.xml</code></p> <p>Windows: <code>c:\files\add.xml</code></p> <p>Linux: <code>/files/add.xml</code></p>
10: <i>Queue event for driver</i>	<p>Adds an event to the driver queue</p> <p><i>Enter filename of XDS event document:</i></p> <p>Examples:</p> <p>NetWare: <code>sys:\files\add.xml</code></p> <p>Windows: <code>c:\files\add.xml</code></p> <p>Linux: <code>/files/add.xml</code></p>
11: <i>Check object password</i>	<p>Validates that an object's password in the connected system is associated with a driver. It matches the object's eDirectory password (Distribution Password, used with Universal Password).</p> <p><i>Enter user name:</i></p>
12: <i>Initialize new driver object</i>	<p>Performs an internal initialization of data on a new Driver object. This is only for testing purposes.</p>
13: <i>Password operations</i>	<p>Nine Password options are available. See Table M-3 on page 261 for a description of these options.</p>
14: <i>Cache operations</i>	<p>Five Cache operations exist. See Table M-4 on page 262 for a descriptions of these options.</p>

Options	Description
99: <i>Exit</i>	Exits the driver options.

Figure M-2 Password Operations

```

Select a password operation

1: Set shim password
2: Clear shim password
3: Set Remote Loader password
4: Clear Remote Loader password
5: Set named password
6: Clear named password(s)
7: List named passwords
8: Get passwords state
99: Exit

Enter choice:

```

Table M-3 Password Operations

Operation	Description
1: <i>Set shim password</i>	Sets the application password. This is the password of the user account you are using to authenticate into the connected system with.
2: <i>Clear shim password</i>	Clears the application password.
3: <i>Set Remote Loader password</i>	The Remote Loader password is used to control access to the Remote Loader instance. Enter the Remote Loader password, then confirm the password by typing it again.
4: <i>Clear Remote Loader password</i>	Clears the Remote Loader password so no Remote Loader password is set on the Driver object.
5: <i>Set named password</i>	Allows you to store a password or other pieces of security information on the driver. See Section 8.7, "Storing Driver Passwords Securely with Named Passwords," on page 110 for more information. Lists four prompts: <ul style="list-style-type: none"> ◆ <i>Enter password name:</i> ◆ <i>Enter password description:</i> ◆ <i>Enter password:</i> ◆ <i>Confirm password:</i>

Operation	Description
6: <i>Clear named passwords</i>	<p>Clears a specified named password or all named passwords that are stored on the Driver object: <i>Do you want to clear all named passwords? (yes/no)</i>.</p> <p>If you enter Yes, all Named Passwords are cleared. If you enter No, you are prompted to specify the password name that you want to clear.</p>
7: <i>List named passwords</i>	<p>Lists all named passwords that are stored on the Driver object. It lists the password name and the password description.</p>
8: <i>Get password state</i>	<p>Lists if a password is set for:</p> <ul style="list-style-type: none"> ◆ Driver Object password ◆ Application password ◆ Remote loader password <p>The dxcmnd utility enables you to set the Application password and the Remote Loader password. You cannot set the Driver Object password with this utility. It displays whether the password has been set.</p>
99: <i>Exit</i>	<p>Exits the current menu and takes you back to the Driver options.</p>

Figure M-3 *Cache Operations*

```

Enter choice: 14

Select a cache operation

1: Get driver cache limit
2: Set driver cache limit
3: View cached transactions
4: Delete cached transactions
99: Exit

Enter choice:

```

Table M-4 *Cache Operations*

Operation	Description
1: <i>Get driver cache limit</i>	<p>Displays the current cache limit that is set for the driver.</p>
2: <i>Set driver cache limit</i>	<p>Sets the driver cache limit in kilobytes. A value of 0 is unlimited.</p>

Operation	Description
3: <i>View cached transactions</i>	<p>A text file is created with the events that are stored in cache. You can select the number of transactions to view.</p> <ul style="list-style-type: none"> ◆ <i>Enter option token (default=0):</i> ◆ <i>Enter maximum transactions records to return (default=1):</i> ◆ <i>Enter name of file for response:</i>
4: <i>Delete cached transactions</i>	<p>Deletes the transactions stored in cache.</p> <ul style="list-style-type: none"> ◆ <i>Enter position token (default=0):</i> ◆ <i>Enter event-id value of first transaction record to delete (optional):</i> ◆ <i>Enter number of transaction records to delete (default=1):</i>
99: <i>Exit</i>	Exits the current menu and takes you back to the Driver options.

Figure M-4 Log Event Operations

```

Select a log events operation

1: Set driver set log events
2: Reset driver set log events
3: Set driver log events
4: Reset driver log events
99: Exit

Enter choice:

```

Table M-5 Log Events Operations

Operation	Description
1: <i>Set driver set log events</i>	<p>Allows you to log driver set events through Novell Audit. You can select 49 items to log. See Table M-6 on page 264 for a list of these options.</p> <p>Type the number of the item you want to log. After the items are selected, enter 99 to accept the selections.</p>
2: <i>Reset driver set log events</i>	Resets all log event options.
3: <i>Set driver log events</i>	<p>Allows you to log driver events through Novell Audit. You can select 49 items to log. See Table M-6 on page 264 for a list of these options.</p> <p>Type the number of the item you want to log. After the items are selected, enter 99 to accept the selections.</p>

Operation	Description
4: <i>Reset driver log events</i>	Resets all of the log event options.
99: <i>Exit</i>	Exits the log events operations menu.

Table M-6 *Driver Set and Driver Log Events*

Options

- 1: Status success
 - 2: Status retry
 - 3: Status warning
 - 4: Status error
 - 5: Status fatal
 - 6: Status other
 - 7: Query elements
 - 8: Add elements
 - 9: Remove elements
 - 10: Modify elements
 - 11: Rename elements
 - 12: Move elements
 - 13: Add-association elements
 - 14: Remove-association elements
 - 15: Query-schema elements
 - 16: Check-password elements
 - 17: Check-object-password elements
 - 18: Modify-password elements
 - 19: Sync elements
 - 20: Pre-transformed XDS document from shim
 - 21: Post input transformation XDS document
 - 22: Post output transformation XDS document
 - 23: Post event transformation XDS document
 - 24: Post placement transformation XDS document
 - 25: Post create transformation XDS document
 - 26: Post mapping transformation <inbound> XDS document
 - 27: Post mapping transformation <outbound> XDS document
-

Options

- 28: Post matching transformation XDS document
 - 29: Post command transformation XDS document
 - 30: Post-filtered XDS document <Publisher>
 - 31: User agent XDS command document
 - 32: Driver resync request
 - 33: Driver migrate from application
 - 34: Driver start
 - 35: Driver stop
 - 36: Password sync
 - 37: Password request
 - 38: Engine error
 - 39: Engine warning
 - 40: Add attribute
 - 41: Clear attribute
 - 42: Add value
 - 43: Remove value
 - 44: Merge entire
 - 45: Get named password
 - 46: Reset Attributes
 - 47: Add Value - Add Entry
 - 48: Set SSO Credential
 - 49: Clear SSO Credential
 - 50: Set SSO Passphrase
 - 51: User defined IDs
 - 99: Accept checked items
-

Table M-7 Job Scheduler Operations

Options	Description
1: <i>Get available job definitions</i>	Allows you to select an existing job. <i>Enter the job number:</i> <i>Do you want to filter the job definitions by containment? Enter Yes or No</i> <i>Enter name of the file for response:</i> Examples: NetWare: <code>sys:\files\user.log</code> Windows: <code>c:\files\user.log</code> Linux: <code>/files/user.log</code>
2: <i>Operations on specific job object</i>	Allows you to perform operations for a specific job.

M.2 Command Line Mode

The command line mode allows you to use script or batch files. [Table M-8 on page 266](#) lists the different options that are available.

To use the command line options, decide which items you want to use and string them together.

Example: `dxcmd -user admin.headquarters -host 10.0.0.1 -password n0vell -start test.driverset.headquarters`

This example command starts the driver.

Table M-8 Command Line Options

Option	Description
Configuration	
<code>-user <user name></code>	Specify the name of a user with administrative rights to the drivers you want to test.
<code>-host <name or IP address></code>	Specify the IP address of the server where the driver is installed.
<code>-password <user password></code>	Specify the password of the user specified above.
<code>-port <port number></code>	Specify a port number, if the default port is not used.
<code>-q <quiet mode></code>	Displays very little information when a command is executed.
<code>-v <verbose mode></code>	Displays detailed information when a command is executed.

Option	Description
-s <stdout>	Writes the results of the <code>dxcmd</code> command to <code>stdout</code> .
-? <show this message>	Displays the help menu.
-help <show this message>	Displays the help menu.
Actions	
-start <driver dn>	Starts the driver.
-stop <driver dn>	Stops the driver.
-getstate <driver dn>	Shows the state of the driver as running or stopped.
-getstartoption <driver dn>	Shows the startup option of the driver.
-setstartoption <driver dn> <disabled manual auto> <resync noresync>	Sets how the driver starts if the server is rebooted. Sets whether the objects are to be resynchronized when the driver restarts.
-getcachelimit <driver dn>	Lists the cache limit set for the driver.
-setcachelimit <driver dn> <0 or positive integer>	Sets the cache limit for the driver.
-migrateapp <driver dn> <filename>	Processes an XML document that contains a query command. Create the XML document that contains a query command by using the Novell <code>nds.dtd</code> (http://www.novell.com/documentation/idm35/index.html?page=/documentation/idm35/policy_dtd/data/dtdndsoverview.html#dtdndsoverview).
-setshimpassword <driver dn> <password>	Sets the application password. This is the password of the user account you are using to authenticate into the connected system with.
-clearshimpassword <driver dn> <password>	Clears the application password.
-setremoteloaderpassword <driver dn> <password>	Sets the Remote Loader password. The Remote Loader password is used to control access to the Remote Loader instance.
<clearremoteloaderpassword <driver dn>	Clears the Remote Loader password.

Option	Description
-sendcommand <driver dn> <input filename> <output filename>	<p>Processes an XDS command document.</p> <p>Specify the XDS command document as the input file.</p> <p>Examples:</p> <p>NetWare: sys:\files\user.xml</p> <p>Windows: c:\files\user.xml</p> <p>Linux: /files/user.log</p> <p>Specify the output filename to see the results.</p> <p>Examples:</p> <p>NetWare: sys:\files\user.log</p> <p>Windows: c:\files\user.log</p> <p>Linux: /files/user.log</p>
-sendevent <driver dn> <input filename>	<p>Submits a document to the driver's Subscriber channel, bypassing the driver cache. The document is processed ahead of anything that might be in the cache at the time of the submission. It also means that the submission fails if the driver is not running.</p>
-queueevent <driver dn> <input filename>	<p>Submits a document to the driver's Subscriber channel by queuing the document in the driver cache. The document is processed after anything that might be in the cache at the time of the submission. The submission won't fail if the driver isn't running.</p>
-setlogevents <dn> <integer ...>	<p>Sets Novell Audit log events on the driver. The integer is the option of the item to log. See Table M-6 on page 264 for the list of the integers to enter.</p>
-clearlogevents <dn>	<p>Clears all Novell Audit log events that are set on the driver.</p>
-setdriverset <driver set dn>	<p>Associates a driver set with the server.</p>
-cleardriverset	<p>Clears the driver set association from the server.</p>
-getversion	<p>Shows the version of Identity Manager that is installed.</p>
-initdriver object <dn>	<p>Performs an internal initialization of data on a new Driver object. This is only for testing purposes.</p>
-setnamedpassword <driver dn> <name> <password> [description]	<p>Sets named passwords on the driver object. You specify the name, the password, and the description of the named password.</p>
-clearnamedpassword <driver dn> <name>	<p>Clears a specified named password.</p>
-startjob <job dn>	<p>Starts the specified job.</p>

Option	Description
-abortjob <job dn>	Aborts the specified job.
-getjobrunningstate <job dn>	Returns the specified job's running state.
-getjobenabledstate <job dn>	Returns the specified job's enabled state.
-getjobnextruntime <job dn>	Returns the specified job's next run time.
-updatejob <job dn>	Updates the specified job.
-clearallnamedpasswords <driver dn>	Clears all named passwords set on a specific driver.

If a command is executed successfully, it returns a zero. If the command returns anything other than zero, it is an error. For example, 0 means success, and -641 means invalid operation. -641 is an eDirectory error code. [Table M-9 on page 269](#) contains other values for specific command line options.

Table M-9 *Command Line Option Values*


Command Line Option	Values
-getstate	0- stopped 1- starting 2- running 3- shutting down 11- get schema Anything else that is returned is an error.
-getstartoption	0- disabled 1- manual 2- auto Anything else that is returned is an error.
-getcachelimit	0- unlimited Anything else that is returned is an error.
-getjobrunningstate	0- stopped 1- running Anything else that is returned is an error.
-getjobenabledstate	0- disabled 1- enabled 2- configuration error Anything else that is returned is an error.

Command Line Option	Values
-getjobnextruntime	Returns the next scheduled time for the job in eDirectory time format (number of seconds since 00:00:00 Jan 1, 1970 UTC).

Properties of the JDBC Driver


N

This section is a reference for all fields on the driver's property pages as displayed in iManager and Designer. Sometimes fields display differently in iManager than in Designer.

The information is presented from the viewpoint of iManager. If a field is different in Designer for Identity Manager, it is marked with a Designer  icon.

The following figure illustrates property pages in iManager:

Figure N-1 Properties Pages in iManager



- ◆ [Section N.1, “Driver Configuration,” on page 271](#)
- ◆ [Section N.2, “Global Configuration Values,” on page 276](#)
- ◆ [Section N.3, “Named Passwords,” on page 278](#)
- ◆ [Section N.4, “Engine Control Values,” on page 278](#)
- ◆ [Section N.5, “Log Level,” on page 280](#)
- ◆ [Section N.6, “Driver Image,” on page 281](#)
- ◆ [Section N.7, “Security Equals,” on page 281](#)
- ◆ [Section N.8, “Filter,” on page 282](#)
- ◆ [Section N.9, “Edit Filter XML,” on page 282](#)
- ◆ [Section N.10, “Misc,” on page 283](#)
- ◆ [Section N.11, “Excluded Users,” on page 283](#)
- ◆ [Section N.12, “Driver Manifest,” on page 284](#)
- ◆ [Section N.13, “Driver Cache Inspector,” on page 284](#)
- ◆ [Section N.14, “Driver Inspector,” on page 285](#)
- ◆ [Section N.15, “Server Variables,” on page 285](#)

N.1 Driver Configuration

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Configuration*.

In Designer:

- 1 Open a project in the Modeler, then right-click the driver line.
- 2 Click *Properties > Driver Configuration*.
 - ◆ [Section N.1.1, “Driver Module,” on page 272](#)
 - ◆ [Section N.1.2, “Driver Object Password,” on page 273](#)
 - ◆ [Section N.1.3, “Authentication,” on page 273](#)
 - ◆ [Section N.1.4, “Startup Option,” on page 274](#)
 - ◆ [Section N.1.5, “Driver Parameters,” on page 275](#)

N.1.1 Driver Module



The driver module changes the driver from running locally to running remotely or the reverse.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Configuration > Driver Module*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Driver Configuration*.
- 3 Select the *Driver Module* tab.

Option	Description
<i>Java</i>	Used to specify the name of the Java class that is instantiated for the shim component of the driver. This class can be located in the <code>classes</code> directory as a class file, or in the <code>lib</code> directory as a <code>.jar</code> file. If this option is selected, the driver is running locally.
<i>Native</i>	Used to specify the name of the <code>.dll</code> file that is instantiated for the application shim component of the driver. If this option is selected, the driver is running locally.
<i>Connect to Remote Loader</i>	Used when the driver is connecting remotely to the connected system.
 <i>Remote Loader Client Configuration for Documentation</i>	 Includes information on the Remote Loader client configuration when Designer generates documentation on the driver.

N.1.2 Driver Object Password

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Configuration > Driver Object Password > Set Password*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then click *Properties > Driver Configuration*.
- 3 Click *Driver Module > Connect to Remote Loader > Set Password*.

Option	Description
<i>Driver Object Password</i>	Use this option to set a password for the driver object. If you are using the Remote Loader, you must enter a password on this page. Otherwise, the remote driver does not run. The Remote Loader uses this password to authenticate itself to the remote driver shim.

N.1.3 Authentication


The authentication section stores the information required to authenticate to the connected system.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Configuration > Authentication*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Driver Configuration*.
- 3 Click *Authentication*.

Option	Description
Authentication information for server	Displays or specifies the IP address or server name that the driver is associated with
<i>Authentication DN</i>	Specifies the DN of the account that the driver will use for authentication.
or	
 <i>Authentication ID</i>	Example: Administrator

Option	Description
<p><i>Authentication Context</i></p> <p>or</p> <p> <i>Connection Information</i></p>	Specify the IP address or name of the server the application shim should communicate with.
<p><i>Remote Loader Connection Parameters</i></p> <p>or</p> <p> <i>Host name</i></p> <p> <i>Port</i></p> <p> <i>KMO</i></p> <p> <i>Other parameters</i></p>	<p>Used only if the driver is connecting to the application through the Remote Loader. The parameter to enter is</p> <pre>hostname=xxx.xxx.xxx.xxx port=xxxx</pre> <p><i>kmo=certificatename</i>, when the host name is the IP address of the application server running the Remote Loader server and the port is the port the Remote Loader is listening on. The default port for the Remote Loader is 8090.</p> <p>The <i>kmo</i> entry is optional. It is used only when an SSL connection exists between the Remote Loader and the Metadirectory engine.</p> <p>Example: <code>hostname=10.0.0.1 port=8090</code> <code>kmo=IDMCertificate</code></p>
<p><i>Driver Cache Limit (kilobytes)</i></p> <p>or</p> <p> <i>Cache limit (KB)</i></p>	<p>Specify the maximum event cache file size (in KB). If it is set to zero, the file size is unlimited.</p> <p> Click <i>Unlimited</i> to set the file size to unlimited in Designer.</p>
<p><i>Application Password</i></p> <p>or</p> <p> <i>Set Password</i></p>	Specify the password for the user object listed in the <i>Authentication ID</i> field.
<p><i>Remote Loader Password</i></p> <p>or</p> <p> <i>Set Password</i></p>	Used only if the driver is connecting to the application through the Remote Loader. The password is used to control access to the Remote Loader instance. It must be the same password specified during the configuration of the Remote Loader on the connected system.

N.1.4 Startup Option


The Startup Option allows you to set the driver state when the Identity Manager server is started.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Configuration > Startup Option*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Driver Configuration*.
- 3 Click *Startup Option*.

Option	Description
<i>Auto start</i>	The driver starts every time the Identity Manager server is started.
<i>Manual</i>	The driver does not start when the Identity Manager server is started. The driver must be started through Designer or iManager.
<i>Disabled</i>	The driver has a cache file that stores all of the events. When the driver is set to <i>Disabled</i> , this file is deleted and no new events are stored in the file until the driver state is changed to <i>Manual</i> or <i>Auto Start</i> .
 <i>Do not automatically synchronize the driver</i>	This option applies only if the driver is deployed and was previously disabled. If this is not selected, the driver re-synchronizes the next time it is started.

N.1.5 Driver Parameters

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Configuration > Driver Parameters*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Driver Configuration*.
- 3 Click *Driver Parameters*.

Parameter	Description
Driver parameters for server	Species of displays the name or IP address of the server that the driver is associated with.
Edit XML	Opens an editor so that you can edit the configuration file.
Driver Options	
<i>Third-party JDBC driver class name</i>	Verify that the version of the <code>jar</code> or <code>zip</code> file that contains this class is compatible with the target database.
<i>Show database scoping parameters</i>	Controls how much of the database is visible to the JDBC driver.
<i>Table/view names</i>	The names of the tables or views to synchronize. For details, click the information icon.
<i>Time syntax</i>	Returns time-related data types as integers, canonical strings, or Java strings. For details, click the information icon.
<i>State directory</i>	The path of the folder where you want the JDBC driver to store state information. Changing the value while using triggerless publication can force all objects to be resynchronized.
<i>Show connectivity parameters?</i>	Hides or displays connectivity-related parameters.

Parameter	Description
<i>Show compatibility parameters?</i>	Hides or displays parameters that can enhance compatibility with third-party JDBC databases.
Subscriber Options	
<i>Disable subscriber?</i>	Specifies whether the driver ignores events that flow from Identity Manager to the database.
<i>Show primary key parameters?</i>	Hides or displays parameters that control how and when primary key values are generated.
<i>Disable statement-level locking?</i>	Specifies whether database resources are locked on a channel.
<i>Check update counts?</i>	Specifies whether the Subscriber channel checks to see if rows were updated after insert, update, or delete statements were executed against a table.
Publisher Options	
<i>Disable publisher?</i>	Specifies whether the Publisher channel ignores events that flow from the database to the Identity Vault.
<i>Publication mode</i>	Specifies whether the driver uses triggered or triggerless publication.
<i>Event log table name</i>	Names the table where publication events are stored.
<i>Delete processed rows?</i>	Specifies whether processed rows are deleted from the event log table. Performance decreases when processed rows remain in the table.
<i>Optimize updates?</i>	Specifies whether the Publisher channel optimizes update events before sending them to the engine. To optimize all update events, select Yes.
<i>Allow loopback?</i>	Specifies whether events that the Subscriber channel initiates loopback on the Publisher channel.
<i>Disable statement-level locking?</i>	Specifies whether database resources are locked on this channel.
<i>Enable future event processing?</i>	Specifies whether events in the event log table are processed according to their effective date or order of insertion, or whether the database local time is published with each event.
<i>Show polling-related parameters?</i>	Specifies whether to show parameters that control polling behavior.
Heartbeat interval (in minutes)	Specifies the number of minutes that the Publisher channel waits between events before sending a heartbeat document.

N.2 Global Configuration Values

Global configuration values (GCVs) enable you to specify settings for the Identity Manager features such as password synchronization and driver heartbeat, as well as settings that are specific to the function of an individual driver configuration. Some GCVs are provided with the drivers, but you can also add your own.

IMPORTANT: Password synchronization settings are GCVs, but it's best to edit them in the graphical interface provided on the Server Variables page for the driver, instead of the GCV page. The Server Variables page that shows Password Synchronization settings is accessible as a tab as with other driver parameters, or by clicking *Password Management > Password Synchronization*, searching for the driver, and clicking the driver name. The page contains online help for each Password Synchronization setting.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Global Config Values*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Global Configuration Values*.

Table N-1 *Global Configuration Values > Password Configuration*

Option	Description
<i>Application accepts passwords from Identity Manager</i>	If <i>True</i> , allows passwords to flow from the Identity Manager data store to the connected system.
<i>Identity Manager accepts passwords from application</i>	If <i>True</i> , allows passwords to flow from the connected system to Identity Manager.
<i>Publish passwords to NDS password</i>	Use the password from the connected system to set the non-reversible NDS [®] password in eDirectory.
<i>Publish passwords to Distribution Password</i>	Use the password from the connected system to set the NMAST [™] Distribution Password used for Identity Manager password synchronization.
<i>Require password policy validation before publishing passwords</i>	If <i>True</i> , applies NMAST password policies during publish password operations. The password is not written to the data store if it does not comply.
<i>Reset user's external system password to the Identity Manager password on failure</i>	If <i>True</i> , on a publish Distribution Password failure, attempt to reset the password in the connected system by using the Distribution Password from the Identity Manager data store.
<i>Notify the user of password synchronization failure via e-mail</i>	If <i>True</i> , notify the user by e-mail of any password synchronization failures.
<i>Connected System or Driver Name</i>	The name of the connected system, application, or Identity Manager driver. The e-mail notification templates use this value.

N.3 Named Passwords

Identity Manager enables you to store multiple passwords securely for a particular driver. This functionality is referred to as Named Passwords. Each different password is accessed by a key, or name.

You can also use the Named Passwords feature to store other pieces of information securely, such as a user name. To configured Named Passwords, see [Section 8.7, “Storing Driver Passwords Securely with Named Passwords,”](#) on page 110.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Named Passwords*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Named Passwords*.

N.4 Engine Control Values

The engine control values are a means through which certain default behaviors of the Metadirectory engine can be changed. The values can only be accessed if a server is associated with the Driver Set object.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Engine Control Values*.

In Designer:

- 1 In the Modeler, right-click a driver line.
- 2 Select *Properties > Engine Control Values*.
- 3 Click the tooltip icon to the right of the *Engine Controls for Server* field. If a server is associated with the Identity Vault, the Engine Control Values display in the large pane.

Table N-2 *Engine Control Values*

Option	Description
<i>Subscriber channel retry interval in seconds</i>	The Subscriber channel retry interval controls how frequently the Metadirectory engine retries the processing of a cached transaction after the application shim's Subscriber object returns a retry status.

Option	Description
<i>Qualified form for DN-syntax attribute values</i>	The qualified specification for DN-syntax attribute values controls whether values for DN-syntax attribute values are presented in unqualified slash form or qualified slash form. A True setting means the values are presented in qualified form.
<i>Qualified form from rename events</i>	The qualified form for rename events controls whether the new-name portion of rename events coming from the Identity Vault are presented to the Subscriber channel with type qualifiers. For example, CN=. A True setting means the names are presented in qualified form.
<i>Maximum eDirectory replication wait time in seconds</i>	This setting controls the maximum time that the Metadirectory engine waits for a particular change to replicate between the local replica and a remote replica. This only affects operations where the Metadirectory engine is required to contact a remote eDirectory server in the same tree to perform an operation and might need to wait until some change has replicated to or from the remote server before the operation can be completed (for example, object moves when the Identity Manager server does not hold the master replica of the moved object; file system rights operations for Users created from a template.)
<i>Use non-compliant backwards-compatible mode for XSLT</i>	<p>This control sets the XSLT processor used by the Metadirectory engine to a backwards-compatible mode. The backward-compatible mode causes the XSLT processor to use one or more behaviors that are not XPath 1.0 and XSLT 1.0 standards-compliant. This is done in the interest of backward compatibility with existing DirXML[®] style sheets that depend on the non-standard behaviors.</p> <p>For example, the behavior of the XPath “!=” operator when one operand is a node-set and the other operand is other than a node-set is incorrect in DirXML releases up to and including Identity Manager 2.0. This behavior has been corrected; however, the corrected behavior is disabled by default through this control in favor of backward compatibility with existing DirXML style sheets.</p>
<i>Maximum application objects to migrate at once</i>	<p>This control is used to limit the number of application objects that the Metadirectory engine requests from an application during a single query that is performed as part of a Migrate Objects from Application operation.</p> <p>If java.lang.OutOfMemoryError errors are encountered during a Migrate from Application operation, this number should be set lower than the default. The default is 50.</p> <hr/> <p>NOTE: This control does not limit the number of application objects that can be migrated; it merely limits the batch size.</p> <hr/>
<i>Set creatorsName on objects created in Identity Vault</i>	<p>This control is used by the Identity Manager engine to determine if the creatorsName attribute should be set to the DN of this driver on all objects created in the Identity Vault by this driver.</p> <p>Setting the creatorsName attribute allows for easily identifying objects created by this driver, but also carries a performance penalty. If not set, the creatorsName attribute defaults to the DN of the NCP[™] Server object that is hosting the driver.</p>

Option	Description
<i>Write pending associations</i>	<p>This control determines whether the Identity Manager engine writes a pending association on an object during Subscriber channel processing.</p> <p>Writing a pending association confers little or no benefit but does incur a performance penalty. Nevertheless, the option exists to turn it on for backward compatibility.</p>
<i>Use password event values</i>	<p>This control determines the source of the value reported for the nspmDistributionPassword attribute for Subscriber channel Add and Modify events.</p> <p>Setting the control to False means that the current value of the nspmDistributionPassword is obtained and reported as the value of the attribute event. This means that only the current password value is available. This is the default behavior.</p> <p>Setting the control to True means that the value recorded with the eDirectory event is decrypted and is reported as the value of the attribute event. This means that both the old password value (if it exists) and the replacement password value at the time of the event are available. This is useful for synchronizing passwords to certain applications that require the old password to enable setting a new password.</p>
<i>Enable password synchronization status reporting</i>	<p>This control determines whether the Identity Manager engine reports the status of Subscriber channel password change events.</p> <p>Reporting the status of Subscriber channel password change events allows applications such as the Identity Manager User Application to monitor the synchronization progress of a password change that should be synchronized to the connected application.</p>

N.5 Log Level

Each driver set and each driver has a log level field where you can define the level of errors that should be tracked. The level you indicate here determines which messages are available to the logs. By default, the log level is set to track error messages. (This also includes fatal messages.) To track additional message types, change the log level.


Novell® recommends that you use Novell Audit instead of setting the log levels. See the *Identity Manager 3.5.1 Logging and Reporting* guide.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Log Level*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Driver Log Level*.

Option	Description
<i>Use log settings from the DriverSet</i>	If this is selected, the driver logs events as the options are set on the Driver Set object.
<i>Log errors</i>	Logs just errors
<i>Log errors and warnings</i>	Logs errors and warnings
<i>Log specific events</i>	Logs the events that are selected. Click the  icon to see a list of the events.
<i>Only update the last log time</i>	Updates the last log time.
<i>Logging off</i>	Turns logging off for the driver.
<i>Turn off logging to DriverSet, Subscriber and Publisher logs</i>	If selected, turns all logging off for this driver on the Driver Set object, Subscriber channel, and the Publisher channel.
<i>Maximum number of entries in the log (50-500)</i>	Number of entries in the log. The default value is 50.

N.6 Driver Image

Allows you to change the image associated with the driver. You can browse and select a different image from the default image.

The image associated with a driver is used by the Identity Manager Overview plug-in when showing the graphical representation of your Identity Manager configuration. Although storing an image is optional, it makes the overview display more intuitive.

NOTE: The driver image is maintained when a driver configuration is exported.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Image*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > iManager Icon*.

N.7 Security Equals

Use the Security page to view or change the list of objects that the driver is explicitly security equivalent to. This object effectively has all rights of the listed objects.

If you add or delete an object in the list, the system automatically adds or deletes this object in that object's "Security Equal to Me" property. You don't need to add the [Public] trustee or the parent

containers of this object to the list, because this object is already implicitly security equivalent to them.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Security Equals*.

Designer does not list the users the driver is security equals to.

N.8 Filter

Launches the Filter editor. You can edit the Filter from this tab.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Filter*.

In Designer:

- 1 In an open project, click the *Outline* tab (Outline view).
- 2 Select the driver you want to manage the filter for, then click the plus sign to the left.
- 3 Double-click the *Filter* icon to launch the Filter Editor.

N.9 Edit Filter XML

Allows you to edit the filter directly in XML instead of using the Filter Editor.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Filter*.

In Designer:

- 1 In an open project, click the *Outline* tab (Outline view).
- 2 Select the driver you want to manage the filter for, then click the plus sign to the left.
- 3 Double-click the *Filter* icon to launch the Filter Editor, then click *XML Source* at the bottom of the Filter Editor.

N.10 Misc


Allows you to add a trace level to your driver. With the trace level set, DSTrace displays the Identity Manager events as the Metadirectory engine processes the events. The trace level affects only the driver it is set for. Use the trace level for troubleshooting issues with the driver when the driver is deployed. DSTrace displays the output of the specified trace level.

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Misc*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Trace*.

Option	Description
<i>Trace level</i>	Increases the amount of information displayed in DSTrace. Trace level 1 shows errors, but not the cause of the errors. If you want to see password synchronization information, set the trace level to 5.
<i>Trace file</i>	When a value is set in this field, all Java information for the driver is written to the file. The value for this field is the path for that file. As long as the file is specified, Java information is written to this file. If you do not need to debug Java, leave this field blank.
<i>Trace file size limit</i>	Allows you to set a limit for the Java trace file. If you set the file size to Unlimited, the file grows in size until there is no disk space left.
<i>Trace name</i>	Driver trace messages are prepended with the value entered in this field.
 <i>Use setting from Driver Set</i>	This option is only available in Designer. It allows the driver to use the same setting that is set on the Driver Set object.

N.11 Excluded Users

Use this page to create a list of users or resources that are not replicated to the application. Novell recommends that you add all objects that represent an administrative role to this list (for example, the Admin object).

In iManager:

- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Excluded Users*.

Designer does not list the excluded users.

N.12 Driver Manifest

The driver manifest is like a resumé for the driver. It states what the driver supports, and includes a few configuration settings. The driver manifest is created by default when the Driver object is imported. A network administrator usually does not need to edit the driver manifest.

In iManager:

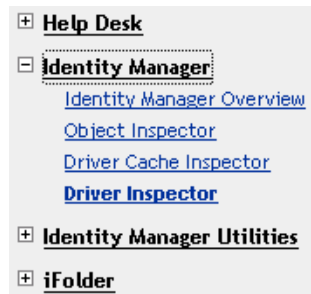
- 1 Click *Identity Manager > Identity Manager Overview*, then click *Search* to search for the driver set that is associated with the driver.
- 2 Browse to the driver, then click the upper right corner of the driver icon.
- 3 Click *Edit Properties > Driver Manifest*.

In Designer:

- 1 Open a project in the Modeler.
- 2 Right-click the driver line, then select *Properties > Driver Manifest*.

N.13 Driver Cache Inspector

Figure N-2 *The Link to the Driver Cache Inspector*



The Driver Cache Inspector page uses a table format to display information about the cache file that stores events while the driver is stopped.

- ♦ **Driver:** A link to run the *Driver Overview* on the driver that is associated with this cache file.
- ♦ **Driver Set:** A link to run the *Driver Set Overview* on the driver set that holds the driver.
- ♦ **Driver's cache on:** Lists the server object that contains this instance of the cache file.
- ♦ **Start/Stop Driver icons:** Displays the current state of the driver and allows you to start or stop the driver.
- ♦ **Edit icon:** Enables you to edit the properties of the currently selected Server object.
- ♦ **Delete:** Deletes the selected items from the cache file.
- ♦ **Refresh:** Enables you to re-read the cache file and refresh the displayed information.
- ♦ **Show:** Limits the number of items to be displayed. The options are:
 - ♦ 25 per page
 - ♦ 50 per page
 - ♦ 100 per page
 - ♦ Other: Enables you to specify a desired number.

- ◆ **Actions:** Enables you to perform actions on the entries in the cache file. Click *Actions* to expand the menu, which includes:
 - ◆ **Expand All:** Expands all of the entries displayed in the cache file.
 - ◆ **Collapse All:** Collapses all of the entries displayed in the cache file.
 - ◆ **Go To:** Enables you to access a specified entry in the cache file. Specify the entry number, then click *OK*.
 - ◆ **Cache Summary:** Summarizes all events stored in the cache file.

N.14 Driver Inspector

The Driver Inspector page displays information about objects associated with the driver. Designer does not have this option.

- ◆ **Driver:** A link to run the *Driver Overview* on the driver that is being inspected.
- ◆ **Driver Set:** A link to run the *Driver Set Overview* of the driver set that holds the driver.
- ◆ **Delete:** Deletes the associations of the selected objects.
- ◆ **Refresh:** Enables you to re-read all of the objects associated with the driver and refresh the displayed information.
- ◆ **Actions:** Enables you to perform actions on the objects associated with the driver. Click *Actions* to expand the menu, which includes:
 - ◆ **Show All Associations:** Displays all objects associated with the driver.
 - ◆ **Filter for Disabled Associations:** Displays all the driver's associated objects that have a Disabled state.
 - ◆ **Filter for Manual Associations:** Displays all the driver's associated objects that have a Manual state.
 - ◆ **Filter for Migrate Associations:** Displays all the driver's associated objects that have a Migrate state.
 - ◆ **Filter for Pending Associations:** Displays all the driver's associated objects that have a Pending state.
 - ◆ **Filter for Processed Associations:** Displays all the driver's associated objects that have a Processed state.
 - ◆ **Filter for Undefined Associations:** Displays all the driver's associated objects that have an Undefined state.
 - ◆ **Association Summary:** Displays the state of all objects associated with the driver.
- ◆ **Object DN:** Displays the DN of the associated objects.
- ◆ **State:** Displays the association state of the object.
- ◆ **Object ID:** Displays the value of the association.

N.15 Server Variables

This page lets you enable and disable Password Synchronization and the associated options for the selected driver.

When setting up Password Synchronization, consider both the settings on this page for an individual driver and the Universal Password Configuration options in your password policies.

This page lets you control which password Identity Manager updates directly, either the Universal Password for an Identity Vault, or the Distribution Password used for password synchronization by Identity Manager.

However, Novell Modular Authentication Service (NMAS) controls whether the various passwords inside the Identity Vault are synchronized with each other. Password Policies are enforced by NMAS, and they include settings for synchronizing Universal Password, NDS Password, Distribution Password, and Simple Password.

To change these settings in iManager:

- 1 In iManager, select *Passwords > Password Policies*.
- 2 Select a password policy, then click *Edit*.
- 3 Select *Universal Password*.

This option is available from a drop-down list or a tab, depending on your version of iManager and your browser.

- 4 Select *Configuration Options*, make changes, then click *OK*.


NOTE: Enabling or disabling options on this page corresponds to values of True or False for certain global configuration values (GCVs) used for password synchronization in the driver parameters. Novell recommends that you edit them here in the graphical interface, instead of on the GCVs page. This interface helps ensure that you don't set conflicting values for the password synchronization GCVs.

Option	Description
<i>Identity Manager accepts password (Publisher Channel)</i>	<p>If this option is enabled, Identity Manager allows passwords to flow from the connected system driver into the Identity Vault data store.</p> <p>Disabling this option means that no <code><password></code> elements are allowed to flow to Identity Manager. They are stripped out of the XML by a password synchronization policy on the Publisher channel.</p> <p>If this option is enabled, and the option below it for Distribution Password is disabled, a <code><password></code> value coming from the connected system is written directly to the Universal Password in the Identity Vault if it is enabled for the user. If the user's password policy does not enable Universal Password, the password is written to the NDS Password.</p>

Option	Description
<i>Use Distribution Password for password synchronization</i>	<p>To use this setting, you must have a version of eDirectory that supports Universal Password, regardless of whether you have enabled Universal Password in your password policies.</p> <p>If this option is enabled, a password value coming from the connected system is written to the Distribution Password. The Distribution Password is reversible, which means that it can be retrieved from the Identity Vault data store for password synchronization. It is used by Identity Manager for bidirectional password synchronization with connected systems. For Identity Manager to distribute passwords to connected systems, this option must be enabled.</p> <p>NMAS and Password policies control whether the Distribution Password is synchronized with other passwords in the Identity Vault. By default, the Distribution Password is the same as the Universal Password in the Identity Vault.</p> <p>If the password in the Identity Vault is to be independent of Password Synchronization, so that Identity Manager is a conduit only for synchronizing passwords among connected systems, change this default setting. In the Universal Password Configuration Options in a Password policy, disable <i>Synchronize Universal Password with Distribution Password</i>. This use of Identity Manager Password Synchronization is also referred to as “tunneling.”</p>
<i>Accept password only if it complies with user's Password Policy</i>	<p>To use this setting, users must have a Password policy assigned that has Universal Password enabled, and Advanced Password Rules enabled and configured.</p> <p>If this option is chosen, Identity Manager does not write a password from this connected system to the Distribution Password in the Identity Manager data store or publish it to connected systems unless the password complies with the user's Password policy.</p> <p>By using the notification option that is also on this page, you can inform users when a password is not set because it is not compliant.</p>

Option	Description
<i>If password does not comply, ignore Password Policy on the connected system by resetting user's password to the Distribution Password</i>	<p>This option lets you enforce Password policies on the connected system by replacing a password that does not comply. If you select this option, and a user's password on the connected system does not comply with the user's Password policy, Identity Manager resets the password on the connected system by using the Distribution Password from the Identity Vault data store.</p> <p>Keep in mind that if you do not select this option, user passwords can become out-of-sync on connected systems.</p> <p>By using the notification option that is also on this page, you can inform users when a password is not set or reset. Notification is especially helpful for this option. If the user changes to a password that is allowed by the connected system but rejected by Identity Manager because of the Password policy, the user won't know that the password has been reset until the user receives a notification or tries to log in to the connected system with the old password.</p> <hr/> <p>NOTE: Consider the connected system's password policies when deciding whether to use this option. Some connected systems might not allow the reset because they don't allow you to repeat passwords.</p>
<i>Always accept password; ignore Password Policies</i>	<p>If you select this option, Identity Manager does not enforce the user's Password policy for this connected system. Identity Manager writes the password from this connected system to the Distribution Password in the Identity Vault data store, and distributes it to other connected systems, even if the password does not comply with the user's Password policy.</p>
<i>Application accepts passwords (Subscriber Channel)</i>	<p>If you select this option, the driver sends passwords from the Identity Vault data store to this connected system. This also means that if a user changes the password on a different connected system that is publishing passwords to the Distribution Password in the Identity Vault data store, the password is changed on this connected system.</p> <p>By default, the Distribution Password is the same as the Universal Password in the Identity Vault, so changes to the Universal Password made in the Identity Vault are also sent to the connected system.</p> <p>If you want the password in the Identity Vault to be independent of Password Synchronization, so that Identity Manager is a conduit only for synchronizing passwords among connected systems, you can change this default setting. In the Universal Password Configuration Options in a password policy, disable <i>Synchronize Universal Password with Distribution Password</i>. This use of Password Synchronization is also referred to as "tunneling."</p>
<i>Notify the user of password synchronization failure via email</i>	<p>If you select this option, e-mail is sent to the user if a password is not synchronized, set, or reset. The e-mail that is sent to the user is based on an e-mail template. This template is provided by the Password Synchronization application. However, for the template to work, you must customize it and specify an e-mail server to send the notification messages.</p> <hr/> <p>NOTE: To set up e-mail notification, select <i>Passwords > Edit Email Templates</i>.</p>

Documentation Updates



The documentation was updated on the following dates:

- ◆ [Section O.1, “October 10, 2007,” on page 289](#)

O.1 October 10, 2007

Updates were made to the following sections. The changes are explained below.

O.1.1 Upgrading the JDBC Driver

Location	Change
Section 4.2, “Upgrading after Identity Manager Is Installed,” on page 46	Change the driver name from Delimited Text to JDBC.
