
UnboundID® Synchronization Server

3.0.3 Administration Guide



Version **3.0.3**
March 3, 2011

Copyright

Copyright © 2011 UnboundID Corporation
All rights reserved

This document constitutes an unpublished, copyrighted work and contains valuable trade secrets and other confidential information belonging to UnboundID Corporation. None of the foregoing material may be copied, duplicated or disclosed to third parties without the express written permission of UnboundID Corporation.

This distribution may include materials developed by third parties. Third-party URLs are also referenced in this document. UnboundID is not responsible for the availability of third-party web sites mentioned in this document. UnboundID does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. UnboundID will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

"UnboundID" is a registered trademark of UnboundID Corporation. UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group. All other registered and unregistered trademarks in this document are the sole property of their respective owners.

The contents of this publication are presented for information purposes only and is provided "as is". While every effort has been made to ensure the accuracy of the contents, the contents are not to be construed as warranties or guarantees, expressed or implied, regarding the products or services described herein or their use or applicability. We reserve the right to modify or improve the design or specifications of such products at any time without notice.

UnboundID Corp
13809 Research Blvd
Suite 500
Austin, TX 78750
www.UnboundID.com

Contents

Chapter 1 Introduction	1
Overview of the Synchronization Server	1
The Synchronization Problem	2
The UnboundID Advantage	3
Common Synchronization Use Cases	3
Use Case: Synchronization during Directory Server Migrations	4
Use Case: Advanced Replication	4
Use Case: Synchronizing with Active Directory	8
Use Case: Realistic Test Environments	9
Use Case: Synchronization with Databases	9
Use Case: Synchronization through Proxy Servers	9
Synchronization Server: How It Works	9
Point-to-Point Bidirectional Synchronization	9
Architecture	10
Modes of Operation	12
Data Transformations	13
Configuration Model	13
Sync Control Flow Scenarios	17
A Synchronization Example	19
Available Tools Summary	19
Summary	20
Chapter 2 Installing the Synchronization Server	21
Before You Begin	21
Supported Operating Platforms	22
Software Requirements: Java 6	22
Unpacking the Installation Packages	23
Installing the UnboundID Synchronization Server	23
Installing the Synchronization Server	23
Installing the Synchronization Server in Non-Interactive Mode	26
Running the Synchronization Server	27
Automatically Starting the Synchronization Server at Boot Time	27
Stopping the Synchronization Server	28
Restarting the Synchronization Server	28
Uninstalling the Synchronization Server	29
Uninstalling the Synchronization Server in Interactive Command-Line Mode	29
Uninstalling the Synchronization Server in Non-Interactive Mode	30
Installing the Sync Management Console	30
Uninstalling the Console	33

Updating the Synchronization Server	33
Reverting an Update	34
Upgrading the Console	35
Installing a Redundant Failover Server	35
Setting the --XX:-UseSSE42Intrinsics Java Property	37

Chapter 3 Configuring the Synchronization Server 39

Pre-Deployment Checklist	39
External Servers	40
Sync Pipes	40
Sync Classes	40
Creating Administrators	41
Configuring the Synchronization Server	42
Configuring the Synchronization Server Using the Command-Line Wizard	42
About the Sync User Account	43
Assumptions	43
Configuring the Synchronization Server Using the Command-Line Wizard	44
Preparing the External Servers	46
Configuring the Sync Pipes and its Sync Classes	46
Review the Configuration and Apply the Changes	47
Configuring the Attribute Map and Mapping	48
Completing the Bidirectional Deployment	49
Configuring the Synchronization Server Using the Management Console	50
Configuring the External Servers	50
Configuring the Sync Pipe	56
Configuring the Sync Class	61
Starting the Global Sync Configuration	67
Using the dsconfig Configuration Tool	68
Using dsconfig in Interactive Command-Line Mode	68
Using dsconfig in Non-Interactive Command-Line Mode	69
Using dsconfig in Batch Mode	69
Configuring the Synchronization Server Using dsconfig	70
Configuring External Servers Using dsconfig Interactive Mode	71
Configuring the Sync Source Using dsconfig Interactive Mode	71
Configuring the Sync Destination Using dsconfig Interactive Mode	72
Configuring a Sync Pipe Using dsconfig Interactive Mode	73
Configuring Sync Class Using dsconfig Interactive Mode	74
Starting the Global Sync Configuring Using dsconfig Interactive Mode	75
Preparing the Synchronization Server for External Server Communication	75
Preparing External Servers: If the Admin Does Not Have Root Access on DSEE External Servers	76
Using Bulk Resync on the Synchronization Server	79
Testing Attribute and DN Maps using Resync	80
Verifying the Synchronization Configuration using Resync	80
Populating an Empty Sync Destination Topology Using Resync	81
Populating an Empty Sync Destination Topology Using Translate-LDIF	82
Setting the Synchronization Rate Using Resync	82
Synchronizing a Specific List of DNs	83
Controlling Real Time Synchronization	84
Starting Real Time Synchronization Globally	85
Pausing Synchronization	86
Setting Startpoints	86
Scheduling a Realtime Sync as a Task	88
Configuring Attribute Maps	88
Configuring an Attribute Map Using dsconfig Interactive	89
Configuring an Attribute Mapping Using dsconfig Interactive	90

Configuring an Attribute Mapping Using dsconfig Non-Interactive	90
Configuring the Directory Server Backend for Synchronizing Deletes	91
Configuring DN Maps	91
Configuring a DN Map Using dsconfig Interactive	92
Configuring a DN Map Using dsconfig Non-Interactive Mode	93
Configuring Fractional Replication	93
Managing Failover Behavior	95
Conditions that Trigger Immediate Failover	96
Configuration Properties that Control Failover Behavior	97
max-operation-attempts	98
response-timeout	98
max-failover-error-code-frequency	99
max-backtrack-replication-latency	99
Changing the max-backtrack-replication-latency Property	100

Chapter 4 Configuring Synchronization with Microsoft Active Directory Systems 101

Before You Begin	101
Configuring Active Directory Synchronization	102
Preparing the External Servers	104
Configuring the Sync Pipes and its Sync Classes	104
Configuring the Password Encryption Component	107
Installing the UnboundID Password Sync Agent	108
Supported Platforms	109
Before You Begin	109
Installing the Password Sync Agent	110
Upgrading the Password Sync Agent	110
Uninstalling the Password Sync Agent	111
Manual Configuration for Advanced Users	111

Chapter 5 Configuring Synchronization with Relational Databases 113

Overview	113
About the DBSync Process	113
About the Server SDK	114
About the DBSync Example	115
Example DS Entries	116
About the Overall DBSync Configuration Process	117
Downloading the Software Packages	117
Creating the JDBC Adapter Script	118
About Groovy	119
Implementing a Scripted JDBC Sync Source	119
Implementing a Scripted JDBC Sync Destination	121
Configuring the Database for Synchronization	122
Pre-Configuration Checklist	124
Configuring the Directory-to-Database Sync Pipe	125
General Tips When Syncing to a Database Destination	125
Step 1. Creating the Directory-to-Database Sync Pipe	126
Step 2. Configuring the Sync Pipe and Sync Classes	129
Step 3. Fine-tuning the Sync Classes	132
Step 4. Configuring the Attribute Mappings	136
Step 5. Run the Resync Tool to Test the Configuration	138
Step 6. Set the Startpoint in the Changelog	138
Step 7. Run the Resync Tool to Populate Data into One Endpoint	139
Step 8. Start the Sync Pipe	139
Step 9. Debugging the Configuration	139

Configuring the Database-to-Directory Sync Pipe	141
General Tips When Syncing from a Database Source	141
Chapter 6 Syncing Through Proxy Servers	145
Features	145
How It Works	146
About the Get Changelog Batch and Get Server ID Controls	147
About the Directory Server and Directory Proxy Server Tokens	148
Changelog Tracking in Entry-Balanced Deployments	149
About the Overall Sync-through-Proxy Configuration Process	150
Configuring Sync Through Proxy	151
Configure the Directory Servers	151
Configuring the Proxy Servers	153
Configuring the Destination Proxy Deployment	155
Configuring the Synchronization Server	157
Set the Proxy-Server and Use-Changelog-Batch-Request Properties	160
Run prepare-external-server on the Backend Set of Directory Servers	161
Testing and Starting the Configuration	161
Chapter 7 Managing Logging and Monitoring	163
Working with Logs	163
Types of Log Publishers	163
Default Synchronization Server Logs	164
Viewing the List of Log Publishers	165
Sync Log Messaging Types	166
Creating New Log Publishers	166
Configuring Log Rotation	168
Configuring Log Retention	168
Working with Administrative Alert Handlers	169
Configuring the JMX Connection Handler and Alert Handler	170
Configuring the SMTP Alert Handler	171
Running the Status Tool	171
Monitoring the Synchronization Server	175
Profiling Server Performance Using the Periodic Stats Logger	176
Configuring a Custom Periodic Stats Logger	179
Chapter 8 Troubleshooting the Synchronization Server	185
About Synchronization Troubleshooting	185
Working with the Troubleshooting Tools	186
Troubleshooting Process Flow	187
Using the Sync Log	187
Troubleshooting Sync Failures	189
Troubleshooting "Entry Already Exists" Failures	189
Troubleshooting "No Match Found" Failures	191
Troubleshooting "Failure at Resource" Failures	192
Working with the Collect Support Data Tool	194
Available Tool Options	196
General Index	199

1 Introduction

Overview

The UnboundID® Synchronization Server is a point-to-point, directory-centric synchronization server between source and destination topologies comprising of UnboundID® Directory Server, UnboundID® Directory Proxy Server (3.x), Alcatel-Lucent® 8661 Directory Server, Alcatel-Lucent® 8661 Directory Proxy Server (3.x), Sun™ Directory Server Enterprise Edition (DSEE 6.x, 7.x), Sun™ Directory Server (5.2 patch 3 or higher), Microsoft® Active Directory®, Oracle® Database (10g, 11g), and Microsoft® SQL Server (2005, 2008) systems. The Synchronization Server has a low cost of ownership with minimal administrative and hardware expenditures to provide a high performance synchronization solution.

This chapter presents a general overview of the Synchronization Server:

- [Overview of the Synchronization Server](#)
- [The Synchronization Problem](#)
- [The UnboundID Advantage](#)
- [Common Synchronization Use Cases](#)
- [Synchronization Server: How It Works](#)
- [Configuration Model](#)
- [Sync Control Flow Scenarios](#)
- [A Synchronization Example](#)
- [Available Tools Summary](#)
- [Summary](#)

Overview of the Synchronization Server

The UnboundID Synchronization Server is an efficient, pure Java-base server that provides high-throughput, low-latency, and bidirectional real-time synchronization between two end-point topologies consisting of directory servers, directory proxy servers, and/or Relational Database Management Systems (RDBMS) systems. Designed to run on inexpensive hardware with little administrative maintenance (i.e., backups are not required), the UnboundID Synchronization Server provides an effective cost-per-performance solution for synchronizing data between LDAP and RDBMS directory topologies.

The Synchronization Server includes the following key features:

- High performance and availability with built-in redundancy to help ensure no downtime.
- Dataless, virtual architecture for a small-memory footprint and easy maintenance.
- Hassle-free setup that allows you to transform and map attribute names, values, and DNs between endpoints. For directory server endpoints, this benefit allows you to make schema and directory information tree changes without the added costs of custom coding and scripting.
- Data flexibility and security, allowing you to replicate data and use advanced replication features in fractional, local data, filtered, or sub-tree replication scenarios.
- Multi-vendor directory server support including the UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun Directory Server Enterprise Edition (DSEE 6.x, 7.x), Sun Directory Server (5.2 patch 3 or higher), and Microsoft Active Directory.
- Relational Database Management System (RDBMS) support including Oracle Database (10g, 11g), and Microsoft SQL Server (2005, 2008) systems.
- Directory Proxy Server support including the UnboundID Directory Proxy Server (3.x) and the Alcatel-Lucent Directory Proxy Server (3.x).

The Synchronization Problem

Synchronization is the process of maintaining data consistency among applications, directories, and data sources in a networked environment. System administrators who use a directory as a writable user repository must ensure that the directory be exposed to all of its applications. However, exposing the user repository becomes problematic when consolidating applications and systems. Often these administrators find that many synchronization solutions lack features, such as, writable partial replicas, or the ability to synchronize with other directory systems and Relational Database Management System (RDBMS) sources.

Most companies employ a meta-directory or a virtual directory synchronization strategy as follows:

I. Meta-Directory (Datafull Approach). The meta-directory solution aggregates all data from the various sources and makes it available for its applications in a centralized directory. The centralized directory then can be updated with those changes pushed back out to the original data sources.

While the meta-directory approach may appear to be an easy-to-manage solution, there are some fundamental flaws that limit its ability for cost-effectiveness and performance:

- **Scalability Limitations.** To maintain a combined view of the data from its sources, the centralized meta-directory is often a bottleneck for any updates that must go through a single directory server instance when synchronizing from one endpoint to another.
- **Functionality Limitations.** The meta-directory solution must often integrate disparate company directories into a single distributed enterprise directory. Integrating data mismatches (e.g., schema variances, privilege differences, etc.) require additional solutions that limit ease-of-use.
- **Administrative and Hardware Cost Limitations.** Because a meta-directory stores a shadow copy of all of the source data that will be synchronized, it requires a large storage and memory footprint. This hardware requirement leads to additional hardware costs and increases the administrative burden of managing backups. The meta-directory solution also has difficulty in providing instantaneous failover between redundant instances.

II. Virtual Directory (Dataless Approach). Virtual Directories provide a consolidated view of the data without actually creating a physical centralized repository for directory information. When an application requests data from the virtual directory, the directory assembles the data and delivers it to the application in real time. However, to achieve synchronization between two backend directory topologies, virtual directories require that all applications update data through the virtual directory exclusively. This scenario prevents client applications from directly modifying the backend directory instances.

The UnboundID Advantage

Synchronization can be a challenging problem when integrating multiple data sources. UnboundID Corporation has a proven track record of successful deployments combined with many years of extensive synchronization experience to solve the problem.

The UnboundID Synchronization Server uses a dataless approach that synchronizes changes directly from the data sources in the background, so that applications can continue to update its data sources directly. The Synchronization Server does not store any data from the endpoints themselves, thereby reducing hardware and administration costs. The server's replication mechanisms also make it easy to fail over from the main synchronization server to its redundant instances.

Common Synchronization Use Cases

The UnboundID Synchronization Server is designed to synchronize data across independent directory server topologies that has its source and destination endpoints using the following vendor directory topologies:

- UnboundID Directory Server topologies
- UnboundID Directory Proxy Server (3.x) topologies
- Alcatel-Lucent 8661 Directory Server
- Alcatel-Lucent 8661 Directory Proxy Server (3.x) topologies
- Sun Directory Server Enterprise Edition (DSEE 6.x, 7.x) topologies
- Sun Directory Server (5.2 patch 3 or higher) topologies
- Microsoft Active Directory topologies
- Relational Database Management Systems (RDBMS) using Oracle (10g, 11g), Microsoft SQL Server (2005, 2008)

The typical deployment scenarios that require synchronization services involve synchronizing data during directory server migrations, replicating with advanced features during normal operations, synchronizing with Active Directory systems, performing real-time testing by obfuscating production data, synchronizing with database systems, and synchronizing through proxy servers.

Use Case: Synchronization during Directory Server Migrations

Directory Server migrations from one system to another can be complicated due to the mismatches in functionality from one system to another (for example, replication limitations, schema mismatches and others). Additionally, if problems arise during a migration, reverting the process can be especially difficult. The UnboundID Synchronization Server solves the migration/reversion problem by allowing you to leave the source deployment untouched, while a separate, synchronized topology of targeted servers is installed and tested. Modifications generated by an application in either topology are immediately synchronized to the other topology and are available to all applications. Once all of the applications have been tested against the new installation, the source directory servers can be phased out.

The general procedure for a migration (for example, from Sun Directory Server 5.x to UnboundID Directory Server) is as follows:

1. Leave the Sun Directory Server 5.x in place.
2. Set up synchronization from the Sun Directory Server 5.x to the UnboundID Directory Server, or vice-versa.
3. Gradually migrate applications and data to the UnboundID Directory Server. You can use the UnboundID Proxy Server in front of the Sun Directory Server 5.x and the UnboundID Directory Server topologies to redirect some client applications to a particular topology.
4. In the event of a rare migration failure, the migration can be reverted back to the Sun Directory Server 5.x if required.

Use Case: Advanced Replication

The UnboundID Synchronization Server provides advanced features that extend the replication capability of its directory servers. The Synchronization Server supports either fractional replication, local data replication, filtered replication, subtree replication, or a combination to replicate parts of a Directory Information Tree (DIT) or subsets of entries. Traditionally, repli-


cation creates exact replicas of servers, including the same DIT structure, entries, and attributes. However, in many cases, replica servers need to store a subset of entries, a subset of attributes, or in some case extra attributes, compared to the full DIT in the primary master servers. Because all servers do not need full copies of the data, the extended replication features of the Synchronization Server improves the overall performance of the directory service and reduces hardware costs.

To provide an example scenario, a large telecommunications company is managing data replication between three divisions: billing, web, and network. The directory server for each division contains the same subscriber information. While the billing division is the authoritative source for the subscriber information, each division has its own unique directory structure. Each division replicates data between its own local servers using replication agreements that accommodate the division's unique DIT and schema. The Synchronization Server can synchronize data across division boundaries using the different replication methods to address the needs of each division.

Fractional Replication

Fractional replication is a form of partial replication that allows a subset of attributes to be replicated to another replica. This feature is often used to reduce replication bandwidth by including only the data that is needed. For example, if a replica only performs user authentications, then replication can be configured to only propagate the `uid` and `userpassword` attributes for a password policy, reducing the database size at the replica and the network traffic needed to keep the server in sync to this server. Furthermore, changes due to password policy attributes, such as, account lockouts, can be replicated back to the main master servers. The UnboundID Synchronization Server supports fractional replication to any type of server.

FIGURE 1-1. Fractional Replication



Returning again to the telecom company example, a subscriber can change their email address in a variety of ways:

- Calling the billing department
- Going to a retail store
- Logging on to the web site
- Using IVR on their telephone


No matter where the email address is changed, it needs to be reflected everywhere. If a subscriber changes his or her email address by calling the billing department, the Synchronization Server uses fractional replication to replicate only the updated email attribute of the subscriber's entry across the servers in other departments.

Local Data Replication

For fractional replication, application-specific repositories generally require directories that contain less data than their primary master servers. For local data replication, replicas tend to have more data than the primary master servers. For example, some applications need to store large amounts of data in a user entry, such as, an XML blob of preference information, a sound file, or an image. Although the data could be required by only one application, the data itself can impact the server performance for all applications. The Synchronization Server can keep this local data isolated to only a few servers dedicated to this application without burdening the master corporate servers.

Returning to the large telecom company example, we can imagine that the web department uses a portal server and web applications that are not used by the other departments. The user preference information stored on the user entries in the web department's directory server is not replicated back to the other departments. Instead, this information is replicated only between the servers in the web department.


FIGURE 1-2. Local Data Replication



Filtered Replication

Filtered replication is a form of partial replication, where a replica contains only selected entries as determined by an LDAP filter. This feature allows directory instances to replicate only specific subtrees of a DIT, determined by base DN's returned using inclusion or exclusion filters. Applications that create an application-specific entry can be restricted to only those directories used by the application and not to every replica in the topology.

FIGURE 1-3. Subtree Replication



For example, the telecom example has an extranet directory server and a corporate directory server. The extranet directory server is used by the sales staff that are on the road to authenticate and use their email over the web. To access their email remotely, an employee must have a `web-access-enabled=true` flag set on their entry. Using filtered replication, the Synchroni-


zation Server can look for entries that have this flag set to true and replicate their changes back to the other corporate directories. If an employee changes their password in the extranet directory, this change will be replicated back to the master corporate directory.

Subtree Replication

In subtree replication, a replica contains only the selected entries as determined by a directory branch. This feature allows directory server instances to replicate only specific subtrees of a DIT, which are determined by inclusion or exclusion filtering on the base DN.

For example, the large telecom company acquires a media company. The telecom company adds a subtree of data in its directory server for the media company, and the media company itself has its own on-premise LDAP directory server. Using a subtree replication protocol, data can be replicated between the media company's directory server and the main telecom directory server.


FIGURE 1-4. Filtered Replication



Advanced Replication Combinations

The UnboundID Synchronization Server can support the various combinations of fractional and filtered replication for those applications that require it. For example, imagine that the telecom company operates in several countries in the European Union. Europe itself restricts the disclosure of anything considered personal data. Rules inside individual countries can further restrict the type of data that can be transferred between country boundaries. While this data must remain in the individual country's directory server, the corporate directory still needs to contain as much information as possible. Using fractional replication, only the parts of the entry that can legally cross country borders would be replicated back to the main corporate directory server.

Further, imagine that the main directory for the telecom company, `dc=corp,dc=com` contains subdirectories for each of the European directories, such as, `ou=France`, `ou=Germany`, and `ou=Italy`. Locally, each country has its own unique directory data stored in its own DIT structure, such as, `dc=corp-fr,dc=com` for France. Using subtree replication in combination with fractional replication, the Synchronization Server can replicate changes between a country's subtree in the master directory and each country's local directory, while adhering to local laws about the transfer of personal information.

FIGURE 1-5. Advanced Replication Combinations

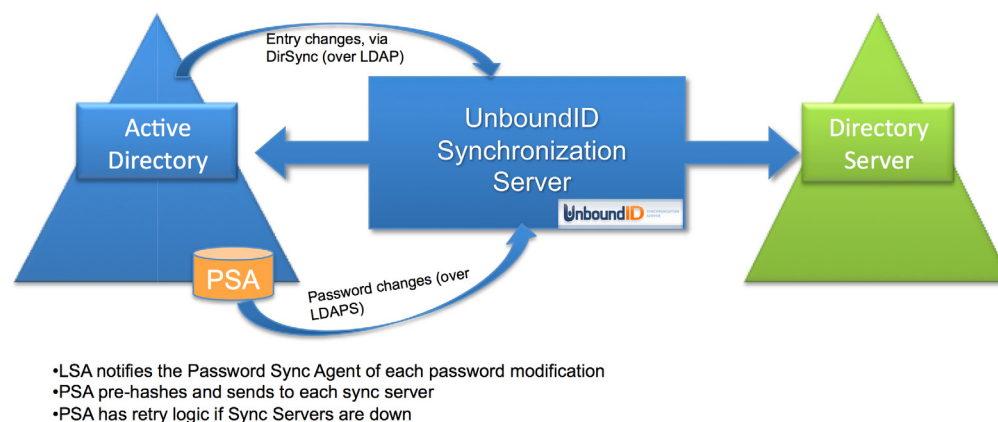
Use Case: Synchronizing with Active Directory

The UnboundID Synchronization Server provides a Microsoft Active Directory mechanism that synchronizes ADD, MODIFY, and DELETE operations for user entries and individual attributes using Active Directory's DirSync control.

If real-time password synchronization is needed, the Synchronization Server also requires that a dedicated component, the UnboundID Password Sync Agent (PSA), be installed on all Active Directory domain controllers. The agent receives password changes from the Local Security Authority (LSA) and immediately hashes them with a secure 160-bit salted secure hash. The agent then sends the hashes to each UnboundID Synchronization Server instance in the topology over a secure LDAPS connection. If the synchronization server instance is down, the agent caches the change and retries synchronization until at least one of the servers has received the updates.

The agent is highly optimized with a small memory footprint. It securely handles sensitive data and uses a small, native DLL on the domain controller, which requires a single restart (Microsoft requirement). Subsequent updates to the DLL do not require a restart.

For more information on Active Directory Sync, see "Configuring Synchronization with Microsoft Active Directory Systems" on page 101.

FIGURE 1-6. Active Directory Sync

Use Case: Realistic Test Environments

In an effort to secure sensitive user data, many companies have their own test environments that use synthetic data that may not be compatible with actual production data. This discrepancy can introduce problems when moving applications from a test environment to a production environment. The UnboundID Synchronization Server is capable of fully synchronizing test or stage servers with production servers while also obfuscating sensitive customer information, such as, social security and phone numbers. The Synchronization Server can sync in real-time or on a nightly basis with little additional performance load on the production servers.

Use Case: Synchronization with Databases

In environments that store data on both directory servers and databases (Oracle 11g, 10g and Microsoft SQL Server 2005, 2008), the UnboundID Synchronization Server can synchronize data on both systems. This solution is more flexible than having to sync through a virtual directory. You can configure the UnboundID Synchronization Server to establish the directory server or the database as the authoritative data source. See “Configuring Synchronization with Relational Databases” on page 113.

Use Case: Synchronization through Proxy Servers

Many data centers use proxy servers in front of a backend set of directory servers in load-balanced and/or entry-balanced deployments. The UnboundID Synchronization Server provides the ability to synchronize data through such proxy deployments. See “Syncing Through Proxy Servers” on page 145.

Synchronization Server: How It Works

The UnboundID Synchronization Server is a lightweight, standalone, 100% Java solution that provides low-latency, highly-available, point-to-point synchronization. The Synchronization Server shares many of the reliable components with the UnboundID Directory Server for easy installation and configuration.

Point-to-Point Bidirectional Synchronization


The UnboundID Synchronization Server performs point-to-point synchronization between a source endpoint and a destination endpoint. An *endpoint* is defined as any Source or Destination topology of directory or database servers comprising of any combination of either UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun Directory Server Enterprise Edition (DSEE 6.x, 7.x), Sun Directory Server (5.2 patch 3 or higher), Microsoft Active Directory, Oracle Database (10g, 11g), and Microsoft SQL Server (2005, 2008) systems.

The Synchronization Server provides that ability to sync data in one directory or bidirectionally between endpoints. For example, in a migration phase from Sun Directory Server to UnboundID Directory Server, you can sync in one direction from the source server to a QA stage server for testing purposes. By definition, the source server is the authoritative endpoint as it generates all of the changes in the system. In a bidirectional synchronization configuration, both sets of endpoints (i.e., the source and the destination) are authoritative for the same set of attributes or for different sets of data. Bidirectional synchronization allows for parallel active installations.

The Synchronization Server also contains no single point of failure, either for detecting changes or for applying changes. The Synchronization Server instances themselves are redundant, so that you can have multiple synchronization server instances running at a time, but only the server with the highest priority is actively syncing changes. Further, the stand-by servers are constantly polling the active server instance to update their persistent state. This state contains the minimum amount of information needed to begin synchronization where the primary server left off, which logically, is the last processed change number for the source server. In the case of a network partition, multiple synchronization servers can synchronize simultaneously without causing problems as they each verify the full entry before making any changes.

Figure 1-7 shows a typical UnboundID Synchronization Server deployment. The Synchronization Server looks for any changes on the main source server on the left and applies those changes to the destination server on the right. In the diagram, the bold lines indicate the primary (active) connections within the synchronization network that show the directional path of the changes. The synchronization servers will communicate with these servers if they are up. The gray lines are possible failover connections if any component is down.

FIGURE 1-7. Synchronization Typology



Architecture

The UnboundID Synchronization Server uses a virtualized, dataless approach and never permanently stores any directory data locally. This approach eliminates the need for backups and large disk requirements. The log files, administrator entries, configuration, sync state informa-

tion are stored as flat files (LDIF format) within the system. No additional database is required.

Change Tracking

To track changes in each endpoint system, the UnboundID Synchronization Server uses the changelog mechanism that is most efficient for each platform.

- For UnboundID Directory Server or Alcatel-Lucent 8661 Directory Server topologies, the UnboundID Synchronization Server uses the instance's LDAP Changelog for modification detection.
- For Sun DSEE or Sun Directory Server topologies, the UnboundID Synchronization Server uses the server's Retro Changelog, which provides a detailed summary of each change applied to the directory.
- For Active Directory, the UnboundID Synchronization Server uses the DirSync control, which polls for object attribute changes.
- For RDBMS systems, the UnboundID Synchronization Server uses a JDBC Adapter Script to interface with a customized RDBMS changelog table. Database triggers on each table record all INSERT, UPDATE, and DELETE operations to the changelog table.

Each directory instance stores a separate entry under `cn=changelog` for every modification made to the directory. The UnboundID Synchronization Server provides full control over the synchronization process by determining which entries are synchronized, how they are correlated to the entries at the destination endpoint, and how they are transformed into the destination schema.

Monitoring and Alert Notifications

The Synchronization Server supports several industry-standard, administrative protocols (seen in Figure 1-8) for monitoring and alert notifications:

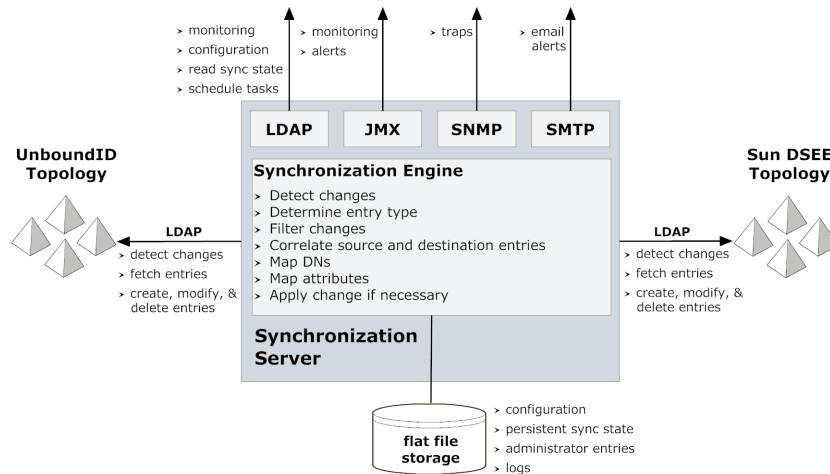
- LDAP. Used for monitoring, configuration, server state, tasks.
- JMX. Used for monitoring and alerts.
- SMTP. Used for email alert notifications.

The UnboundID Synchronization Server provides an administrative alert framework that can be used to notify administrators of any significant warnings, errors, or other noteworthy events that occur in the server. Existing alert notification handlers can notify administrators through log messages, email, or JMX notifications. All administrative alerts are also exposed over LDAP as entries below a base DN "cn=alerts". You can use the persistent search operation to ensure that you are automatically notified over LDAP of any new alerts generated by the server. The administrator can select the admin action for each type of notification based on the severity level or the specific type of alert. For example, it may be desirable to log information about all types of alerts, but only generate email messages. Typical alert events are startup/shutdown, applied configuration changes, or synchronized resources unavailable.

Logging

The UnboundID Synchronization Server provides standard logs (sync, access, error, failed-operations, config-audit.log, debug) to troubleshoot any issues. The server can also be configured for multiple active sync logs. For example, you can log each detected change, each dropped change, each applied change, or each failed change.

FIGURE 1-8. Synchronization Architecture



Modes of Operation

The UnboundID Synchronization Server runs as a standalone Java process with two complementary synchronization modes of operation: real-time or scheduled bulk synchronization mode. The real-time and bulk modes are typically used together when configuring a Synchronization network. Real-time synchronization can continue to synchronize new changes as they are made, while bulk resynchronization operations guarantee that all entries are in-sync at initialization.

Real-time Synchronization Mode

In real-time mode, the UnboundID Synchronization Server polls the source server for changes and synchronizes the destination entries immediately. Once the UnboundID Synchronization Server determines that a detected change should be included in the synchronization, it fetches the full entry from the source. Then, it finds the corresponding entry in the destination endpoint using flexible correlation rules and applies the minimum set of changes to bring the attributes that were modified into sync. The server fetches and compares the full entries to make sure it does not synchronize any stale data from the changelog.

Bulk Synchronization Resync Mode

The UnboundID Synchronization Server's bulk synchronization mode, known as *resynchronization*, streams all entries or those entries matching certain criteria from the source servers and updates the corresponding destination entries. The server does a bulk comparison of the source entries with the destination entries and, depending on the configuration, updates and reports what is out-of-sync.

Resynchronization is principally designed to accomplish any of the following tasks:

- Initially populate a target directory over LDAP.
- Verify that the two endpoints are in-sync.
- Perform scheduled (for example, nightly) synchronization in place of real-time synchronization.
- Recover from a failover by resynchronizing entries that were modified since the last backup was taken.
- Preview what is out-of-sync to validate that the server is behaving as expected. You can use the `--dry-run` option on active systems to validate that sync is operating properly without updating any entries.

Resynchronization also allows for fine control over what can be synchronized. You can control the following items:

- Include or exclude source and destination attributes
- Apply an LDAP filter to only sync entries created since the last time you ran the tool (for example, `createTimeStamp >= 20100311200-0600`).
- Synchronize only creates or only modifications.
- Change the logging verbosity.
- Set a limit on how fast the resynchronization runs (for example, only 2000/second) to limit the impact on endpoint directory servers.

Data Transformations

Data transformations alter the contents of synchronized entries between the source and destination directory server topologies to transparently handle variances in attribute names, attribute values, or DN structures. When the UnboundID Synchronization Server synchronizes entries between a source and a destination server, it can be configured to change the contents of these entries using attribute and DN mappings, so that neither server needs be aware of the transformations.

- **Attribute Mapping.** The Synchronization Server can transparently rename any attribute in the entry to fit the schema definitions from the source endpoint to the destination endpoint. This mapping makes it possible to synchronize information stored in one attribute in one directory server topology to an attribute with a different name in another directory server topology, or to construct an attribute using portions of the source attribute values.
- **DN Mapping.** The Synchronization Server can transparently alter any DNs referenced in the entries. This mapping makes it possible to synchronize data from a topology that uses one DIT structure to a system that uses a different DIT structure.

Configuration Model

The UnboundID Synchronization Server supports a flexible configuration model that is designed for easy installation and maintenance based on a comprehensive set of command-

line tools and a graphical console. The Synchronization Server supports a defined set of configuration parameters that determine how synchronization takes place between directories or databases. The server can be configured remotely or locally with all configuration changes taking effect on-the-fly.

The configuration parameters are listed in Table 1-1 and the configuration model is summarized in Figure 1-9.

TABLE 1-1. UnboundID Synchronization Server Configuration Components

Component	Description
Sync Pipe	Defines a single synchronization path between the source and destination topologies. Every Sync Pipe has one or more Sync Classes that controls how and what is synchronized. Multiple Sync Pipes can run in a single UnboundID Synchronization Server instance.
Sync Source	Defines the directory topology that is the source of the data to be synchronized. When data in the Sync Source changes, it is synchronized to the Sync Destination topology. A Sync Source can reference one or more external servers of the appropriate type (UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun DSEE, Sun Directory Server 5.x, Microsoft Active Directory, Oracle, Microsoft SQL Server).
Sync Destination	Defines the topology of directory servers where changes detected at the Sync Source are applied. A Sync Destination can reference one or more external servers of the appropriate type (UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun DSEE, Sun Directory Server 5.x, Microsoft Active Directory, Oracle, Microsoft SQL Server).
External Server	Defines a single server in a topology of identical, replicated servers to be synchronized. For an LDAP server, you must define the host, port, SSL, bind DN, and bind password. A single external server configuration object can be referenced by multiple Sync Sources and Sync Destinations.
Sync Class	Defines the operation types (e.g., creates, modifies, or deletes) and attributes that are synchronized, how attributes and DNs are mapped, and how source and destination entries are correlated. A source entry is in at most one Sync Class and is determined by a base DN and LDAP filters. A Sync Class can reference zero or more Attribute Maps and DN Maps, respectively. Within a Sync Pipe, a Sync Class is defined for each type of entry that needs to be treated differently. For example, entries that define attribute mappings or entries that should not be synchronized at all. A Sync Pipe must have at least one Sync Class but can refer to multiple Sync Class objects.

TABLE 1-1. UnboundID Synchronization Server Configuration Components

Component	Description
DN Map	<p>Defines mappings when destination DNs differ from source DNs. These mappings allow the use of wild cards for DN transformations. A single wild card ("*") matches a single RDN component and can be used any number of times. The double wild card ("**") matches zero or more RDN components and can be used only once. The wild card values can be used in the <code>to-dn-pattern</code> attribute using <code>{1}</code> and their original index position in the pattern, or <code>{attr}</code> to match an attribute value. For example:</p> <pre>** ,dc=myexample,dc=com->{1},o=example</pre> <p>You can also use regular expressions and attributes from the user entry. For example, the following mapping constructs a value for the <code>uid</code> attribute, which is the RDN, out of the initials (first letter of <code>givenname</code> and <code>sn</code>) and the employee ID (the <code>eid</code> attribute).</p> <pre>uid={givenname:/^(.) (.*)/\$1/s}{sn:/^(.) (.*)/\$1/s}{eid},{2},o=example</pre> <p>The following diagram shows how a nested DIT can be mapped to a flattened structure:</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;"> <p>Nested DIT</p> <pre> graph TD A["dc=example,dc=com"] --> B["st=TX"] B --> C["o=accounts"] C --> D["acctid=geneh"] D --> E["sub=5127516011 acctid: geneh st: TX"] </pre> </div> <div style="flex: 1;"> <p>Flattened DIT</p> <pre> graph TD A["dc=example,dc=com"] --> B["o=accounts"] A --> C["o=subscribers"] B --> D["acctid=geneh"] C --> E["sub=5127516011 acctid: geneh st: TX"] </pre> </div> <div style="flex: 1;"> <p>* matches 1 RDN ** matches 0 or more RDNs {1} substitute first wildcard {2} substitute second wildcard {attr} substitute attr value</p> </div> </div> <div style="margin-top: 10px;"> <p>Subscriber to Flattened Map</p> <pre> from: *,** ,dc=example,dc=com to: {1},o=subscribers,dc=example,dc=com </pre> <p>Subscriber to Nested Map</p> <pre> from: *,** ,dc=example,dc=com to: {1},acctid={acctid},o=accounts,st={st},dc=example,dc=com </pre> <p style="text-align: right;">attributes from subscriber entry</p> </div>
Attribute Map & Attribute Mappings	<p>A Sync Class can reference zero or more DN maps. Multiple Sync Classes share the same DN Map.</p> <p>Defines a mapping when destination attributes differ from source attributes. An Attribute Map is a collection of Attribute Mappings. There are three types of Attribute mappings:</p> <ul style="list-style-type: none"> • Direct mapping. Attributes are directly mapped to another attribute: For example, <code>employeenumber->employeeid</code> • Constructed Mapping. Destination attribute values are derived from source attribute values and static text. For example: <pre>{givenname} . {sn}@example.com->mail</pre> • DN Mapping. Attributes are mapped for attributes that store DNs. You can reference the same DN mappings that map entry DNs. For example, you could have a <code>manager</code> attribute that has the value <code>uid=jdoe,ou=People,dc=example,dc=com</code>. <p>A Sync Class can reference multiple Attribute Maps. Multiple Sync Classes can share the same Attribute Map.</p>

Figure 1-9 shows a summary of the configuration model and how the configuration objects can be referenced.

FIGURE 1-9. Configuration Model Referenced Objects

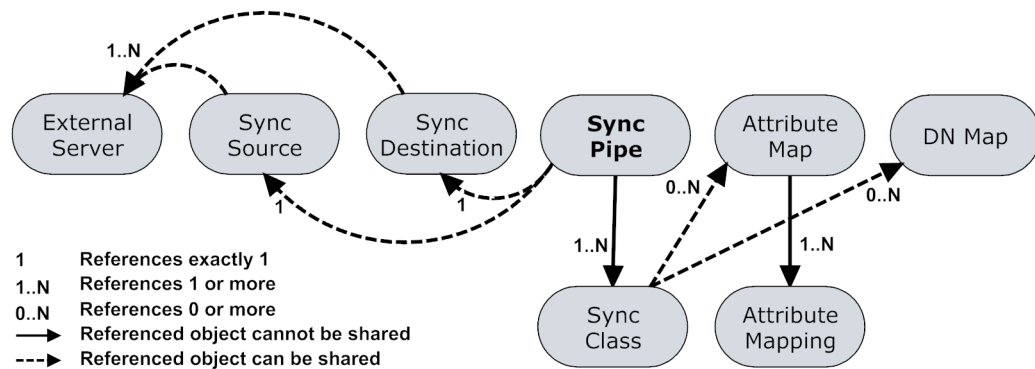
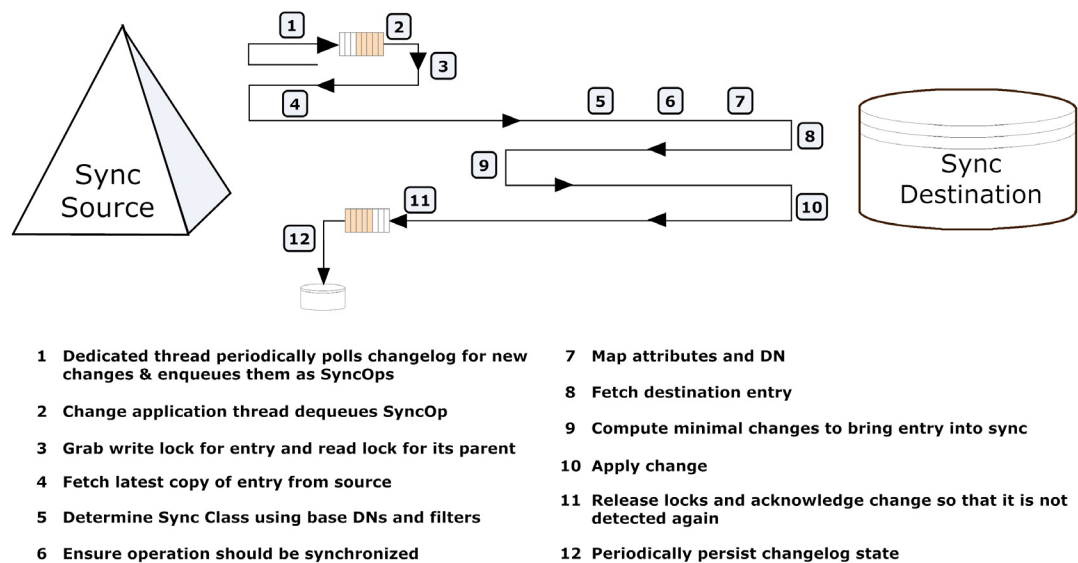


Figure 1-10 shows the typical Synchronization Change Flow that occurs in a process. Although not pictured, the changes are processed in parallel, which increases throughput and offsets network latency.

FIGURE 1-10. Synchronization Change Flow



Sync Control Flow Scenarios

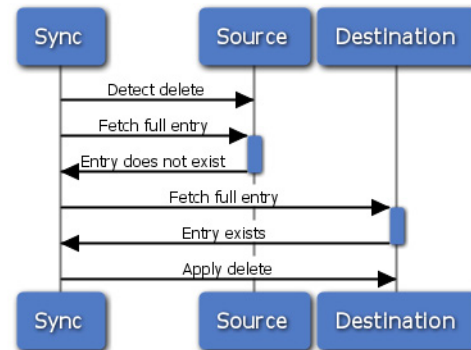
The Synchronization Server processes changes by fetching the most up-to-date, full entries from both sides and then compares them. The processing flow differs depending on the type of Sync Server change (ADD, MODIFY, DELETE, MODDN) that is requested. Table 1-2 shows the control flow diagrams for the sync operations, especially for those cases when a MODIFY or a DELETE operation is dropped. The sync logs record all completed and failed operations.

TABLE 1-2. Sync Server Operational Control Flow


<p>Standard Modify</p> <ol style="list-style-type: none"> 1. Detect change from the changelog table on the source. 2. Fetch the entry or table rows from affected tables on the source. 3. Perform any mappings and compute the equivalent destination entry by constructing an equivalent LDAP entry or equivalent table row. 4. Fetch the entry or table rows from affected tables on the destination. 5. Diff the computed destination entry and actual destination entry. 6. Apply the minimal set of changes to the destination to bring it in sync. 	<pre> sequenceDiagram participant Sync participant Source participant Destination Sync->>Source: Detect change Sync->>Source: Fetch full entry Note over Sync: Apply mappings Sync->>Destination: Fetch full entry Note over Sync: Diff the entries Sync->>Destination: Apply minimal changes </pre>
<p>Standard Add</p> <ol style="list-style-type: none"> 1. Detect change from the changelog table on the source. 2. Fetch the entry or table rows from affected tables on the source. 3. Perform any mappings and compute the equivalent destination entry by constructing an equivalent LDAP entry or equivalent table row. 4. Fetch the entry or table rows from affected tables on the destination. 5. The entry or table row does not exist on the destination. 6. Create the entry or table row. 	<pre> sequenceDiagram participant Sync participant Source participant Destination Sync->>Source: Detect change Sync->>Source: Fetch full entry Note over Sync: Apply mappings Sync->>Destination: Fetch full entry Destination-->>Sync: Entry does not exist Sync->>Destination: Create entry </pre>

TABLE 1-2. Sync Server Operational Control Flow**Standard Delete**


1. Detect delete from the changelog table on the source.
2. Fetch the entry or table row from affected tables on the source.
3. Perform any mappings and compute the equivalent destination entry by constructing an equivalent LDAP entry or equivalent table row.
4. Fetch the entry or table row from affected tables on the destination.
5. The entry or table row exists on the destination.
6. Apply the delete on the destination.

**Delete After Source Entry Was Re-Added**

1. Detect delete from the changelog table on the source.
2. Fetch the entry or table rows from affected tables on the source.
3. Entry still exists on the source.
4. Delete request is dropped.

**Modify After Source Entry Is Deleted**

1. Detect change from the changelog table on the source.
2. Fetch the entry or table rows from affected tables on the source.
3. Entry does not exist.
4. Change request is dropped as the source entry no longer exists.



A Synchronization Example

Figure 1-11 shows a synchronization migration example from a Sun Directory Server Enterprise Edition (DSEE) topology to the UnboundID Directory Server topology with a change in the DIT structure to a flattened directory structure. The Sync Pipe connects the Sun Directory Server topology as the Sync Source and the UnboundID Directory Server topology as the Sync Destination. Each endpoint is defined with three external servers in their respective topology. The sync pipe destination has its base DN set to `o=example`, which is used when performing LDAP searches for entries.

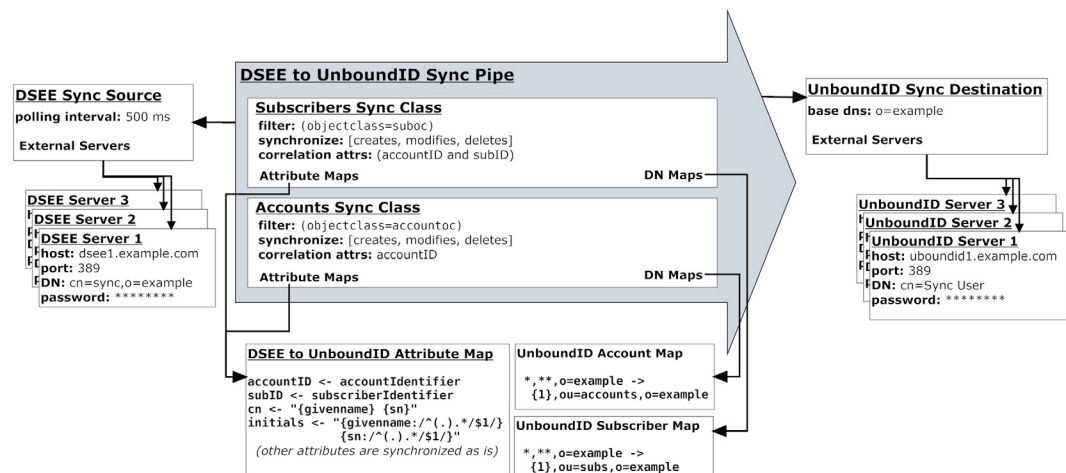
Two sync classes are defined: one for Subscribers, and one for Accounts. Each Sync Class uses a single "Sun DS to UnboundID Attribute Map" that has four attribute mappings defined. All other attributes are mapped as is.

Each sync class also defines its own DN Maps. For example, the Accounts Sync Class uses a DN Map, called UnboundID Account Map, that is used to flatten a hierarchical DIT, so that the Account entries appear directly under `ou=accounts`. The DN Map is as follows:

```
*,**,o=example -> {1},ou=accounts,o=example
```

With this mapping, if an entry DN has `uid=jsmith,ou=people,o=example`, then "*" matches `uid=jsmith`, "**" matches `ou=people`, and {1} matches `uid=jsmith`. Thus, `uid=jsmith,ou=people,o=example` gets mapped to `uid=jsmith,ou=accounts,o=example`. A similar map is configured for the Subscribers Sync Class.

FIGURE 1-11. A Typical Synchronization Topology Configuration

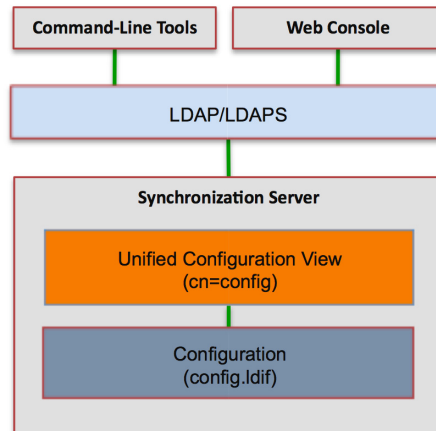


Available Tools Summary

The UnboundID Synchronization Server supports a flexible configuration framework that stores its settings in a flat file for a small memory footprint and easy access. Administrators can access the configuration using the UnboundID® Sync Management Console or using the

full suite of command-line tools in the **bin** directory for UNIX® or Linux® systems and the **bat** directory for Microsoft® Windows® systems. For detailed information and examples of the command-line utilities, see the *UnboundID Synchronization Server Command Line Tool Reference*.

FIGURE 1-12. UnboundID Synchronization Server Configuration Framework



You can view the Help information for each command-line tool by typing `--help` with the command (or type the short forms, `-?` and `-H`).

```
$ bin/resync -H
```

For those utilities that support additional subcommands (`dsconfig`, `dsframework`), you can also get more detailed subcommand information by typing `--help` with the subcommand:

```
$ bin/dsconfig list-log-publishers --help
```

Summary

The UnboundID Synchronization Server provides an excellent "total cost versus high performance" solution for your synchronization requirements. The server provides point-to-point, bidirectional synchronization with immediate failover and is ideal for large production environments that depend on highly available performance. This parallel, high-throughput, 100% Java solution requires little administrative cost and maintenance. UnboundID has leveraged its years of directory and synchronization expertise in the Identity Management industry to provide the ideal solution to meet your data center and enterprise synchronization needs.

2 Installing the Synchronization Server

Overview

You can begin the setup process using one of the UnboundID Synchronization Server's easy-to-use installation modes. Then, you can configure the Synchronization Server using the `create-sync-pipe-config` tool, the `dsconfig` command-line tool, or the UnboundID Sync Management Console if you prefer a graphical interface. Other instructions are provided to install redundant failover servers.

This chapter presents the various installation options and procedures available to administrators:

- [Before You Begin](#)
- [Unpacking the Installation Packages](#)
- [Installing the UnboundID Synchronization Server](#)
- [Installing the Synchronization Server in Non-Interactive Mode](#)
- [Running the Synchronization Server](#)
- [Stopping the Synchronization Server](#)
- [Uninstalling the Synchronization Server](#)
- [Installing the Sync Management Console](#)
- [Uninstalling the Console](#)
- [Updating the Synchronization Server](#)
- [Updating the Synchronization Server](#)
- [Upgrading the Console](#)
- [Setting the `--XX:-UseSSE42Intrinsics` Java Property](#)

Before You Begin

To begin the installation process, obtain the latest zip release bundle for the Synchronization Server and the Sync Management Console from UnboundID and unpack it in a folder of your choice.

Supported Operating Platforms

The UnboundID® Synchronization Server is a pure Java application and is certified VMWare Ready™. It is intended to run within the Java Virtual Machine on any Java 6 Standard Edition (SE) or Enterprise Edition (EE) certified platform including, but not limited, to the following:

- ❑ RedHat® Enterprise Linux® 5.5 (natively or running under VMware® ESX 4.0 Update 2)
- ❑ Oracle® Solaris™ 10 update 9
- ❑ RedHat®/CentOS 5.5 & 6
- ❑ Windows Server® 2003 & 2008
- ❑ Novell® SUSE® 11.3

Software Requirements: Java 6

For optimized performance, the UnboundID Synchronization Server requires the Java 6 for 32-bit and 64-bit architectures, respectively, depending on your system requirements. We recommend JDK1.6.0_20.

Even if your system already has Java installed, you might want to create a separate Java 6 installation for use by the UnboundID Synchronization Server to ensure that updates to the system-wide Java installation do not inadvertently impact the Synchronization Server. This setup requires that the JDK, rather than the JRE, for both the 32-bit and 64-bit versions or both depending on your operating system be installed on your system.

On Solaris™ systems, if you want to use the 64-bit version of Java, you need to install both the 32-bit and 64-bit versions. The 64-bit version of Java on Solaris is not a full stand-alone installation, but instead relies on a number of files provided by the 32-bit installation. Therefore, the 32-bit version should be installed first, and then the 64-bit version installed in the same location with the necessary additional files.

On other platforms (Linux, Microsoft Windows), the 32-bit and 64-bit versions of Java each contain complete installation packages. If you only want to run the 64-bit version of Java, installing the 32-bit JDK is not necessary. If you want both versions installed on your Linux or Windows platforms, then they should be installed in separate directories, because the files cannot coexist in the same directory as they can on Solaris systems.

Important

We have found that an existing Sun JVM bug (i.e., use of `String.indexOf()` intrinsic in SSE 4.2) in JDK1.6u21 and JDK1.6u22 causes promotion failures that lead to extended stop-the-world garbage collection pauses. We recommend that JDK1.6u20 be used with the Synchronization Server and that the use of JDK1.6u21 and JDK1.6u22 *not* be implemented. If your current system uses JDK1.6u18 to JDK1.6u20, then you must also configure an additional Java property for the `start-sync-server` command that bypasses this bug. See “Setting the `--XX:-UseSSE42Intrinsics` Java Property” on page 37.

Unpacking the Installation Packages

To begin the installation process, you can obtain the latest zip release bundle from UnboundID and unpack it in a folder of your choice. In this example, the release bundle unpacks in the `UnboundID-Sync` directory.

```
$ unzip UnboundID-Sync-<release-version>.zip
```

You can now install the Synchronization Server.

Installing the UnboundID Synchronization Server

One of the strengths of UnboundID Synchronization Server is the ease with which you can install a synchronization server instance using the `setup` tool. To install a stand-alone synchronization server instance, run the `setup` tool in one of the two following modes: interactive command-line mode or non-interactive command-line mode.

- **Interactive Command-Line Mode.** Interactive command-line mode prompts for information during the installation process. To run the installation in this mode, use the `setup` command.
- **Non-interactive Command-Line Mode.** Non-interactive command-line mode is designed for quick command-line entry or for use in install scripts to automate the process. To run the installation in this mode, the `setup` tool must be run with the `--no-prompt` option as well as all other arguments required to define the appropriate initial configuration.

All of the steps that are executed during the installation and configuration process should be performed while logged into the system as the user or role under which the Synchronization Server will run.

Installing the Synchronization Server

The `setup` tool provides an interactive text-based interface to install a Synchronization Server instance. You can install the Synchronization Server by entering the required input as presented by the prompts.

To Install the Synchronization Server in Interactive Command-Line Mode

1. Go to the installation directory if you are not already there.

```
$ cd UnboundID-Sync
```

2. Use the `setup` command to install the Synchronization Server.

```
$ ./setup
```

Note

If your **JAVA_HOME** environment variable is set to an older version of Java, you must explicitly specify the path to the Java 6 JDK installation during setup. You can either set the **JAVA_HOME** environment variable with the Java 6 JDK path or execute the **setup** command in a modified Java environment using the **env** command.

```
$ env JAVA_HOME=/java/jdk1.6.0 ./setup
```

3. Read the UnboundID End-User License Agreement. If you agree to its terms, type **yes** to continue.
4. Next, you will be prompted to add this server to an existing Synchronization Server topology. If you are adding this Synchronization Server instance to an existing topology, type **yes**. (See “Updating the Synchronization Server” on page 33. for more information on installing redundant servers.) Otherwise, press Enter or Return to accept the default (**no**).
5. Type the initial root user DN and password for the Synchronization Server, or press Enter or Return to accept the default (**cn=Directory Manager**), and then type and confirm the root user password.
6. Next, type the listener port number of the Synchronization Server, or press Enter or Return to accept the default port (**389**). If you did not log in as a root user, you will see port 1389.
7. Next, you will be prompted if you want to enable SSL or StartTLS, type **yes** to enable one or both. Otherwise, press Enter or Return to accept the default (**no**).

Note

If you plan to synchronize passwords with Active Directory systems, you must configure the Synchronization Server to accept SSL connections. The Synchronization Server accepts JKS, PKCS#12 certificates, and PKCS#11 tokens. Afterwards, you need to install the UnboundID Password Sync Agent.

If you typed **yes** for SSL, StartTLS, or both, you will be prompted for the SSL port number and certificate options. For the SSL port, enter a port number of your choice, or accept the default. If you use the Java or the PKCS#12 key store, you will be asked for the key store path, and the key store PIN. If you use the PKCS#11 token, you will be asked for only the key store PIN. Enter the options as seen below if you want to enable SSL and StartTLS.

8. By default, the Synchronization Server listens on all available network interfaces for client connections. If this is acceptable, you can skip this step. If you want to limit the client connections to specific host names or IP addresses, type **yes** at the prompt, and then, enter the host name or IP address. You will be prompted again to enter another host name or IP address. When you are done, press Enter or Return to continue.
9. Next, you will be prompted if you want to allocate the maximum amount of memory to the JVM heap for maximized performance. This option should only be selected if the Synchronization Server is the primary application and no other processes consume a significant amount of memory. Type **yes** if you want to tune the memory allocated to the JVM for maximized performance. If you do not want to maximize the memory, press Enter or Return to accept the default (**no**).

If you selected JVM memory tuning, enter the maximum amount of memory to be allocated to the Synchronization Server. For example, you can enter '64m', '1g', or '16g', etc.

10. Next, type **yes**, or press Enter or Return to accept the default to start the Synchronization Server after the configuration has completed. If you plan to configure additional settings or import data, you can type **no** to keep the server in shutdown mode.
11. On the Setup Summary page, confirm the configuration, and press Enter or Return to set up the Synchronization Server. The configuration is recorded in the `/server-root/logs/setup.log` file.

12. At this point, you have the following options depending on your specific configuration:

- Build a configuration if the command-line wizard was not employed.
- Determine if a bulk import or resync is necessary.
- Determine if a `realtime-sync set-startpoint` is necessary.
- Enable syncing with `realtime-sync start`.

See “Configuring the Synchronization Server Using the Command-Line Wizard” on page 42 to continue.

Installing the Synchronization Server in Non-Interactive Mode

You can run the `setup` command in non-interactive mode to automate the installation process using a script or to run the command directly from the command line. If there is a missing or incorrect argument, the `setup` tool fails and aborts the process.

To Install the Server in Non-Interactive Mode

Use `setup` with the `--no-prompt` option. The command uses the bind DN (`cn=Directory Manager`) and default LDAP port 389.

```
$ env JAVA_HOME=/sync/java ./setup --no-prompt \  
  --rootUserPassword "password" --acceptLicense
```

If you have already configured a trust store, you can also use the `setup` tool to enable security. The following example enables security, both SSL and StartTLS. It also specifies a JKS key-store and truststore that define the server certificate and trusted CA. The passwords for the keystore files are defined in the corresponding `.pin` files, where the password is written on the first line of the file. The values in the `.pin` files will be copied to the `server-root/config` directory in the `keystore.pin` and `truststore.pin` files.

Note that the password to the private key within the key store is expected to be the same as the password to the key store. If this is not the case, the private key password can be defined within the Sync Management Console or `dsconfig` by editing the Trust Manager Provider standard configuration object.

```
$ env JAVA_HOME=/sync/java ./setup --cli --no-prompt \  
  --bindDN "cn=Directory Manager" \  
  --rootUserPassword "password" --ldapPort 389 \  
  --enableStartTLS --ldapsPort 636 \  
  --useJavaKeystore /path/to/devkeystore.jks \  
  --keyStorePasswordFile /path/to/devkeystore.pin \  
  --certNickName server-cert \  
  --useJavaTrustStore /path/to/devtruststore.jks \  
  --trustStorePasswordFile /path/to/devtruststore.pin \  
  --acceptLicense
```

Running the Synchronization Server

To start the Synchronization Server, run the `bin/start-sync-server` command on UNIX or Linux systems, or `bat\start-sync-server` on Microsoft Windows systems. The `start-sync-server` command starts the Synchronization Server as a background process when no other options are specified.

To Start the Synchronization Server

Go to the installation directory, and then use `bin/start-sync-server`.

```
$ bin/start-start-sync-server
```

To Start the Synchronization Server as a Foreground Process

1. Type `bin/start-sync-server` with the `--nodetach` option to launch the Synchronization Server as a foreground process.

```
$ bin/start-sync-server --nodetach
```

2. You can stop the process by pressing `ctrl+c` in the terminal window where the server is running or by running the `stop-sync-server` utility from another window.

Automatically Starting the Synchronization Server at Boot Time

By default, the UnboundID Synchronization Server does not start automatically when the system is booted. Instead, you must manually start it with the `bin/start-sync-server` command. To configure the Synchronization Server to start automatically when the system boots, use the `create-rc-script` tool to create a run control (RC) script.

To Configure the Synchronization Server to Start at Boot

1. Create the startup script.

```
$ bin/create-rc-script --outputFile unboundid-sync.sh --userName sync
```

2. As a root user, move the generated `unboundid-sync.sh` script into the `/etc/init.d` directory, and create symlinks to it from the `/etc/rc3.d` (starting with an “S” to ensure that the server is started) and `/etc/rc0.d` (starting with a “K” to ensure that the server is stopped).

```
# mv unboundid-sync.sh /etc/init.d
# ln -s /etc/init.d/unboundid-sync.sh /etc/rc3.d/S50-unboundid-sync.sh
# ln -s /etc/init.d/unboundid-sync.sh /etc/rc0.d/K50-unboundid-sync.sh
```

Note

Some Linux implementations may not like the "-" in the scripts. If your scripts do not work, try renaming the scripts without the dashes. You can also try symlinking the S50* file into the `/etc/rc3.d` and/or `/etc/rc5.d` directory, based on whatever runlevel the server enters when it starts. Some Linuxes do not even use `init.d`-style startup scripts, so depending on the flavor of Linux that you are using, you might have to put the script somewhere else or use some other mechanism for having it launched at startup.

3. Log out as root, and re-assume the **sync** role if you are on a Solaris system.

Stopping the Synchronization Server

The Synchronization Server provides a simple shutdown script, `bin/stop-sync-server`, to stop the server. You can run it manually from the command line or within a script.

To Stop the Synchronization Server

Change to the installation directory if you are not there, and then use the **stop-sync-server** tool.

```
$ bin/stop-sync-server
```

To Schedule a Synchronization Server Shutdown

Use the `bin/stop-sync-server` tool with the `--stopTime YYYYMMDDhhmmss` option to schedule a server shutdown. The Synchronization Server schedules the shutdown and sends a notification to the `server.out` log file. The following example sets up a shutdown task that is scheduled to be processed on October 24, 2008 at 3:45 P.M. local time. The server uses UTC time if the provided timestamp includes a trailing "Z", for example, 2008102415450Z.

```
$ bin/stop-sync-server --port 389 \
  --bindDN "cn=Directory Manager" --bindPassword secret \
  --stopTime 20081024154500
```

Restarting the Synchronization Server

You can restart the Synchronization Server using the **stop-sync-server** command with the `--restart` or `-R` option. Running the command is equivalent to shutting down the server, exiting the JVM session, and then starting up again. To avoid destroying and re-creating the JVM, use an internal restart, which can be issued over LDAP. The internal restart will keep the same Java process and avoid any changes to the JVM options.

To Restart the Synchronization Server

Go to the installation directory. Using a loop back interface, run the `stop-sync-server` command with the `-R` or `--restart` options.

```
$ bin/stop-sync-server --restart --hostname 127.0.0.1 --port 389 \  
--bindDN "cn=Directory Manager" --bindPassword secret
```

Uninstalling the Synchronization Server

The Synchronization Server provides an `uninstall` command-line utility for quick and easy removal of the code base. You can uninstall the Synchronization Server using one of the following modes:

- **Interactive Command-Line Mode.** Interactive command-line mode is a text-based interface that prompts the user for input. You can start the command using the `bin/uninstall` command. The utility prompts you for input if more data is required.
- **Non-Interactive Command-Line Mode.** Non-interactive mode suppresses progress information from being written to standard output during processing, except for fatal errors. This mode is convenient for scripting and is invoked using the `bin/uninstall` command with the `--no-prompt` option.

Note

For stand-alone installations with a single Synchronization Server instance, you can also manually remove the Synchronization Server by stopping the server and recursively deleting its directory and subdirectories using:

```
$ rm -rf UnboundID-Sync
```

Uninstalling the Synchronization Server in Interactive Command-Line Mode

Interactive mode uses a text-based, command-line interface to help you remove your instance. If the `uninstall` tool cannot remove all of the synchronization server files, the server generates a message with a list of the files and directories that must be manually deleted. The `uninstall` command must be run as either the root user or the same user (or role) that installed the Synchronization Server.

To Uninstall the Synchronization Server in Interactive Command-Line Mode

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

2. Use the `uninstall` command.

```
$ ./uninstall
```

3. Select the components to be removed. If you want to remove all components, press Enter or Return to accept the default. Otherwise, enter **1** to remove specific components.
4. If your synchronization server is running, press Enter or Return to shutdown the server before continuing the uninstall process.
5. Complete the uninstall, and view the logs for any remaining files. Manually remove any remaining files or directories, if required.

Uninstalling the Synchronization Server in Non-Interactive Mode

The **uninstall** tool provides a non-interactive mode that suppresses progress information during processing, except for fatal errors. You can use this option by using the **--no-prompt** option. Another useful argument is the **--forceOnError** option that continues the uninstall process when an error is encountered. If an option is incorrectly entered or if a required option is omitted and if the **--forceOnError** option is not used, the command will fail and abort.

To Uninstall the Synchronization Server in Non-Interactive Mode

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

2. Use **uninstall** with the **--remove-all** option to remove all of the Synchronization Server's libraries. The **--quiet** option suppresses output information and is optional.

```
$ ./uninstall --remove-all --no-prompt --quiet --forceOnError
```

3. If any files or directories remain, manually remove them.

To Uninstall Selected Components in Non-Interactive Mode

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

2. Use **uninstall** with the **--log-files** option to remove the Synchronization Server's backup files. Use the **--help** or **-H** option to view the other options available to remove specific components.

```
$ ./uninstall --log-files --no-prompt --quiet --forceOnError
```

Installing the Sync Management Console

UnboundID provides a web-based graphical Sync Management Console that administrators can use to configure and monitor the Synchronization Server. UnboundID distributes the con-

sole in zip distribution together with the Synchronization Server zip file. The supplied war file should work with any J2EE 5 web container that supports Java Server Faces version 1.2. The following installation instructions are specific to the Apache Tomcat Application Server version 6 but can be easily adapted to other web containers.

Note

When managing multiple instances in a highly-available environment, the console automatically discovers other instances in the topology by reading the topology information stored under the `cn=admin data` entry.

To Install the Console

1. Download and install Apache Tomcat version 6 or later. For example, download `apache-tomcat-<version>.zip` from <http://tomcat.apache.org/>, and then unzip the file in a location of your choice.
2. Download the Sync Management Console zip file, `UnboundID-Sync-web-console-<version>.zip` and unzip the file on your server. You should see the following files:

```
3RD-PARTY-LICENSE.TXT
LICENSE.TXT
README
dsconsole.war
```

3. If the Synchronization Server is listening for LDAP connections on port 389 on the same machine as the console, then copy `dsconsole.war` to `apache-tomcat-<version>/webapps` and go to step 7.
4. Create a `sync` directory in `apache-tomcat-<version>/webapps`. Then, copy the `dsconsole.war` file to `apache-tomcat-<version>/webapps/sync`.

```
$ mkdir /apache-tomcat-<version>/webapps/sync
$ cp dsconsole.war /apache-tomcat-<version>/webapps/sync
```

5. Go to the `apache-tomcat-<version>/webapps/sync` directory to extract the contents of the console. The `jar` command is included with the JDK.

```
$ cd /apache-tomcat-<version>/webapps/sync
$ jar xvf dsconsole.war
```

6. Edit the `sync/WEB-INF/web.xml` file to point to the correct directory instance. Change the host and port to match your directory. For example, if your system's host name is `server1.example.com` and the LDAP listener port is `1389`, change those lines in the `<context-param>` section. The defaults are `localhost` for the host name and `389` for the LDAP port.

```
<context-param>
  <param-name>host</param-name>
  <param-value>localhost</param-value>
</context-param>
<context-param>
  <param-name>port</param-name>
  <param-value>389</param-value>
</context-param>
```

7. With the out-of-the-box configuration, Tomcat times out sessions after 30 minutes of inactivity forcing the user to log back in again. This can be changed on an application server wide basis by editing `apache-tomcat-<version>/conf/web.xml`, and updating the value of this configuration parameter:

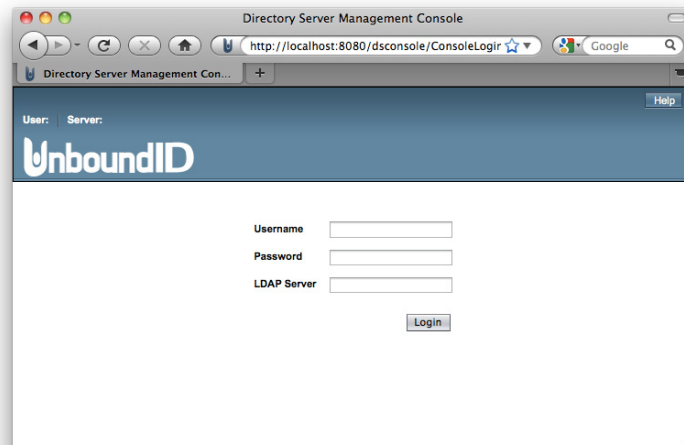
```
<session-config>
  <session-timeout>120</session-timeout>
</session-config>
```

Tomcat will expire sessions after the specified number of minutes. Changing the value to 120, for example, will extend the expiration to two hours. Changes to this setting might not take effect until the Tomcat instance is restarted, so consider changing the value before starting Tomcat for the first time.

8. Start the Synchronization Server if it is not already running (`bin/start-sync-server` on UNIX /LINUX or `bat\start-sync-server` on Windows), and then start the Sync Management Console using the `apache-tomcat-<version>/bin/startup.sh` script. Use `shutdown.sh` to stop the application server. (On Windows, use `startup.bat` and `shutdown.bat`.) Note that the `JAVA_HOME` environment variable must be set to specify the location of the Java installation to run the server.

```
$ env JAVA_HOME=/sync/java bin/startup.sh
Using CATALINA_BASE:   /apache-tomcat-6.0.18
Using CATALINA_HOME:   /apache-tomcat-6.0.18
Using CATALINA_TMPDIR: /apache-tomcat-6.0.18/temp
Using JRE_HOME:        /sync/java
```

9. Open a browser to `http://hostname:8080/sync`. The application server listens to port 8080 for HTTP requests. Enter the bind DN, bind DN password, and server hostname or IP address and port to log on (for example, `server1.example.com:389`).



Uninstalling the Console

You can manually remove the console by deleting the **webapps/sync** directory. Make sure to close the console prior to removing it.

To Uninstall the Console

1. Close the Sync Management Console, and shut down the application server. (On Windows, use **shutdown.bat**).

```
$ /apache-tomcat-<version>/bin/shutdown.sh
```

2. Remove the **webapps/sync** directory.

```
$ rm -rf webapps/sync
```

3. Restart the Tomcat instance, if necessary. Alternatively, if no other applications are installed in the Tomcat instance, then the entire Tomcat installation can be removed by deleting the **apache-tomcat-<version>** directory tree.

Updating the Synchronization Server

UnboundID issues new software builds periodically and distributes the software package in zip format. Administrators can use the Synchronization Server's **update** utility to upgrade the current server code with the latest features and bug fixes. To update the Synchronization Server to a newer version, download the build package, and then unzip the new server package on the same host as the server that you wish to update. Before upgrading a server, you should ensure that it is capable of starting without severe or fatal errors.

During an update process, the updater checks a manifest file that contains an MD5 checksum of each file in its original state when installed from zip. Next, it compares the checksum of the new server files to that of the old server. Any files that have different checksums will be updated. For files that predate the manifest file generation, the file is backed up and replaced. The updater also logs all file changes in the history directory to tell what files have been changed.

For schema updates, the **update** tool preserves any custom schema definitions (**99-user.ldif**). For any default schema element changes, the updater will warn the user about this condition and then create a patch schema file and copy it into the server's schema directory. For configuration files, the **update** tool preserves the configuration file, **config.ldif**, unless new configuration options must be added to the Synchronization Server.

Once the updater finishes its processing, it checks if the newly updated server starts without any fatal errors. If an error occurs during the update process, the **update** tool reverts the server root instance to the server state prior to the update.

The updater also upgrades the Password Synchronization Agent plug-in to its latest version automatically. Any software updates to the PSA plug-in will be included with the new Synchronization Server zip file.

To Update the Synchronization Server

Assume that an existing version of the Synchronization Server is stored at **UnboundID-Sync-old**, which you want to update.

1. Make sure you have complete, readable backup of the existing system before upgrading the Synchronization Server build. Also, make sure you have a clear backout plan and schedule.
2. Download the latest version of the UnboundID Synchronization Server software and unzip the file. You can use the **unzip -f/--file** option to specify a location on the server. For this example, let's assume the new server is located in the **UnboundID-Sync-new** directory.
3. Check the version number of the newly downloaded synchronization server instance using the **--version** option on any command-line utility. For example, you should see the latest revision number.

```
$ UnboundID-Sync-new/setup --version
UnboundID Synchronization Server 3.0.0
Build 20100416162859Z
Revision 5905
```

4. Use the **update** tool of the newly unzipped build to update the Synchronization Server code. Make sure to specify the Synchronization Server instance that you are upgrading with the **--serverRoot** option.

```
$ UnboundID-Sync-new/update --serverRoot UnboundID-Sync-old
```

Note

The UnboundID Synchronization Server provides a web console, the UnboundID Synchronization Management Console, to configure and monitor the server. If you update the Synchronization Server version, you should also update the Sync Management Console.

5. View the log file to see which files were changed. The log file is located in the **<server-root>/history** directory. The file will be labelled with the Synchronization Server version number and revision.

```
$ view <server-root>/history/1272307020420-3.0.0.6011/update.log
```

Reverting an Update

Once the Synchronization Server has been updated, you can revert to the most recent version (one level back) using the **revert-update** tool. The **revert-update** tool accesses a log of file actions taken by the updater in order to put the filesystem back to its prior state. If you have run multiple updates, you can run the tool multiple times to revert to each prior update sequentially. You can only revert back one level. For example, if you ran the update twice since first

installing the Synchronization Server, you can run the update command to revert to its previous state, then run the update command again to return to its original state.

To Revert to the Most Recent Server Version

Use `revert-update` in the server root directory revert back to the most recent version of the server.

```
$ UnboundID-Sync-new/revert-update
```

Upgrading the Console

The `update` utility does not automatically upgrade the Sync Management Console. Therefore, you must upgrade the console manually by replacing the `web.xml` file.

To Upgrade the Console

1. Shut down the application server.
2. In the current deployment of the Sync Management Console, move the `<webapps>/sync/WEB-INF/web.xml` file to another temporary location.
3. Download and deploy the latest version for the Sync Management Console. Follow steps 2–6 in the section “To Install the Console” on page 31.
4. Assuming you had not renamed the war file when you originally deployed the Sync Management Console, replace the newly deployed Sync Management Console’s `web.xml` to `<webapps>/sync/WEB-INF/web.xml`.
5. Start the application server.

Installing a Redundant Failover Server

The UnboundID Synchronization Server supports multiple redundant failover servers that automatically become active when the main synchronization server is down for any reason. Only one Synchronization Server instance is active at any time, but multiple redundant servers can be present in the topology in a configurable prioritized order.

Before you install a redundant failover server, you must have already installed and configured a Synchronization Server instance. When installing the redundant server, the installer will copy the first Synchronization Server’s configuration, including external server setup, sync pipes, sync classes, DN and attribute maps.

To Install a Redundant Failover Server

Before you install the redundant failover server, you should already have an existing Synchronization Server instance configured and running.

1. Unpack the UnboundID Synchronization Server zip build. Make sure you name the unpacked directory to something other than the first server instance directory.

```
$ unzip UnboundID-Sync-<version>.zip -d sync2
```

2. Go to the installation directory if you are not already there.

```
sync2.example.com> $ cd UnboundID-Sync
```

3. Following steps 2–12 in “Installing the UnboundID Synchronization Server” on page 23, except in step 4, type **yes** to add the server to an existing topology. If you are using the **setup** tool in non-interactive mode, use the following command:

```
$ ./setup --localHostName sync2.example.com --ldapPort 8389 \  
--masterHostName sync1.example.com --masterPort 7389 \  
--masterUseNoSecurity --acceptLicense --rootUserPassword password \  
--no-prompt
```

The secondary server is now ready to take over as a primary server in the event of a failover. As a result, no realtime-sync invocations are needed for this server.

4. Verify the configuration by using the **bin/status** tool. Note the Priority Index associated with each Synchronization Server instance. The Synchronization Server with the lowest **priority-index** number has the highest priority.

```
$ bin/status --bindPassword secret
```

```
... (status output) ...
```

```
          --- Sync Topology ---  
Host:Port                : Status      : Priority Index  
-----  
sync1.example.com:389 (this server) : Active      : 1  
sync2.example.com:389              : Unavailable  : 2
```

5. To change the Priority Index of the Synchronization Server, use **bin/dsconfig**.

```
$ bin/dsconfig set-external-server-prop \  
--server-name sync2.example.com:389 \  
--set sync-server-priority-index:1
```

Note

To change the priority index interactively, use **bin/dsconfig**. First, enable the Advanced Objects menu. Next, on the UnboundID Synchronization Server configuration console main menu, select External Server, select View and Edit, then select the Synchronization Server instance. Finally, on the Sync Server External Server menu, select the **sync-server-priority-index** property and change it to a value of your choice. Remember, the lower priority-index number has the higher priority (e.g., "1" has the highest priority).

Setting the `--XX:-UseSSE42Intrinsics` Java Property

We have found that a known Sun (now Oracle) JVM bug in versions 1.6u21 and 1.6u22 (i.e., use of implementation `String.indexOf()` intrinsic in SSE 4.2) results in promotion failures that lead to extended stop-the-world garbage collection pauses. We highly recommend that users use JDK 1.6u20 with the Synchronization Server and not install 1.6u21 and 1.6u22. The use of JDK 1.6u20 also has an additional requirement that the Java property, `--XX:-UseSSE42Intrinsics` be set for the `start-sync-server` command in the directory server's Java properties file. The Java property avoids the `String.indexOf()` bug.

To Configure the `UseSSE42Intrinsics` Java Property

1. Open the `config/java.properties` file with a text editor.

2. Add the `--XX:-UseSSE42Intrinsics` flag to the `start-sync-server` command.

```
# These JVM arguments can be used to run the Synchronization Server with an aggressive memory tuning:
# start-sync-server.java-args=-d64 -server -Xmx2g -Xms2g
-XX:+UseConcMarkSweepGC -XX:+CMSConcurrentMTEnabled -XX:+CMSParallelRemarkEnabled
-XX:+CMSParallelSurvivorRemarkEnabled -XX:+CMSScavengeBeforeRemark
-XX:RefDiscoveryPolicy=1 -XX:ParallelCMSThreads=1
-XX:CMSMaxAbortablePrecleanTime=3600000 -XX:CMSInitiatingOccupancyFraction=80
-XX:+UseParNewGC -XX:+UseMembar -XX:+UseBiasedLocking -XX:+UseCompressedOops
-XX:+HeapDumpOnOutOfMemoryError -XX:-UseSSE42Intrinsics
```

3. Save the file.
4. Run `bin/dsjavaproperties` to apply the change. If this command is not run, you will see a warning to run the command.

```
$ bin/dsjavaproperties
```

3 Configuring the Synchronization Server

Overview

The UnboundID Synchronization Server provides a comprehensive suite of command-line tools and a graphical Sync Management Console that accesses the underlying Synchronization Server configuration framework. The configuration is stored as a flat file (LDIF format) in the `cn=config` branch. Administrators can use the tools to configure a single server instance or server groups remotely or locally. All configuration changes to the server and their equivalent reversion commands are recorded in the `config-audit.log`.

Before setting up the Synchronization Server, review the section “Configuration Model” on page 13 to understand the important components of the Synchronization Server.

This chapter presents the following topics:

- [Pre-Deployment Checklist](#)
- [Creating Administrators](#)
- [Configuring the Synchronization Server](#)
- [Configuring the Synchronization Server Using the Command-Line Wizard](#)
- [Configuring the Synchronization Server Using the Management Console](#)
- [Using the dsconfig Configuration Tool](#)
- [Configuring the Synchronization Server Using dsconfig](#)
- [Preparing the Synchronization Server for External Server Communication](#)
- [Using Bulk Resync on the Synchronization Server](#)
- [Controlling Real Time Synchronization](#)
- [Configuring Attribute Maps](#)
- [Configuring DN Maps](#)
- [Configuring Fractional Replication](#)
- [Managing Failover Behavior](#)

Pre-Deployment Checklist

Prior to any deployment, you must determine the configuration parameters necessary for your Synchronization topology. Answer the following questions and record them prior to configuring your Synchronization Server instance(s).

External Servers

External Server Type. What type of external servers are you using in the Synchronization topology: UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Sun Directory Server (5.x and above), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun Directory Server Enterprise Edition (DSEE 6.x, 7.x), Microsoft Active Directory, Oracle (10g, 11g), Microsoft SQL Server (2005, 2008).

LDAP Connection Settings. What is the host, port, bind DN, and bind password for each external server instance(s) that you want included in the Synchronization topology?

Security and Authentication Settings. If the external server instance uses a secure connection, does it use SSL or StartTLS? What authentication method does the external server use: none, simple, external (i.e., SASL mechanisms)? If you are synchronizing to or from an Active Directory system, you must establish an SSL or StartTLS connection to the Synchronization Server.

Sync Pipes

A *Sync Pipe* defines a single synchronization path between the Source and Destination targets. You will need one Sync Pipe for each point-to-point synchronization path that you define for a topology of source servers to a topology of destination servers. Answer the following questions.

Sync Source. Which external server is the Sync Source for the Synchronization topology? You can define a priority order if more than one external server is defined as a Sync Source for failover purposes.

Sync Destination. Which external server is the Sync Destination for the Synchronization topology? You can define a priority order if more than one external server is defined as a Sync Destination for failover purposes.

Sync Classes

For each Sync Pipe defined, you must define one or more Sync Classes. A Sync Class defines how attributes and DNs are mapped and how Source and Destination entries are correlated. Questions required to define a Sync Class are as follows:

Evaluation Ordering. If you will be defining more than one Sync Class, what is the evaluation order of each Sync Class?

Sync Classes are evaluated according to the `evaluation-order-index` property and the criteria used to identify the first matching Sync Class. When there is an overlap between criteria used to identify a Sync Class, the Sync Class with the most specific criteria will be used first.

Base DNs. Are entries in the Sync Class only under specific base DNs?

Include Filters. What are the search filters to be used to search for entries in the Sync Source?

Synchronized Entry Operations. Which types of operations on entries should be synchronized: creates, modifications, and/or deletes?

DNs. What are the differences between the DNs from the Sync Source topology to the Sync Destination topology? Are there structural differences in terms of the Directory Information Tree (DIT) between the Sync Source and the Sync Destination? For example, does the Sync Source use a Nested DIT versus a Flattened DIT? Does the Sync Destination use a corresponding DIT as the Sync Source (i.e., a Nested DIT versus a Flattened DIT)?

Destination Correlation Attributes. Which set of Sync Destination attributes in an entry should be used to correlate an entry in the Sync Source? For example, the default option is to use the `dn` attribute to correlate entries, but you could specify that the `dn` and `uid` attributes be used, or the `cn` and `employeeNumber` attributes, or others, depending on how the entries are structured in the Sync Source and Sync Destination, respectively.

Attributes. What are the differences between the attributes from the Sync Source to the Sync Destination? Some questions related to attributes are as follows:

- **Automatically Mapped Source Attributes.** Are there attributes that can be automatically synchronized with the same name at the Sync Source to Sync Destination? For example, can you set direct mappings for `cn`, `uid`, `telephoneNumber`, or for all attributes?
- **Non-Auto Mapped Source Attributes.** Are there some attributes that should not be automatically mapped from the Sync Source to Sync Destination? For example, the Sync Source may have an attribute, `employee`, while the Sync Destination may have a corresponding attribute, `employeeNumber`. If an attribute is not automatically mapped, then an Attribute Mapping must be provided if it is to be synchronized.
- **Attribute Mappings.** How are attributes mapped from the Sync Source to the Sync Destination? (For example, are they mapped directly, mapped based on attribute values, or mapped based on attributes that store DN values?)

Creating Administrators

The UnboundID Sync Management Console does not persistently store any credentials for authenticating to the Synchronization Server but uses the credentials provided by the user when logging in. When managing multiple synchronization server instances, the provided credentials must be valid for each instance. Therefore, assuming you have multiple synchronization servers—the main server and a failover—if you change an admin user on the main synchronization server instance, you must make the same change on the other server instance. Likewise, if you have multiple synchronization servers, you must make any changes manually at each server instance.

To Create an Administrator

To log into the console, you can either use a root user DN or create a new administrator user ID. The `dsframework` command can be used to create a user ID, for example:

```
$ bin/dsframework create-admin-user --hostname server1.example.com \  
--port 1389 --bindDN "cn=Directory Manager" \  
--bindPassword secret --userID someAdmin --set password:secret
```

Once you have set up a new admin account, the administrator can log in to the Sync Management Console using the user ID short form "someAdmin" or the full DN, "cn=someAdmin,cn=Administrators,cn=Admin Data".

Configuring the Synchronization Server

The UnboundID Synchronization Server configuration can be accessed and modified in the following ways:

- **Using Synchronization Command-Line Tools.** The UnboundID Synchronization Server provides a command-line tool to quickly configure a Synchronization Server topology: **create-sync-pipe-config**. The tool records the configuration in a batch file and optionally allows you to apply the configuration to a local Synchronization Server. Once the batch file is applied to the server, the Synchronization Server records those changes in the `<server-root>/sync-pipe-cfg.txt` file, which can be re-applied to other servers. Another tool, **resync**, is used to verify that everything is in-sync after synchronization has started or used in bulk synchronization mode to initially populate a target directory or database. The **realtime-sync** tool is used to start synchronization immediately, at a specified point at a changelog event, or at a specified time duration ago.
- **Using the Sync Management Console.** The UnboundID Synchronization Server provides a web-based console for graphical server management and monitoring. The console provides equivalent functionality as the **dsconfig** command for viewing or editing configurations. All configuration changes using this tool are recorded in `logs/config-audit.log`, which also has the equivalent reversion commands should you need to back out of a configuration.
- **Using the dsconfig Command-Line Tool.** The **dsconfig** tool is a text-based menu-driven interface to the underlying configuration. The tool runs the configuration using three operational modes: interactive command-line mode, non-interactive command-line mode, and batch mode. All configuration changes made using this tool are recorded in `logs/config-audit.log`. If you are configuring a Sync Pipe from scratch, we recommend using the **create-sync-pipe-config** tool as it will lead you through the steps necessary to define each Sync Pipe component.

Configuring the Synchronization Server Using the Command-Line Wizard

The UnboundID Synchronization Server provides useful command-line tools to configure your synchronization topology. To set up your topology, use the **create-sync-pipe-config**

tool, which is a command-line wizard that will prompt you for input when setting up your Sync Pipes and Sync Classes. After you have configured your topology, you can verify that everything is in-sync using the **resync** tool.

About the Sync User Account

During the configuration process, the Synchronization Server sets up a Sync User Account DN on each external server. The account (by default, **cn=Sync User**) is used exclusively by the Synchronization Server to communicate with the endpoint external servers. The entry is important in that it contains the credentials (DN and password) used by the Synchronization Server to access the source and target servers. The Sync User account resides in different entries depending on the targeted system:

- For UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), the Sync User Account resides in the configuration entry (e.g., **cn=Sync User, cn=Root DNs, cn=config**).
- For Sun Directory Server and Sun DSEE, the Sync User account resides under the base DN in the **userRoot** backend (e.g., **cn=Sync User, dc=example, dc=com**). We also recommend that the Sync User account NOT be in the **cn=config** branch for Sun Directory Server and DSEE machines. If it resides there, delete it, and then add it to the normal backend (**dc=example, dc=com**) and update the configuration in the Sync Server.
- For Microsoft Active Directory servers, the Sync User account resides in the Users container (e.g., **cn=Sync User, cn=Users, DC=adsync, DC=unboundid, DC=com**).
- For Oracle and Microsoft SQL Servers, the Sync User account is a login account (SyncUser) with the sufficient privileges (for example, Resource and Connect) to access the tables to be synchronized.

Although in most cases, modifications to this account will never take place, you can ensure that the entry never gets synchronized by setting up an optional Sync Class if your Sync User account resides in the **userRoot** backend (Sun Directory Server or Sun DSEE) or Users container (Microsoft Active Directory). For example, you can configure this Sync Class to have all **CREATE**, **MODIFY**, and **DELETE** operations set to false, so that the Sync User Account never gets synchronized with the other user entries.

Assumptions

The following example shows a bidirectional synchronization deployment between a Sun Directory Server 5.x topology and an UnboundID Directory Server topology. The example assumes that you have two replicated topologies configured: the first endpoint topology consists of two Sun Directory Server LDAP servers (version 5.2 patch 4): the main server and one failover. Both Sun Directory Servers 5.x have their Retro Change logs enabled and contains the full DIT that will be synchronized to the second endpoint. The second endpoint topology consists of two UnboundID Directory Servers (version 3.x): the main server and one failover.

Both UnboundID Directory Servers have their change logs enabled and contain entries similar to the first endpoint servers, except that it uses a **mail** attribute, instead of an **email** attribute.

Important

For UnboundID Directory Server and Alcatel-Lucent 8661 Directory Server systems, you must configure the **changelog-deleted-entry-include-attribute** property on the changelog backend. This property allows for the proper synchronization of DELETE operations that occur with this endpoint server. For more information, see “Configuring the Directory Server Backend for Synchronizing Deletes” on page 91.

Configuring the Synchronization Server Using the Command-Line Wizard

The general process to configure a Synchronization Server is to first define the external servers in the topology and then define the Sync Pipe(s) and its associated Sync Classes. You can use the **create-sync-pipe-config** tool to set up your Sync Pipes and Sync Classes. For bidirectional deployments, you will need to configure two Sync Pipes, one for each directional path (Sun DS 5.2 to UnboundID DS and vice-versa).

In the following example, we will also set up a simple attribute map that maps an **email** attribute on the first endpoint servers to the **mail** attribute on the second endpoint servers. In typical cases like these, you need to set up a specific attribute mapping from **email** to **mail** for the Sun to UnboundID Sync Pipe and also have both of these source attributes be excluded for automatic mapping. If you do not exclude the source attribute, the Synchronization Server will attempt to create an **email** attribute on the second endpoint topology, which could fail if the **email** attribute is not present in the destination server’s schema. Conversely, you have to create a specific **mail** to **email** mapping and auto-exclude the source attribute on the UnboundID to Sun Sync Pipe going the other direction.

For this example, you will define two sync classes: one to handle the customized **email** to **mail** attribute mapping; the other, to handle all other cases (called the default sync class). Next, you will use the **dsconfig** command to create the specific attribute mapping. After that, you can run the **resync** command to test the mappings. Finally, you can start synchronization using the **realtime-sync** command.

To Configure the Synchronization Server Using the Command-Line Wizard

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

2. Start the Synchronization Server.

```
$ bin/start-sync-server
```

3. From the **bin** directory, use the **create-sync-pipe-config** tool to set up the Synchronization sync pipes. The tool will start the command-line wizard and walk you through the steps to configure your Sync Pipes.

```
$ bin/create-sync-pipe-config
```

4. On the Initial Synchronization Configuration Tool menu, press Enter or Return (yes) to continue the configuration.
5. On the Synchronization Directory menu, select if the Synchronization topology will be one-way (1) or bidirectional (2). In this example, enter 2 for bidirectional.
6. On the First Endpoint Type menu, select the type of Directory Server for the first endpoint. The available options are seen below. In this example, type 5 to specify the Sun Directory Server assuming that we are using the Sun Directory Server 5.x.

```
>>>> First Endpoint Type
```

```
Enter the type of data store for the source endpoint:
```

- 1) UnboundID Directory Server
- 2) UnboundID Proxy Server
- 3) Alcatel-Lucent Directory Server
- 4) Alcatel-Lucent Proxy Server
- 5) Sun Directory Server
- 6) Microsoft Active Directory
- 7) Microsoft SQL Server
- 8) Oracle Database
- 9) Generic JDBC

- b) back
- q) quit

```
Enter choice [1]: 5
```

7. On the First Endpoint Name menu, type a name for the Endpoint Server, or accept the default ("Sun Directory Server"). In this example, type "Sun DS 5.2".
8. On the Base DN's menu, type the base DN on the first endpoint topology where the entries will be searched. In this example, accept the default (**dc=example,dc=com**).

If you have other base DN's, type the DN or press Enter when finished. If you enter another base DN, make sure that it does not overlap with the other base DN(s).
9. On the Server Security menu, select the server security type. The available options are 1 for None (LDAP), 2 for SSL, and 3 for StartTLS. In this example, press Enter to accept the default (**None**).
10. On the First Endpoint Servers menu, type the host name and listener port number for the First Endpoint Server, or accept the default (port 389). Make sure that the endpoint servers are online and running. The server will perform a test connection to the server. If the server is unresponsive, you will be asked to retry contacting the server, discard the server, or keep the server.
11. After entering the first server, enter the hostname and listener ports of the additional servers in the endpoint topology. The server will also perform a test connection to this server. If the server is unresponsive, you will be asked to retry contacting the server, discard the server, or keep the server.

At this stage, you can enter more servers, remove the existing servers, or press Enter when you are finished entering the servers.

12. Next, you will be prompted to enter the Sync User account DN for the endpoint servers. This step will ask you to enter a Sync User Account DN (**cn=Sync User**) and password.
13. At this point, you must set up the servers in the Second Endpoint topology. The steps are similar to steps 6–12. Select the option for UnboundID, and then set up the two external endpoint servers and Sync User Account DN.

Preparing the External Servers

1. After you have configured the first and second endpoint topologies, the Synchronization Server will prompt you to "prepare" each external server by testing the connection to each server. This step entails determining if each external server has the necessary privileges to communicate and to transfer data during synchronization. If an error occurs, the Synchronization Server will prompt you to re-configure the specific connection parameter.

Using the Sync User Account DN, the server verifies the base DNs, and enables and checks the changelog on the external server. If the maximum age of the changelog has not been set, you will also be prompted for a value.

2. Repeat step 1 to prepare the other external servers.

Note	
------	--

If your endpoint servers have no base entries or data, the command cannot create the **cn=Sync User**, **cn=Root DNs**, **cn=config** account. In this specific case, you can select 2 (Abandon the Operation) to continue, then create the base entry on the destination servers.

Configuring the Sync Pipes and its Sync Classes

1. Continuing the **create-sync-pipe-config** session, you will be prompted to create a name for the Sync Pipe on the Sync Pipe Name menu. Type a descriptive name to identify the Sync Pipe or accept the default. Because this example is bidirectional, the following step is setting up a Sync Pipe path from the Sun DS 5.2 endpoint to the UnboundID Directory Server endpoint. In a later step, you will need to define another Sync Pipe from UnboundID DS to Sun DS.
2. On the Sync Class Definitions menu, type **yes** if you want to create a custom Sync Class or press Enter to accept the default (no). A Sync Class defines the operation types (e.g., creates, modifies, or deletes) and attributes that are synchronized, how attributes and DNs are mapped, and how source and destination entries are correlated. In this example, create a basic sync class for the **email** to **mail** attribute mapping, which will exclude the source attribute from automatic synchronization. Later in the procedure, you will need to configure the **email** to **mail** attribute mapping using the **dsconfig** tool.
3. Next, you will be prompted to create a Sync Class name. Enter a name for the new Sync Class. For this example, enter "SunDS>UBID".
4. On the Base DNs for Sync Class menu, enter one or more base DNs if you want to synchronize specific subtrees of a DIT. Entries outside of the specified base DNs will be excluded

from synchronization. Make sure the base DN's do not overlap in any way. In this example, press Enter to accept the default (**no**) as we will not restrict any entries during the synchronization process.

5. On the Filters for Sync Class menu, you can define one or more LDAP search filters to restrict specific entries for synchronization. Those entries that do not match the filters will be excluded from synchronization. In this example, press Enter to accept the default (**no**).
6. Next, on the Synchronized Attributes for Sync Class menu, specify which attributes will be automatically mapped from one system to another. You can select the following options: 1 to Synchronize all attributes, 2 to Specify attributes to synchronize, 3 to Specify attributes to exclude from synchronization. In this example, assume that the Sun Directory Server endpoint has an **email** attribute that needs to be mapped to a **mail** attribute in the target endpoint servers. A specific attribute mapping will be configured in a later step. In this example, we will exclude the source attribute (**email**) from being auto-mapped to the target servers by selecting option 3 (Specify attributes to exclude from synchronization).
7. On the Operations for Sync Class menu, select the operations that will be synchronized for the Sync Class (1 for creates, 2 for deletes, 3 for modifies, 4 for none), or press Enter to accept the default ("1, 2, 3"). You can enter a comma-separated list of numbers that correspond to the operation. For these example, press Enter to accept the default (creates, deletes, modifies).
8. Next, define a default or "catch-all" Sync Class that specifies how the other entries are processed. In the following example, press Enter to continue, the system will create a Sync Class called "Default Sync Class".
9. On the Default Sync Class Operations menu, specify the operations that the default Sync Class (1 for creates, 2 for deletes, 3 for modifies, 4 for none) will handle during synchronization. In this example, press Enter to accept the default (1, 2, 3). You have successfully defined one sync pipe that goes from Sun Directory Server to UnboundID Directory Server.
10. At this stage, you must define a Sync Pipe going from the UnboundID Directory Server to the Sun Directory Server. Repeat the previous steps 4–9. When you create a sync class, make sure to create a **UBID>SunDS** sync class, and then exclude the **mail** attribute from being synchronized to the other endpoint servers.

Review the Configuration and Apply the Changes

1. Review the Sync Pipe Configuration Summary, and then, press Enter to accept the default ("write configuration"), which records the commands in a batch file (**sync-pipe-cfg.txt**). The batch file can be re-used to set up other Sync topologies.

```
>>>> Configuration Summary
```

```
Sync Pipe: Sun DS 5.2 to UnboundID DS
Source: Sun DS 5.2
Type: Sun Directory Server
Access Account: cn=Sync User,cn=Root DNs,cn=config
Base DN: dc=example,dc=com
Servers: sun-ds1.example.com:21389, sun-ds2.example.com:22389

Destination: UnboundID DS
```

```

Type: UnboundID Directory Server
Access Account: cn=Sync User,cn=Root DNs,cn=config
Base DN: dc=example,dc=com
Servers: UnboundID.example.com:23389, UnboundID.example.com:24389

Sync Classes:
SunDS>UBID
Base DN:
Filters:
DN Map: None
Synchronized Attributes: all except: email
Operations: Creates,Deletes,Modifies
DEFAULT
Operations: Creates,Deletes,Modifies

Sync Pipe: UnboundID DS to Sun DS 5.2
Source: UnboundID DS
Type: UnboundID Directory Server
Access Account: cn=Sync User,cn=Root DNs,cn=config
Base DN: dc=example,dc=com
Servers: UnboundID.example.com:23389, UnboundID.example.com:24389

Destination: Sun DS 5.2
Type: Sun Directory Server
Access Account: cn=Sync User,cn=Root DNs,cn=config
Base DN: dc=example,dc=com
Servers: sun-ds1.example.com:21389, sun-ds2.example.com:22389

Sync Classes:
UBID>SunDS
Base DN:
Filters:
DN Map: None
Synchronized Attributes: all except: mail
Operations: Creates,Deletes,Modifies
DEFAULT
Operations: Creates,Deletes,Modifies

w) write configuration
b) back
q) quit

Enter choice [w]:

```

2. Apply the configuration changes to the local Synchronization server instance. Once you have applied the changes to the server, you can review the configuration in the `<server-root>/sync-pipe-cfg.txt` file.
3. Connect to the Synchronization Server using the LDAP Connection Parameters: `host name`, `port`, `user bind DN` and `bind DN` password. The configuration is recorded to the `<server-root>/sync-pipe-cfg.txt`.

You have successfully configured the initial Sync Pipes for your system. The next step will be to configure the attribute mappings using the `dsconfig` command.

Configuring the Attribute Map and Mapping

The following section continues from the previous example by defining an attribute map that has a mapping from the `email` attribute in the source servers to a `mail` attribute in the target servers. You must ensure that both attributes are valid in the target servers and are present in their respective schemas.

1. On the Synchronization Server, run the `dsconfig` command to create an attribute map for the "SunDS>UBID" sync class for the "Sun DS 5.2 to UnboundID DS" sync pipe, and then run the second `dsconfig` command to apply the new attribute map to the Sync Pipe and Sync Class.

```
$ bin/dsconfig create-attribute-map \  
--map-name "SunDS>UBID Attr Map" \  
--set "description:Attribute Map for SunDS>UBID Sync Class" \  
--port 7389 --bindDN "cn=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt  
  
$ bin/dsconfig set-sync-class-prop \  
--pipe-name "Sun DS 5.2 to UnboundID DS" \  
--class-name "SunDS>UBID" \  
--set "attribute-map:SunDS>UBID Attr Map" \  
--port 7389 --bindDN "cn=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

Note

You can use `dsconfig` in interactive mode. The attribute map and attribute mapping options appear in the Synchronization Server Configuration Console main menu.

2. Next, create an attribute mapping (from `email` to `mail`) for the new attribute map.

```
$ bin/dsconfig create-attribute-mapping \  
--map-name "SunDS>UBID Attr Map" --mapping-name mail --type direct \  
--set "description:Email>Mail Mapping" --set from-attribute:email \  
--port 7389 --bindDN "cn=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

3. Because this example shows how to set up a bidirectional deployment, repeat steps 1–2 to create an attribute map for the `UBID>SunDS` sync class for the UnboundID DS to Sun DS 5.2 sync pipe, and create an attribute mapping that maps `mail` to `email`.

```
$ bin/dsconfig create-attribute-map --map-name "UBID>SunDS Attr Map" \  
--set "description:Attribute Map for UBID>SunDS Sync Class" \  
--port 7389 --bindDN "cn=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt  
  
$ bin/dsconfig set-sync-class-prop \  
--pipe-name "UnboundID DS to Sun DS 5.2" --class-name "UBID>SunDS" \  
--set "attribute-map:UBID>SunDS Attr Map" \  
--port 7389 --bindDN "cn=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt  
  
$ bin/dsconfig create-attribute-mapping \  
--map-name "UBID>SunDS Attr Map" --mapping-name email --type direct \  
--set "description:Mail>Email Mapping" --set from-attribute:mail \  
--port 7389 --bindDN "cn=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

Completing the Bidirectional Deployment

At this stage, you have configured the Sync Pipes, Sync Classes, and Attribute Mappings necessary for your synchronization topology.

To Complete the Bidirectional Deployment

1. Next, run the bulk synchronization command `resync` to test the attribute mapping. For more information, see “Using Bulk Resync on the Synchronization Server” on page 79.

```
$ bin/dsconfig --pipe-name "Sun DS 5.2 to UnboundID DS" \  
--sourceSearchFilter "(uid=user.0)" --dry-run \  
--logFilePath logs/resync.log --logLevel debug
```

2. Finally, start the synchronization process using the **realtime-sync** command. For more information, see “Controlling Real Time Synchronization” on page 84.

```
$ bin/realtime-sync start --pipe-name "Sun DS 5.2 to UnboundID DS" \  
--pipe-name "UnboundID DS to Sun DS 5.2" \  
--port 389 --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret
```

You have successfully completed your bidirectional Synchronization deployment.

Configuring the Synchronization Server Using the Management Console

The UnboundID Synchronization Server provides a graphical web application tool, UnboundID Sync Management Console, which accesses the server's underlying configuration. The Sync Management Console provides functionally equivalent to the **dsconfig** command-line tool in addition to monitoring and server information.

Note

Like the **dsconfig** tool, all changes made using the Sync Management Console are recorded in **logs/config-audit.log**.

Configuring the External Servers

External servers are the specific servers that should be included in the Synchronization topology. You must specify the LDAP connection and security parameters necessary to send requests to these servers. External servers can be either UnboundID Directory Servers, UnboundID Directory Proxy Servers (3.x), Alcatel-Lucent 8661 Directory Servers, Alcatel-Lucent 8661 Directory Proxy Servers (3.x), Sun Directory Server 5.x, Sun Directory Server Enterprise Edition (DSEE 6.x, 7.x), Microsoft Active Directory, Oracle (10g, 11g), or Microsoft SQL Server (2005, 2008).

To Define the External Servers

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

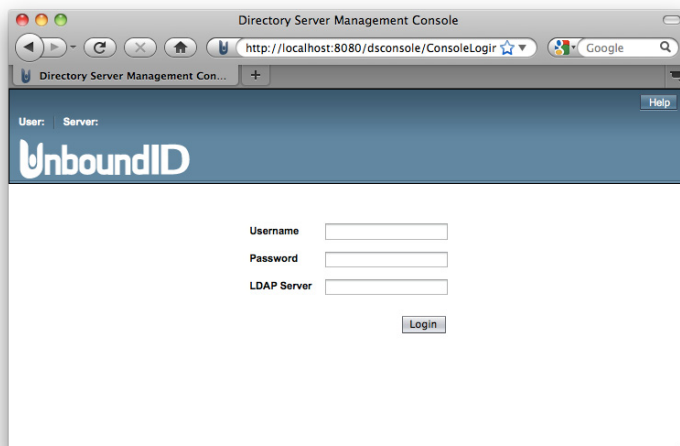
2. Start the Synchronization Server.

```
$ bin/start-sync-server
```

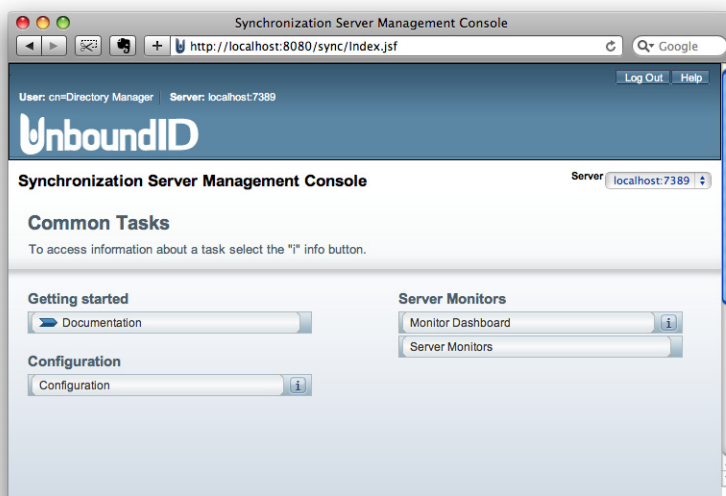
3. Start the Apache Tomcat, or equivalent, application server.

```
$ /apache-tomcat-<version>/bin/startup.sh
```

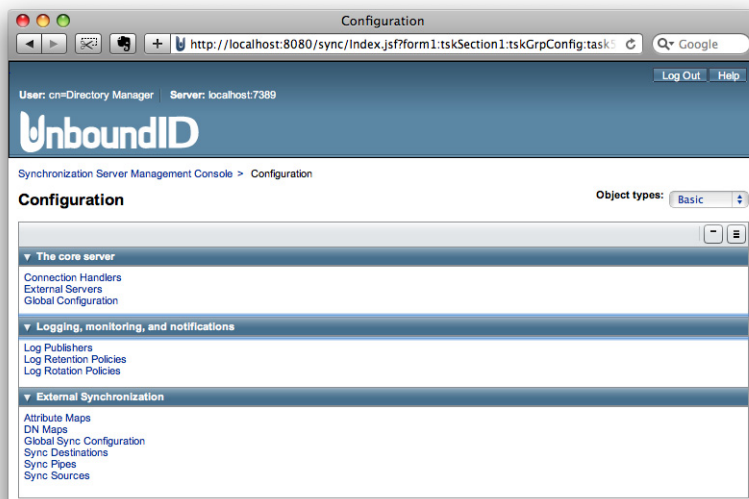
4. Open a browser to **http://hostname:8080/sync**. The application server listens on port 8080 for HTTP requests.
5. Type the root user DN (or any authorized administrator user name) and password, and the server hostname or IP address and port to log on (for example, server1.example.com:389).



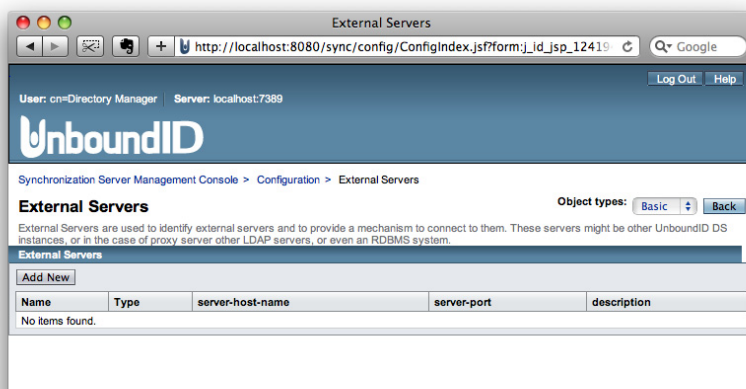
6. On the Synchronization Server Management Console, click **Configuration**.



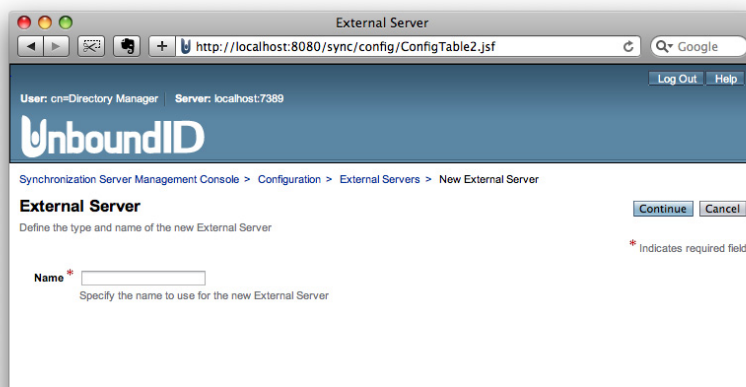
7. Under "The core server," click **External Servers** to identify all of the servers that will be synchronized.



8. Click the **Add New** button to define the first server in the topology.



9. Type a name for the external server, and then click **Continue**. The name can be any label that will help you identify the server.



10. On the Type drop-down menu, select the type of external server that you are defining. In this example, select Sun DS Sync Source.

External Server

User: cn=Directory Manager Server: localhost:7389

UnboundID

Synchronization Server Management Console > Configuration > External Servers > New External Server

External Server

Define the type and name of the new External Server

Continue Cancel

* Indicates required field

Name *
Specify the name to use for the new External Server

Type *
Specify the type to use for the new External Server

11. Type the hostname for the external server, and then click **Continue**.

External Server

User: cn=Directory Manager Server: localhost:7389

UnboundID

Synchronization Server Management Console > Configuration > External Servers > New External Server

External Server

Define the type and name of the new External Server

Continue Cancel

* Indicates required field

Name *
Specify the name to use for the new External Server

Type *
Specify the type to use for the new External Server

Server Host Name
The host name or IP address of the target LDAP server.

12. Type the external server's connection parameters that were configured when the server was first installed. If security and authentication settings were configured for the external server, define them on this page. When completed, click **Confirm then Save**.

The screenshot shows the 'sun-ds-1' configuration page in the UnboundID Synchronization Server Management Console. The page is titled 'sun-ds-1' and includes a breadcrumb trail: 'Synchronization Server Management Console > Configuration > External Servers > sun-ds-1'. The user is 'cn=Directory Manager' and the server is 'localhost:7389'. The page contains several configuration fields:

- Description:** A text area for a description of the external server.
- Server Host Name:** A text field containing 'sun-ds1.example.com'.
- Server Port:** A text field containing '21389'.
- Bind DN:** A text field containing 'cn=Directory Manager'.
- Password:** A button labeled 'Set Password'.
- Connection Security:** A dropdown menu set to 'none'.
- Authentication Method:** A dropdown menu set to 'simple'.
- Allowed Operation:** Two lists, 'Available' and 'Selected', with buttons to move items between them. The 'Selected' list contains: abandon, add, bind, compare, delete, extended, modify, modify-dn, search.
- Key Manager Provider:** A dropdown menu set to 'None'.
- Trust Manager Provider:** A dropdown menu set to 'None'.

At the bottom of the page are buttons: 'Show Advanced', 'Confirm then Save', 'Save Now', and 'Cancel'.

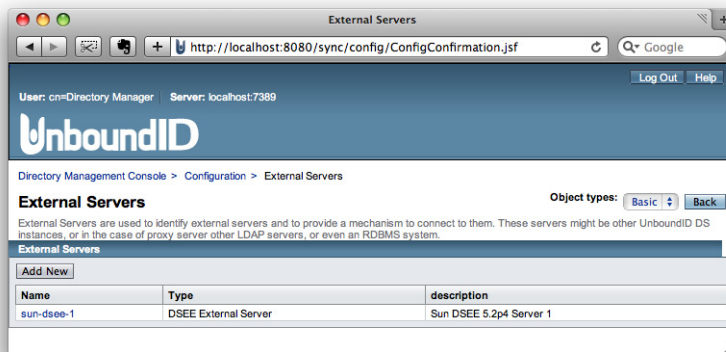
13. Click **Apply** to save the settings for this external server. The equivalent **dsconfig** command-line instruction is displayed to recreate the external server in a scripted installation or to quickly define similar external servers from the command line.

The screenshot shows the 'Confirm changes to sun-ds-1' page in the UnboundID Synchronization Server Management Console. The page is titled 'Confirm changes to sun-ds-1' and includes a breadcrumb trail: 'Synchronization Server Management Console > Configuration > External Servers > sun-ds-1 > Confirm changes to sun-ds-1'. The user is 'cn=Directory Manager' and the server is 'localhost:7389'. The page contains a yellow button labeled 'Click Apply to create 'sun-ds-1''. Below this is a section titled 'Equivalent dsconfig Command' with the following command:

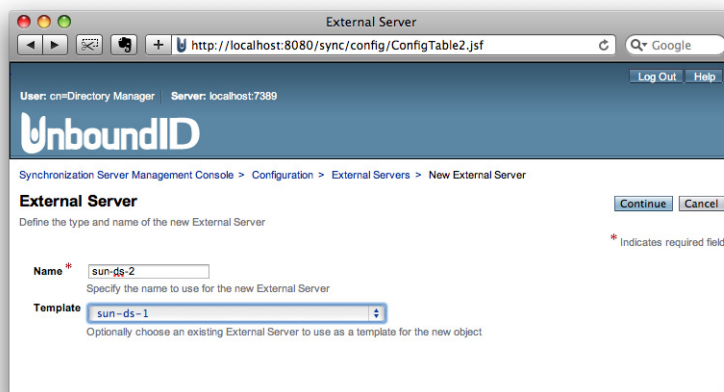
```
dsconfig create-external-server --server-name sun-ds-1 --type dsee --set server-host-name:sun-ds1.example.com --set server-port:21389 --set "bind-dn:cn=Directory Manager" --set "password:AAA/ko5XLMsP4PcGSxJ6fdxL5VeaJts8IA="
```

At the bottom of the page are buttons: 'Apply' and 'Back'.

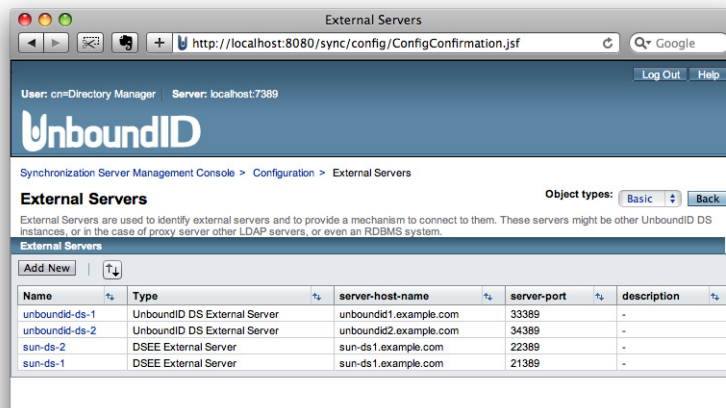
14. Click **Add New** to define another external server.



15. Once you have defined one external server, you can use the settings defined for the first server as a template for the next one. On the External Server page, type the name of the new external server, select the first server on the Template drop-down menu, and then click **Continue**. In this example, only the server name and host name has changed for the second server.



16. At this stage, repeat steps 8–13 to define the other external source and target servers. You will only need to change the description, host name, and port number for each server.



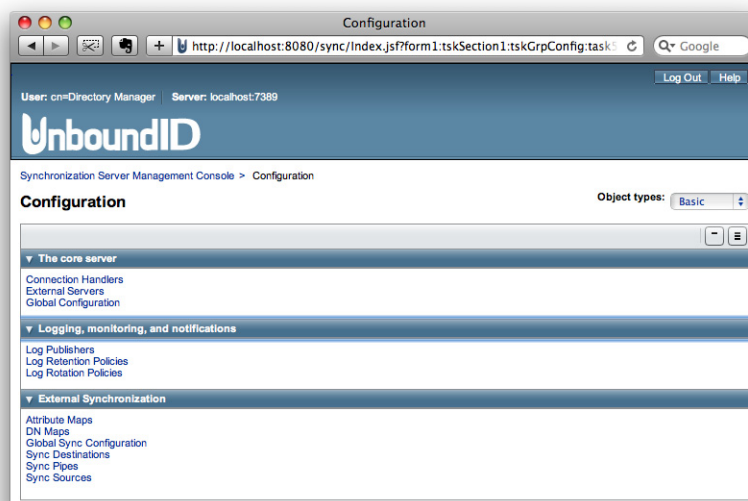
You have successfully defined the external servers in the Synchronization topology.

Configuring the Sync Pipe

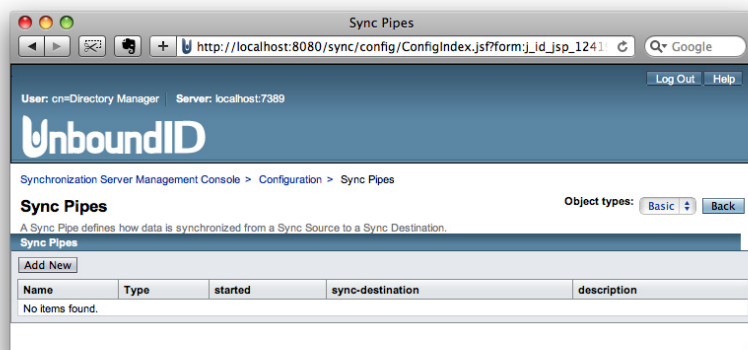
Next, you will need to configure how synchronization is processed between the Source and Destination topologies. The next two sections present information on how to set up the Sync Pipe and Sync Class.

To Configure the Sync Pipe Using the Management Console

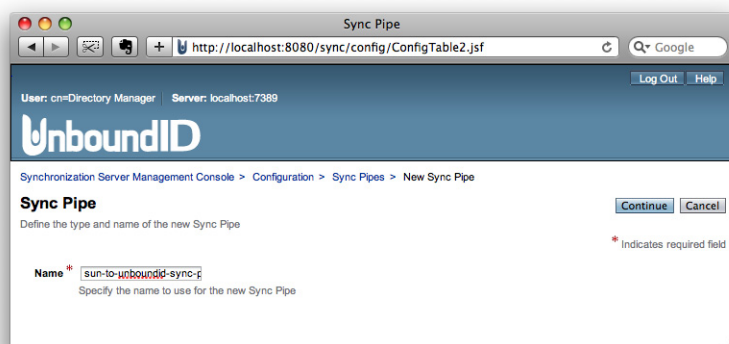
1. On the Configuration page, click **Sync Pipes**.



2. Click **Add New** to define a Sync Pipe.

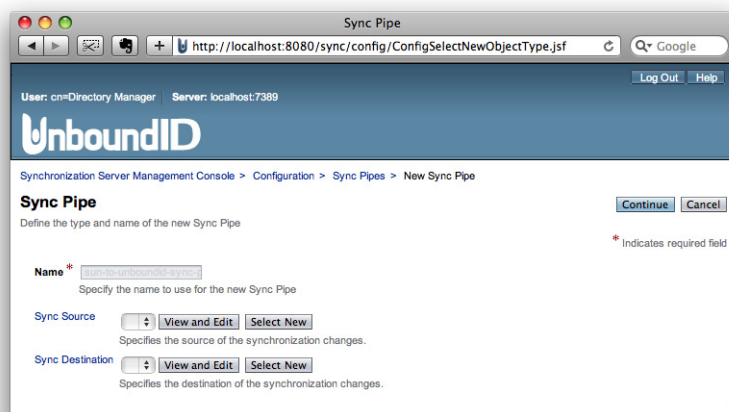


3. Type the name of the Sync Pipe, and then click **Continue**.



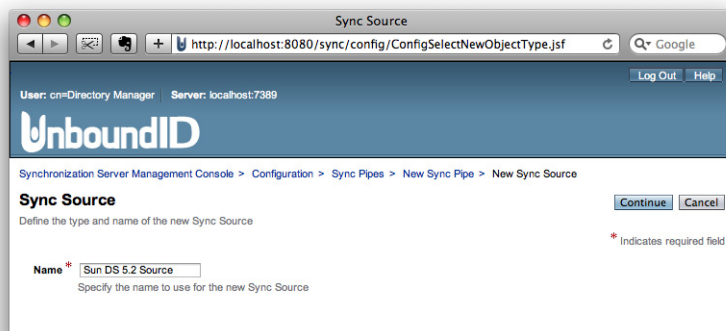
The screenshot shows the 'Sync Pipe' configuration page in the UnboundID Synchronization Server Management Console. The browser address bar shows 'http://localhost:8080/sync/config/ConfigTable2.jsf'. The page header includes 'User: cn=Directory Manager' and 'Server: localhost:7389'. The breadcrumb trail is 'Synchronization Server Management Console > Configuration > Sync Pipes > New Sync Pipe'. The main heading is 'Sync Pipe' with a subtext 'Define the type and name of the new Sync Pipe'. There are 'Continue' and 'Cancel' buttons. A 'Name' field is present with the text 'sun-to-unboundid-sync-p' and a red asterisk indicating it is a required field. A note below the field says 'Specify the name to use for the new Sync Pipe'. A legend indicates '* Indicates required field'.

4. For the Sync Source, click **Select New**.



The screenshot shows the 'Sync Pipe' configuration page, similar to the previous one, but with additional options. The 'Sync Source' and 'Sync Destination' fields are visible, each with a dropdown menu and a 'Select New' button. The 'Sync Source' field has a note 'Specifies the source of the synchronization changes.' and the 'Sync Destination' field has a note 'Specifies the destination of the synchronization changes.' The 'Continue' and 'Cancel' buttons are still present. The 'Name' field remains the same. The legend '* Indicates required field' is also present.

5. Type a name for the Sync Source to identify the topology, and then click **Continue**.



The screenshot shows the 'Sync Source' configuration page in the UnboundID Synchronization Server Management Console. The browser address bar shows 'http://localhost:8080/sync/config/ConfigSelectNewObjectType.jsf'. The page header includes 'User: cn=Directory Manager' and 'Server: localhost:7389'. The breadcrumb trail is 'Synchronization Server Management Console > Configuration > Sync Pipes > New Sync Pipe > New Sync Source'. The main heading is 'Sync Source' with a subtext 'Define the type and name of the new Sync Source'. There are 'Continue' and 'Cancel' buttons. A 'Name' field is present with the text 'Sun DS 5.2 Source' and a red asterisk indicating it is a required field. A note below the field says 'Specify the name to use for the new Sync Source'. A legend indicates '* Indicates required field'.

- On the Type drop-down menu, select the Sync Source type. In this example, select "Sun DS Sync Source."

- In the **New** field, type the base DN to be used for synchronization searches, and then click **Add**. The base DN defines the scope of the searches that the Synchronization Server processes for its change flow. You can specify more than one base DN, but the base DN's must not overlap another base DN (i.e., they cannot be sub-branches of another base DN).
- In the **Server** section, select the server(s) to be used as the Sync Source. The order of the servers is important as it determines the priority order of the Synchronization source. Specifically, the Synchronization Server will connect to the first server when detecting changes as long as it is available.

Click **Add All** if the default order is acceptable. For example, in the graphic below, the **sun-ds-1** server is used in preference to the **sun-ds-2** server. If you want to select the second server as the higher priority, click the server link, and then click **Add**. Then, move the other server to the Select column. Click **Continue** when done.

- In the **Ignore Changes by DN** field, type a DN for which the Synchronization Server should ignore any modifications by the user DN, and then click **Add**. This function serves as a form of loopback detection from the Destination target to the Source target when using bidirectional synchronization. During loop back, the Synchronization Server ignores any modifications made by the user, except for any deletion operations.

Click **Confirm then Save** when done. For example, you can specify the DN of the synchronization user, `uid=user.0,ou=People,dc=example,dc=com`.

Sun DS 5.2 Source

User: cn=Directory Manager Server: localhost:7389

Synchronization Server Management Console > Configuration > Sync Pipes > New Sync Pipe > Sun DS 5.2 Source

Sun DS 5.2 Source [Confirm then Save](#) [Save Now](#) [Cancel](#)

An DSEE Sync Source defines the source of a Sync Pipe that is topology of Sun Directory Server instances.

Description:

A description for this Sync Source

Base DN

New: [Add](#)

Current: [Remove/Edit](#)

Current:

Specifies the base DN of the directory servers referenced by this Sync Source. These base DN's are used as the base of LDAP searches when locating entries. These base DN's must not overlap.

Ignore Changes By DN

New: [Add](#)

Current: [Remove/Edit](#)

Current:

Modifications performed by users with the specified DN will not be synchronized.

Server

Available: [Add >](#) [Add All >>](#) [< Remove](#) [<< Remove All](#)

Selected: sun-ds-1
sun-ds-2

[View and Edit](#) [Select New](#)

Specifies the names of the Sun Directory Servers that should be used as the source of synchronization. The order of values is important as it is used as a priority order -- if a server earlier in the list is available, then it will be used in preference to a server later in the list.

[Confirm then Save](#) [Save Now](#) [Cancel](#)

- Click **Apply** to save the Sync Source configuration.

Confirm changes to Sun DS 5.2 Source

User: cn=Directory Manager Server: localhost:7389

Click Apply to create 'Sun DS 5.2 Source'

Synchronization Server Management Console > Configuration > Sync Pipes > New Sync Pipe > Sun DS 5.2 Source > Confirm changes to Sun DS 5.2 Source

[Apply](#) [Back](#)

Equivalent dsconfig Command

```
dsconfig create-sync-source --source-name "Sun DS 5.2 Source" --type dsee --set "base-dn:dc=example,dc=com" --set server:sun-ds-1 --set server:sun-ds-2
```

11. Repeat steps 5–10 for the Sync Destination, so that you have the Sync Source and Sync Destination defined for the Sync Pipe.

Sync Pipe

User: cn=Directory Manager Server: localhost:7389

Synchronization Server Management Console > Configuration > Sync Pipes > New Sync Pipe

Sync Pipe

Define the type and name of the new Sync Pipe

Name * * Indicates required field

Specify the name to use for the new Sync Pipe

Sync Source

Specifies the source of the synchronization changes.

Sync Destination

Specifies the destination of the synchronization changes.

12. After defining the Sync Destination, enter a description for the Sync Pipe.

Modify the Polling Interval if necessary. The Polling Interval is the amount of time that the Synchronization Server waits between checking the Sync Source for changes. The default time is 500 ms.

Although likely unnecessary, you can change the default number of worker threads if necessary. The number of worker threads should be increased if there is a large network latency between the Sync Source servers and the Sync Destination servers. Click **Confirm then Save** when done.

sun-to-unboundid-sync-pipe

User: cn=Directory Manager Server: localhost:7389

Synchronization Server Management Console > Configuration > Sync Pipes > sun-to-unboundid-sync-pipe

sun-to-unboundid-sync-pipe

A Sync Pipe defines how data is synchronized from a Sync Source to a Sync Destination.

Description

A description for this Sync Pipe

Started ☒

Indicates whether the synchronization for the Sync Pipe is started.

Sync Source

Specifies the source of the synchronization changes.

Sync Destination

Specifies the destination of the synchronization changes.

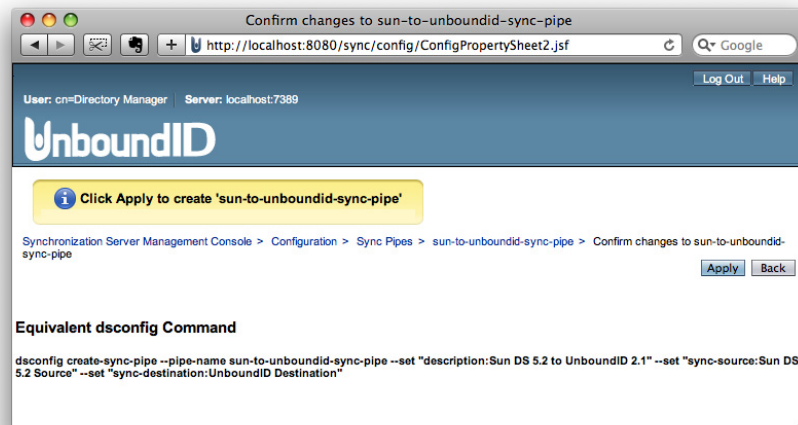
Change Detection Polling Interval

The initial amount of time to wait between polling the Sync Source for changes.

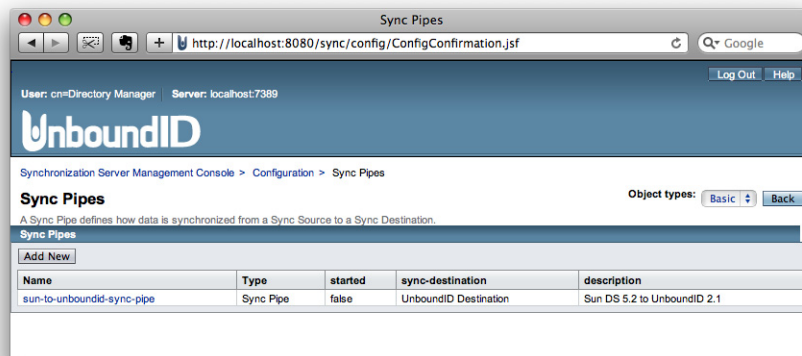
Num Worker Threads

Specifies the number of worker threads that should be used to process synchronization operations.

13. Click **Apply** to save the changes.



14. Repeat steps 1–13 if you want to define a bidirectional Sync pipe from the Sync Destination to Sync Source. Otherwise, click **Back** to define the next Synchronization configuration.



You should now be on the specific Sync Pipe configuration page. At this stage, you are almost done. You must configure the Sync Class, and then enable the Sync Pipe, which is covered in the next section.

Configuring the Sync Class

A Sync Class is defined for each type—or *class*—of entry that should be treated differently by the Synchronization server. This includes what types of changes are synchronized, what attributes are synchronized and how they are mapped, how source and destination entries are correlated, and how DNs are mapped.

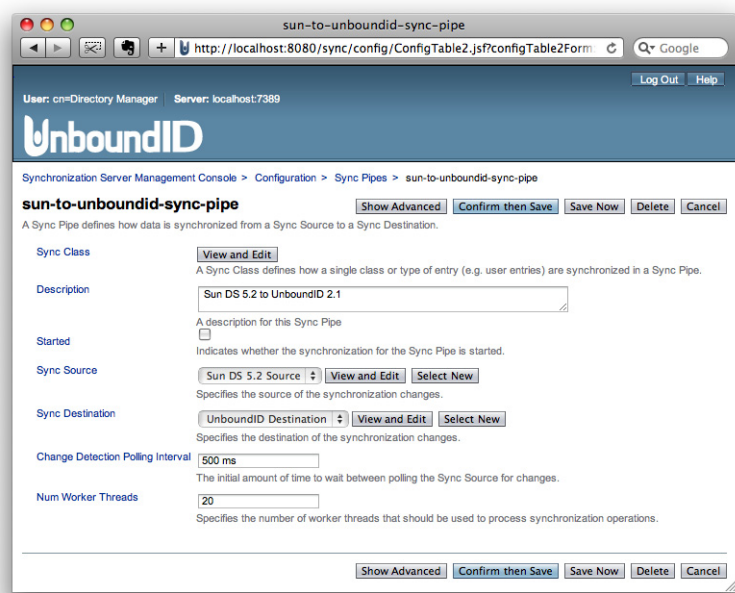
The Sync Class also defines what attributes within the entries should be included or excluded in the synchronization process. When a change to an entry is first detected in a Sync Source, the Sync Pipe evaluates the inclusion criteria (i.e., `include-base-dn` and `include-filter`) to find the first matching Sync Class according to the `evaluation-order-index` property. If a change does not match any Sync Class, then it is discarded. Otherwise, the matching Sync

Class processes any attribute or DN mappings and determines what type of change is synchronized.

Note	If you do not want certain types of entries to be synchronized, then you can define a Sync Class for these attributes and then clear the sync-creates , sync-mods , and sync-deletes boxes.
------	--

To Configure a Sync Class

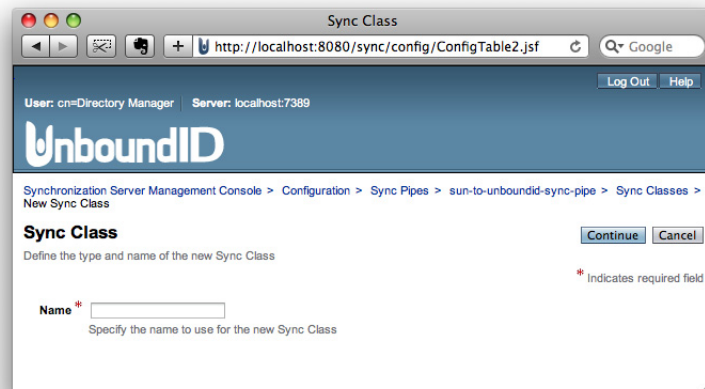
- 1. On the **Sync Pipe** page, click **View and Edit** next to the Sync Class that you want to configure.



- 2. Click **Add New** to define a Sync Class for the Sync Pipe. A Sync Pipe may have more than one Sync Class defined.



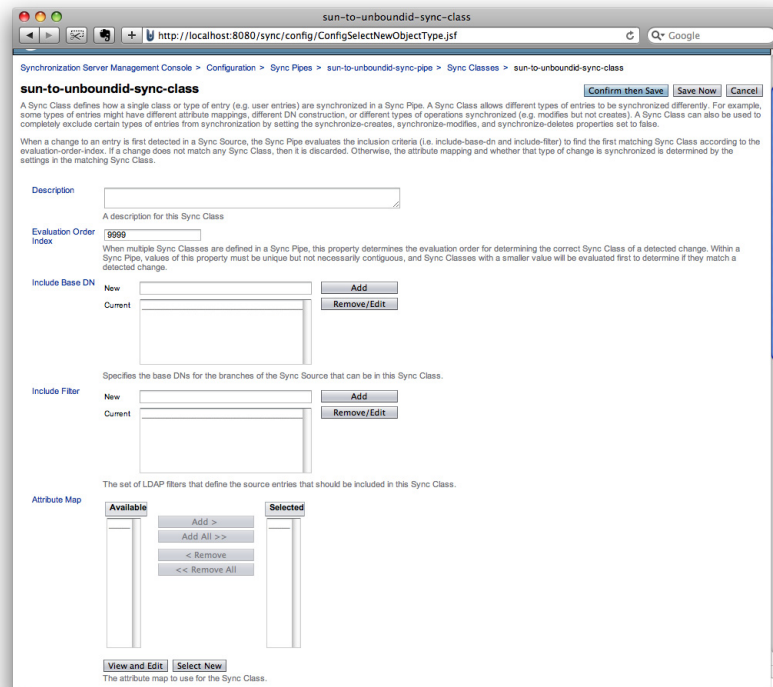
3. Type a name for the Sync Class.



4. On the **Sync Class** page, in the **Description** field, type a general description for the Sync Class. This is an optional step.
5. In the **Evaluation Order Index** field, type the priority ordering for the Sync Class if you have more than one Sync Class configured for the topology. Sync Classes with a smaller evaluation order index are evaluated first. Because this example defines only one Sync Class, the default value of 9999 is used.
6. In the **Include Base DN** field, type the base DN for the branches of the Sync Source that contain entries in this Sync Class. Only entries with this base DN will be included in the Sync Class. This is an optional step.

If no base DN is specified, the location of the entry in the Sync Source is not taken into account when determining if an entry is part of this Sync Class.
7. In the **Include Filter** field, type a search filter that determines which entries are in the Sync Class. If no filter is specified, all entries within the specified included base DN's are included in the Sync Class.
8. In the **Attribute Map** section, click Select New to define a set of attribute mappings from Sync Source to Sync Destination. In this example, the Sync Source (Sun DS 5.2) attributes

map directly to the Sync Destination (UnboundID Directory Server), so no attribute maps require definition. See “Configuring Attribute Maps” on page 88.



9. In the **DN Map** section, click **Select New** to define a set of DN mappings from Sync Source to Sync Destination. In this example, the Sync Source DNs (Sun DS 5.2) map directly to the Sync Destination DNs (UnboundID Directory Server), so no DN maps require definition. See “Configuring DN Maps” on page 91.
10. In the **Auto Mapped Source Attribute** field, type any source attributes that should be automatically mapped to attributes of the same name in the destination target, and then click **Add**. By default, all attributes are mapped automatically.
11. In the **Excluded Auto Mapped Attributes** field, type any source attributes that should not be automatically mapped to attributes in the destination target, and then click **Add**. By default, no attributes are excluded.
12. In the **Destination Correlation Attributes** field, type a comma-separated list of destination attributes that are used to correlate a source entry to a destination entry. For example, the default option is to use the **dn** to correlate entries, but you could specify that the **dn** and **uid** attributes be used to correlate entries, or the **cn** and **employeeNumber** attributes, or others, depending on how the entries are structured in the Sync Source and Sync Destination, respectively.
13. Clear the specific types of changes that you do not want to synchronize:
 - Synchronize Creates
 - Synchronize Modifies
 - Synchronize Deletes

14. When completed, click **Confirm then Save**.

Sync Class 1

http://localhost:8080/sync/config/ConfigSelectNewObjectType.jsf

DN Map

Available: Add >, Add All >>, << Remove, << Remove All

Selected:

View and Edit | Select New

If the DN values at the Sync Destination differ from the DN values at the Sync Source, a DN map can be used to construct the destination DN value using components of the source DN and attributes from the source entry.

Auto Mapped Source Attributes

New: Add
Current: -all- Remove/Edit

The set of source attribute names that should be automatically mapped to a destination attribute with the same name.

Excluded Auto Mapped Attributes

New: Add
Current: Remove/Edit

The set of source attribute names that should not be automatically mapped to a destination attribute with the same name.

Destination Correlation Attributes

New: Add
Current: dn Remove/Edit

Each value of this property is a comma separated list of destination attribute names that are used to correlate a source entry to a destination entry. When multiple attribute lists are provided, the server iterates through the attribute lists and uses the first one that matches a single entry.

Synchronize Creates ☒ This property controls whether or not the creation of a new entry at the Sync Source should be synchronized to the Sync Destination. If set to true, the creation will be synchronized, and otherwise the change will be dropped.

Synchronize Modifies ☒ This property controls whether or not the modification of an entry at the Sync Source should be synchronized to the Sync Destination. If set to true, the modify will be synchronized, and otherwise the change will be dropped.

Synchronize Deletes ☒ This property controls whether or not the deletion of an entry at the Sync Source should be synchronized to the Sync Destination. If set to true, the corresponding entry in the destination will be deleted, and otherwise the change will be dropped.

Confirm then Save | Save Now | Cancel

15. Click **Apply** to complete defining the Sync Class.

Confirm changes to sun-to-unboundid-sync-class

http://localhost:8080/sync/config/ConfigPropertySheet2.jsf

User: cn=Directory Manager Server: localhost:7389 Log Out Help

UnboundID

Click Apply to create 'sun-to-unboundid-sync-class'

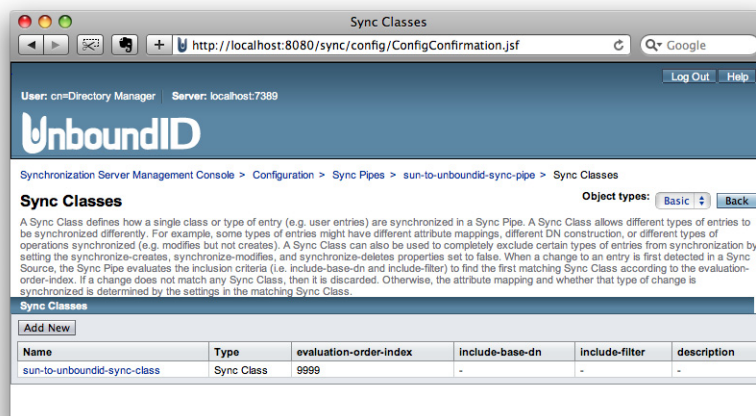
Synchronization Server Management Console > Configuration > Sync Pipes > sun-to-unboundid-sync-pipe > Sync Classes > sun-to-unboundid-sync-class > Confirm changes to sun-to-unboundid-sync-class

Apply Back

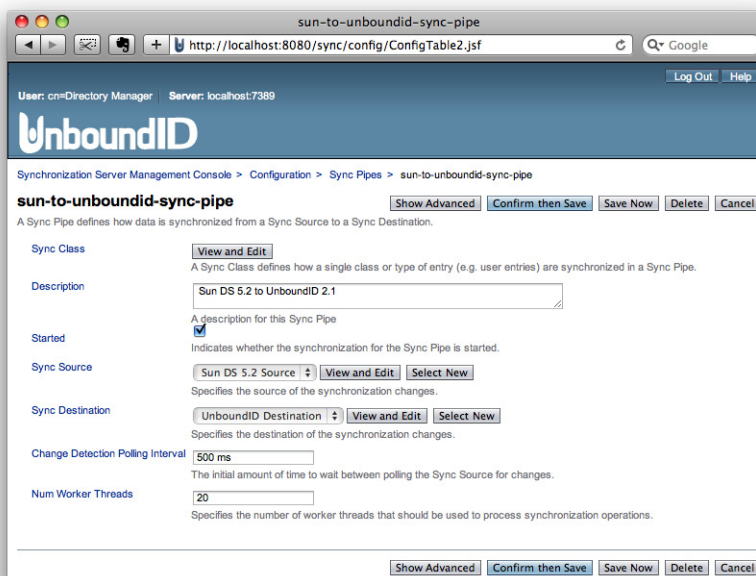
Equivalent dsconfig Command

dsconfig create-sync-class --pipe-name sun-to-unboundid-sync-pipe --class-name sun-to-unboundid-sync-class

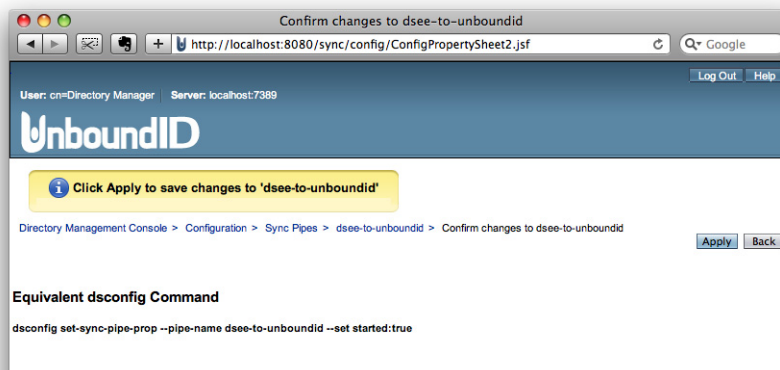
16. On the **Sync Classes** page, click **Back** to return to the Sync Pipe page.



17. On the **Sync Pipe** page, click **Started**, and then click **Confirm then Save**. The **Started** field on the Sync Pipe controls whether a given Sync Pipe is synchronizing.



18. Click **Apply** to save the settings.



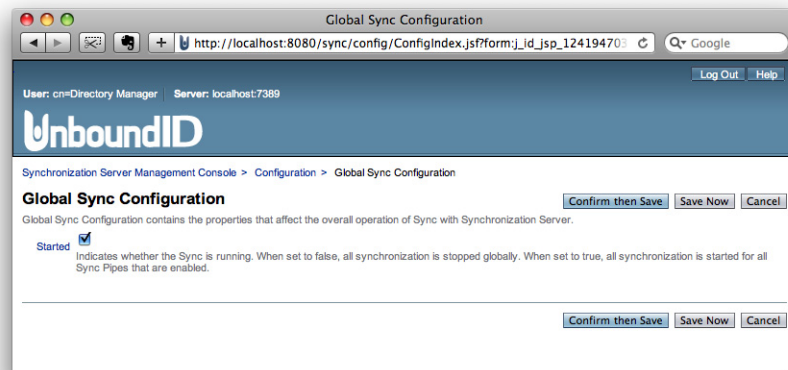
19. Repeat steps 2–18 to create another Sync Class, or log out of the console.

Starting the Global Sync Configuration


After you have configured the Sync Pipe and Sync Class, you must start the Global Sync configuration property (that is, enable synchronization). By starting the Synchronization Server, it starts or stops synchronization for all configured Sync Pipes. Each Sync Pipe must also be started for synchronization to take place.

To Start the Global Sync Configuration

1. On the Configuration page, click **Global Sync Configuration**.
2. Click the **Started** box, and then click **Confirm then Save**.



3. Click **Apply** to save the configuration settings.



4. After you have completed the configuration, run the **prepare-endpoint-server** tool from the command line to ensure that the external servers can communicate with each other. See “Preparing the Synchronization Server for External Server Communication” on page 75 for more information.
5. Next, run the **bin/resync** tool to verify the synchronization configuration. See “Verifying the Synchronization Configuration using Resync” on page 80.

6. Next, run the `bin/realtime-sync` tool to start the startpoint. See “Setting Startpoints” on page 86.

Using the dsconfig Configuration Tool

The `dsconfig` tool is the text-based management script used to configure the underlying Synchronization Server configuration. The tool has three operational modes: interactive mode, non-interactive mode, and batch mode. The tool also offers an enhanced user interface that displays the most common configuration options using object complexity levels. By selecting a complexity level, you see additional menu configuration properties that are normally hidden for the basic user. The tool also has toggling mechanisms to hide or show advanced properties for specific configuration objects.

Using dsconfig in Interactive Command-Line Mode

Running `dsconfig` in interactive command-line mode provides a user-friendly, menu-driven interface for accessing and configuring the UnboundID Synchronization Server. To start `dsconfig` in interactive command-line mode, simply invoke the `dsconfig` script without any arguments. You will be prompted for information about how to connect and authenticate to the Synchronization Server, and then a menu will be displayed of the available operation types.

```
$ bin/dsconfig
```

In some cases, a default value will be provided in square brackets. For example, `[389]` indicates that the default value for that field is port 389. You can press Enter or Return to accept the default. To skip the connection and authentication prompts, provide the information using the command-line options as follows:

```
$ bin/dsconfig --port 389 --bindDN "uid=admin,dc=example,dc=com" \
  --bindPassword secret
```

Many of the `dsconfig` property menus provide additional feature options, such as, the show advanced properties option and display non-interactive equivalent. The Show Advanced Properties option is available on some `dsconfig` interactive property menus and can be invoked by selecting "a" to show any advanced properties in the menu. The advanced properties are hidden to safeguard against potential harm to the server if not properly configured by an experienced user.

```
? )  help
f )  finish - create the new Sun DS External Server
a )  show advanced properties of the Sun DS External Server
d )  display the equivalent dsconfig arguments to create this object
b )  back
q )  quit
```

```
Enter choice [b]:
```

Typically, obtaining the correct non-interactive arguments and properties for each configuration change can be time-consuming. The Display Equivalent Arguments feature allows you to

see an equivalent non-interactive **dsconfig** command of the configuration object. You can see the equivalent command by selecting "d" (seen above).

Note

There are two other ways to get the equivalent **dsconfig** command. One way is to look at the `logs/config-audit.log`. The other method is to configure a property on the Sync Management Console. The console shows the equivalent **dsconfig** command prior to applying the change.

Using dsconfig in Non-Interactive Command-Line Mode

The **dsconfig** non-interactive command-line mode provides a simple way to make arbitrary changes to the Synchronization Server by invoking it on the command line. If you want to use administrative scripts to automate configuration changes, then running **dsconfig** command in non-interactive mode is more convenient. Note, however, that if you plan to make changes to multiple configuration objects at the same time, then the batch mode might be more appropriate.

You can use the **dsconfig** tool to update a single configuration object using command-line arguments to provide all of the necessary information. The general format for the non-interactive command line is:

```
$ bin/dsconfig --no-prompt {subcommand} {globalArgs} {subcommandArgs}
```

The `--no-prompt` argument indicates that you want to use non-interactive mode, which suppresses output information, except when fatal errors occurs. The `{subcommand}` is used to indicate which general action to perform. The `{globalArgs}` argument provides a set of arguments that specify how to connect and authenticate to the Synchronization Server, and the `{subcommandArgs}` argument contains a set of arguments specific to the particular subcommand that you wish to invoke.

Note

Global arguments can appear anywhere on the command line (including before the subcommand, and after or intermingled with subcommand-specific arguments). The subcommand-specific arguments can appear anywhere after the subcommand.

Using dsconfig in Batch Mode

The UnboundID Synchronization Server provides a **dsconfig** batching mechanism that reads multiple **dsconfig** invocations from a file and executes them sequentially. The batch file provides advantages over standard scripting in that it minimizes LDAP connections and JVM invocations that normally occur with each **dsconfig** call. Batch mode is the best method to use with setup scripts when moving from a development environment to test environment, or from a test environment to a production environment. For example, assuming that you created a Synchronization topology as outlined in “Configuring the Synchronization Server Using the Command-Line Wizard” on page 42, you can use the `sync-pipe-config.txt` to apply the same configuration on a failover instance.

```
$ bin/dsconfig --hostname host1 --port 1389 \  
--bindDN "uid=admin,dc=example,dc=com" --bindPassword secret \  
--batch-file /path/to/sync-pipe-config.txt --no-prompt
```

If a **dsconfig** command has a missing or incorrect argument, the command will fail and abort the batch process. Any command that was successfully executed prior to the abort will be applied to the Synchronization Server. The **--no-prompt** option is required with **dsconfig** in batch mode.

You can view the **logs/config-audit.log** file to review the configuration changes made to the Synchronization Server and use them in the batch file. The batch file can have blank lines for spacing and lines starting with a pound sign (#) for comments.

Configuring the Synchronization Server Using dsconfig

You can use the **dsconfig** tool to configure any part of the Synchronization Server. However, you will likely use the tool for more fine-grained adjustments. If you are configuring a Sync Pipe for the first time, you should use the **bin/create-sync-pipe-config** tool as it will guide you through the necessary Sync Pipes creation steps for your system.

To Configure the Synchronization Server Using dsconfig Interactive Command-Line Mode

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```
2. Launch the **dsconfig** tool in interactive command-line mode.

```
$ bin/dsconfig
```
3. On the LDAP Connection Parameters menu, type the Synchronization Server host name, or IP address, or press Enter or Return to accept the default.
4. On the Synchronization Server Connection menu, type the number corresponding to the type of LDAP connection type (1 for LDAP, 2 for SSL, 3 for StartTLS) that you are using on the Synchronization Server, or press Enter or Return to accept the default.
5. Next, type the LDAP listener port number, and then type the user bind DN, and the bind DN password.
6. On the Synchronization Server Configuration Console main menu, enter a number corresponding to a component that you want to configure or edit.

```
>>>> UnboundID Synchronization Server configuration console main menu
```

```
What do you want to configure?
```

- | | |
|------------------------------|-------------------------|
| 1) Attribute Map | 8) Log Publisher |
| 2) Attribute Mapping | 9) Log Retention Policy |
| 3) Connection Handler | 10) Log Rotation Policy |
| 4) DN Map | 11) Sync Class |
| 5) External Server | 12) Sync Destination |
| 6) Global Configuration | 13) Sync Pipe |
| 7) Global Sync Configuration | 14) Sync Source |


```
o) 'Basic' objects are shown - change this
q) quit

Enter choice:
```

Configuring External Servers Using dsconfig Interactive Mode

To set up a Synchronization topology, you must define a single server of a topology of identical, replicated servers to be synchronized. For each Directory Server, you must define the host, port, SSL, bind DN, and bind password. A single external server configuration object can be referenced by multiple Sync Sources and Sync Destinations.

To Configure the External Servers Using dsconfig Interactive Mode

1. On the Synchronization Server Configuration console main menu, type the number corresponding to External Server. In this example, type 5.
2. On the External Server Management menu, type 2 to create a new External Server.
3. Next, select the type of external server. In this example, type 5 for Sun DS External Server.
4. Next, you will be prompted to enter the name for the external server.
5. On the Server-Host-Name Property menu, type the host name of the external server.
6. On the Sun DS External Server Properties menu, change the **server-port**, **bind-dn**, and **password** for the external server. Type the number corresponding to each property, and follow the prompts to enter the values. When completed, type **z** to save and apply the changes.
7. Repeat steps 2–6 to define any additional external servers. The Synchronization Server uses the settings for the first server as a template to create the other external servers. Type 1 to use the first external server as a template for the other external server.
8. Repeat steps 2–7 to create the other external servers that you plan to synchronize.
9. On the External Server Management menu, type 1 to view the list of external servers that you have created.

External Server	Type	server-host-name	server-port
ds-dest1	UnboundID-ds	ds3.example.com	389
ds-dest2	UnboundID-ds	ds4.example.com	389
ds-src1	sun-ds	ds1.example.com	389
ds-src2	sun-ds	ds2.example.com	389

Configuring the Sync Source Using dsconfig Interactive Mode

Sync Sources define the directory topology that is the source of the data to be synchronized. When data in the Sync Source changes, it is synchronized to the Sync Destination topology. Sync Sources can reference one or more external servers of the appropriate type (UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory

Server, Alcatel-Lucent 8661 Directory Proxy Servers (3.x), Sun Directory Server 5.x, Sun DSEE 6.x, 7.x, Microsoft Active Directory, Oracle 10g, 11g, or Microsoft SQL Server 2005, 2008).

To Configure the Sync Source

1. On the Synchronization Server Configuration console main menu, type the number corresponding to the Sync Source. In this example, type **14**.
2. On the Sync Source Management menu, type **2** to create a new Sync Source.
3. On the Sync Source Type menu, enter the number corresponding to the Sync Source type. In this example, type **4** for Sun DS Sync Source.
4. Next, you will be prompted to enter a name for the Sync Source. Enter a unique name for the sync source.
5. On the Base DN Property menu, enter the base DN for the Sync Source. In this example, type **dc=example,dc=com**, and then press Enter when prompted to complete the step.
6. On the Configuring the Server Property menu, select the external servers that will be part of the Sync Source topology. You can enter the number corresponding to the external servers separated by commas. For example, enter **"3,4"** for several external servers.
7. On the Sync Source Properties menu, you can set the **ignore-changes-by-dn** property that specifies the user DN whose modifications on the external server will be ignored during synchronization. This property is useful when using the UnboundID Synchronization Server bidirectionally to limit loop back synchronization changes (modifications) back to the source by the specified user DN (for example, **uid=sync user**). Note: The DN of the user who is performing a delete operation is not normally available in the changelog. So delete operations by these users will not be ignored. Because this example is setting up a one-way sync pipe, you can type **f** to finish.

Configuring the Sync Destination Using dsconfig Interactive Mode

Sync Destinations define the topology of directory servers where changes detected at the Sync Source are applied. Sync Destinations reference one or more external servers of the appropriate type.

To Configure the Sync Destination

1. On the Synchronization Server Configuration console main menu, type the number corresponding to set up the Sync Destination.
2. On the Sync Destination Management menu, type **2** to create a new Sync Destination.

3. Next, on the Sync Destination Type menu, enter the number corresponding to the Sync type (1 for UnboundID Directory Server, 2 for Microsoft Active Directory, 3 for JDBC Sync, 4 for Sun DS Sync). In this example, type **1** for UnboundID Sync Destination.
4. Next, you will be prompted to enter a name for the Sync Destination. Enter a unique name for the Sync Destination.
5. On the Base DN Property menu, enter the base DN for the Sync Destination. In this example, type **dc=example,dc=com**, and then press Enter when prompted to complete the step.
6. On the Server Property menu, select the external servers that will be part of the Sync Destination topology. You can enter the number corresponding to the external servers separated by commas (e.g., "**1,2**").
7. On the Sync Destination Properties menu, type **␣** to save and apply the changes.

Configuring a Sync Pipe Using dsconfig Interactive Mode

A Sync Pipe defines a single synchronization path between the source and destination topologies. Every Sync Pipe has one or more Sync Classes that controls how and what is synchronized. Multiple Sync Pipes can run in a single UnboundID Synchronization Server instance.

Note	Once you have set up a Sync Pipe, remember to start the Sync Pipe for synchronization.
------	--

To Configure a Sync Pipe

1. On the Synchronization Server Configuration console main menu, type the number corresponding to the Sync Pipe. In this example, type **13**.
2. On the Sync Pipe Management menu, type **2** to create a new Sync Pipe.
3. Enter a unique name for the Sync Pipe. A Sync Pipe defines a single synchronization path between the Sync Source and Sync Destination.
4. On the Sync-Source Property menu, select the Sync Source for the Sync Pipe from an existing sync source, or create a new Sync Source if it was not created in an earlier step.
5. On the Sync-Destination Property menu, select the Sync Destination for the Sync Pipe from an existing sync destination, or create a new Sync Destination if it was not created in an earlier step.
6. On the Sync Pipe Properties menu, type **2** to start the Sync Pipe, follow the prompts, and then when done, type **␣** to save and apply the changes. Although the Sync Pipe has started, you must define at least one Sync Class for synchronization to work.
7. Repeat steps 1–6 to create other Sync Pipes. The Synchronization Server can have multiple Sync Pipes in the system. When done, you must define at least one Sync Class for each

Sync Pipe. Within a Sync Pipe, a Sync Class defines each type of entry that needs to be treated differently.

Configuring Sync Class Using dsconfig Interactive Mode

Sync Classes define the operation types (e.g., creates, modifies, or deletes) and attributes that are synchronized, how attributes and DNs are mapped, and how source and destination entries are correlated. A source entry is in at most one Sync Class and is determined by a base DN and LDAP filters. A Sync Class can have multiple Attribute Maps and DN Maps, or none. For each Sync Pipe, a Sync Class is defined for each type of entry that needs to be treated differently.

To Configure a Sync Class for each Sync Pipe

1. On the Synchronization Server Configuration console main menu, type the number corresponding to the Sync Class. In this example, type **11**.
2. On the Sync Class Management menu, type **2** to create a new Sync Class.
3. Select the Sync Pipe that will use the Sync Class. If there is only one Sync Pipe, verify that the existing Sync Pipe is the one that you are configuring, and then press Enter or Return to accept the default.
4. Next, enter a name for the Sync Class that you are defining.
5. On the Sync Class Properties menu, for the **Evaluation Order Index** field, type the priority ordering for the Sync Class if you have more than one Sync Class configured for the topology. Sync Classes with a smaller `evaluation-order-index` property is evaluated first. Because this example defines only one Sync Class, the default value of 9999 is used.
6. For the **Include Base DN** field, type the base DN for the branches of the Sync Source that contain entries in this Sync Class. Only entries with this base DN will be included in the Sync Class. This is an optional step.

If no base DN is specified, the location of the entry in the Sync Source is not taken into account when determining if an entry is part of this Sync Class.
7. For the **Include Filter** field, type a search filter that determines which entries are in the Sync Class. If no filter is specified, all entries within the specified included base DNs are included in the Sync Class.
8. For the **Attribute Map** field, enter an attribute map for the Sync Class. Because this example shows a migration path from Sun Directory Server 5.x to UnboundID Directory Server, you do not need to set up an attribute map, unless you have added new attributes to your schema. See “Configuring Attribute Maps” on page 88.
9. For the **DN Map** field, enter a DN map for the Sync Class. See “Configuring DN Maps” on page 91.

10. On the Sync Class Properties menu, type **£** to save and apply the changes when you have completed configuring the sync class.
11. Repeat steps 1–10 to define another Sync Class for the Sync Pipe.

Starting the Global Sync Configuring Using dsconfig Interactive Mode

After you have set up the Synchronization topology, you must start the Global Sync Configuration, which will use only those Sync Pipes that have been started.

To Start the Global Sync Configuration

1. On the Synchronization Server Configuration console main menu, type the number corresponding to the Global Sync Configuration. In this example, type **7**.
2. On the Global Sync Configuration Management menu, type **1** to view and edit the configuration.
3. On the Global Sync Configuration Properties menu, type **1** to set the **started** property, and then follow the prompts to set the value to **TRUE**.
4. On the Global Sync Configuration Properties menu, type **£** to save and apply the changes.

Preparing the Synchronization Server for External Server Communication

The UnboundID Synchronization Server provides a tool, **prepare-endpoint-server**, that sets up any communication variances that may occur between the Sync Server and the external servers. Typically, directory servers can have different security settings, privileges, and passwords (e.g., for trust stores) configured on the Sync Source that would reject any import of entries in the Sync Destination.

The **prepare-endpoint-server** tool also creates a Synchronization User Account and its privileges on all of the external servers (see “About the Sync User Account” on page 43 for more detailed information). If necessary, you will be prompted for the root or administrator credentials (for example, **uid=admin,dc=example,dc=com**) to set up this user account. The **prepare-endpoint-server** tool also checks if the sync-user account has the proper privileges to access the **firstChangeNumber** and **lastChangeNumber** attributes in the root DSE entry so that it can get the most up-to-date changes to the system. If the Sync User does not have the

proper privileges, the Synchronization Server displays a warning message. You can view any warning or error messages in the `prepare-endpoint-server.log` file in the `logs` directory.

Note	If you created your Synchronization topology using the <code>create-sync-pipe-config</code> tool, then you do not need to run this command separately as it is already part of the process.
------	---

To Prepare the Synchronization Server for External Server Communication

1. Use the `prepare-endpoint-server` tool to prepare the directory server instances on the remote host for synchronization as a data source for the subtree, `dc=example,dc=com`. If the user account is not present on the external server, then the Synchronization Server will create it if it has a parent entry.

```
$ bin/prepare-endpoint-server \  
  --hostname sun-ds1.example.com --port 21389 \  
  --syncServerBindDN "cn=Sync User,dc=example,dc=com" \  
  --syncServerBindPassword secret --baseDn "dc=example,dc=com" \  
  --isSource
```

2. When prompted, enter the bind DN and password to create the user account. This step enables the changelog database and sets the `changelog-maximum-age` property to some recommended value.
3. Repeat steps 1–2 for the other external source servers. Remember to specify the host name and port number of the external server.
4. For the destination servers, repeat steps 2–3 but remember to include the `--isDestination` option. If your destination servers do not have any entries, then a "Denied" message will be generated when creating the `cn=Sync User` entry as no base DN exists.

```
$ bin/prepare-endpoint-server \  
  --hostname UnboundID-ds1.example.com --port 33389 \  
  --syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \  
  --syncServerBindPassword sync --baseDN "dc=example,dc=com" \  
  --isDestination
```

5. Repeat step 4 for the other Destination servers.

Preparing External Servers: If the Admin Does Not Have Root Access on DSEE External Servers

If you are syncing from a Sun DSEE external endpoint server and do not have root access to those machines, you can provide the following manual steps to someone who does have root access on the machines.

To Setup the DSEE External Endpoint Servers

1. Complete the Synchronization Server configuration using the `create-sync-pipe-config` tool or the `dsconfig` commands.

2. Make sure that the Sync User account is created outside of the `cn=config` branch. We have seen problems with DSEE when the Sync User is placed there. If the Sync User is in `cn=config`, delete it, add it to the normal backend (e.g., `dc=example,dc=com`), and then update the configuration in the Sync Server.

For example, create an LDIF file, and save it as "sync-user.ldif".

```
dn: cn=Sync User,dc=example,dc=com
cn: Sync User
givenName: Sync
sn: User
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
userPassword: password
```

Add the following entry to the external DSEE server:

```
$ ldapmodify -h {dsee-host} -p {ldap-port} -D "cn=directory manager" -w {password} \
-a -f sync-user.ldif
```

3. On the DSEE server, perform a search to see what arguments are already set on the retro-changelog plugin. Look for arguments for the attribute, `nsslapd-pluginarg[0-9]`. In the following example, we see that `nsslapd-pluginarg0` and `nsslapd-pluginarg1` are already present, so we need to use `nsslapd-pluginarg2` for any additional settings.

```
$ ldapsearch -h {dsee-host} -p {ldap-port} -D "cn=directory manager" -w {password} \
-b "cn=Retro Changelog Plugin,cn=plugins,cn=config" -s base "(objectclass=*)"
```

```
dn: cn=Retro Changelog Plugin,cn=plugins,cn=config
objectClass: top
objectClass: nsSlapdPlugin
objectClass: ds-signedPlugin
objectClass: extensibleObject
cn: Retro Changelog Plugin
nsslapd-pluginPath: /ds/upc/servers/sunds52/lib/retrocl-plugin.so
nsslapd-pluginInitfunc: retrocl_plugin_init
nsslapd-pluginType: object
nsslapd-plugin-depends-on-type: database
nsslapd-changelogdir: /ds/upc/servers/sunds52/slapd-upc/db/changelog
nsslapd-pluginEnabled: on
nsslapd-changelogmaxage: 3d
nsslapd-pluginarg0: -ignore_attributes
nsslapd-pluginarg1: copyingFrom
nsslapd-pluginId: retrocl
nsslapd-pluginVersion: 5.2_Patch_4
nsslapd-pluginVendor: Sun Microsystems, Inc.
nsslapd-pluginDescription: Retrocl Plugin
ds-pluginSignatureState: valid signature
```

4. On the DSEE server, enable the Retro Changelog Plug-in using the console or command-line tool. Use `ldapmodify` to apply the following LDIF to the server, or you can make the equivalent changes to `dse.ldif` after the server has been shutdown.

The LDIF file enables the retro changelog plugin, sets the max age to three days, and adds the `deletedEntryAttributes` setting into one of the `nsslapd-pluginarg` fields (see below). The `deletedEntryAttributes` attribute is used to ensure that the Sync Server has the proper information for the correlation of deletes against the target system. The attribute

will be used to record `objectclass`, `cn`, `uid`, and `modifiersName` during deletes. You can modify this list of attributes so that the Sync Server can find the corresponding entry in the destination server. In this example, make sure to use the `nsslapd-pluginarg2` attribute name to add the `deletedEntryAttributes` parameters as `nsslapd-pluginarg0` and `nsslapd-pluginarg1` are in use. Finally, save the file as "retro-changelog-enable.ldif"

```
dn: cn=Retro Changelog Plugin,cn=plugins,cn=config
changetype: modify
replace: nsslapd-pluginEnabled
nsslapd-pluginEnabled: on
-
replace: nsslapd-changelogmaxage
nsslapd-changelogmaxage: 3d
-
replace: nsslapd-pluginarg2
nsslapd-pluginarg2: deletedEntryAttributes=objectclass,cn,uid,modifiersName
```

5. Use `ldapmodify` to enable the Retro Changelog Plugin.

```
$ ldapmodify -h {dsee-host} -p {ldap-port} -D "cn=directory manager" -w {password} \
-f retro-changelog-enable.ldif
```

6. Restart DSEE so that the plugin can start recording changes.

7. Create an LDIF file called "sync-dsee-aci.ldif" to add an ACI so that the Sync User can access the changelog and data, respectively.

```
dn: cn=changelog
changetype: modify
add: aci
aci: (targetattr="*")(version 3.0; acl "UnboundID Sync User Access"; allow
(read,search,compare) userdn="ldap:///cn=Sync User,dc=example,dc=com";)

dn: dc=example,dc=com
changetype: modify
add: aci
aci: (targetattr="*")(version 3.0; acl "UnboundID Sync User Read/Write Access";
allow (all) userdn="ldap:///cn=Sync User,dc=example,dc=com";)
```

If the DSEE server is only used as a source, and no modifies will be performed against the server, then the ACI should be as follows:

```
dn: dc=changelog
changetype: modify
add: aci
aci: (targetattr="*")(version 3.0; acl "UnboundID Sync User Access"; allow
(read,search,compare) userdn="ldap:///cn=Sync User,dc=example,dc=com";)

dn: dc=example,dc=com
changetype: modify
add: aci
aci: (targetattr="*")(version 3.0; acl "UnboundID Sync User Read Only Access"; allow
(read,search,compare) userdn="ldap:///cn=Sync User,dc=example,dc=com";)
```

8. Use `ldapmodify` to add the `sync-dsee-aci.ldif` to the DSEE server.

```
$ ldapmodify -h {dsee-host} -p {ldap-port} -D "cn=directory manager" -w {password} \
-f sync-dsee-aci.ldif
```


Using Bulk Resync on the Synchronization Server

The UnboundID Synchronization Server provides a bulk synchronization command-line tool, **resync**, that can be used to verify the Synchronization setup. The **resync** tool operates on a single Sync Pipe at a time and retrieves entries from the Sync Source in bulk, and compares the source entries with the corresponding destination entries. If destination entries are missing or attributes out-of-sync, then the Synchronization Server updates them.

The command provides a **--dry-run** option that can be run to test the matches between Sync Source and Destination but not commit any changes to the target topology. The **resync** tool also provides options to write debugging output to a log with a configurable level of verbosity for testing purposes.

Note

While you can use the **resync** tool to update any mismatched entries, you should use the tool only for relatively small datasets. For large deployments, you can export entries from the Sync Source into an LDIF file, run the **translate-ldif** tool to translate and filter the entries into the destination format, and then import the result LDIF file into the Sync Destination.

Typically, you can use the **resync** tool to verify the synchronization configuration after it has been configured. The command has some important options that can be used to test the configuration:

TABLE 3-1. Useful Resync Command Options

Resync Options	Description
--dry-run	Reports the sync status of the configuration without committing the change to the target topology.
--numPasses	Specifies the number of passes to compare an entry that is out-of-sync to account for synchronization delays. If both Sync Source and Sync Destination are quiescent, then a value of 1 can be provided.
--logFilePath	Specifies the path to the log file that records the details of the resync operation.
--logLevel	Specifies the resync log level that controls the amount of logging. The following levels are available: <ul style="list-style-type: none"> • out-of-sync-summary. Provides a single summary message for each missing or out-of-sync entry • out-of-sync-detailed. Provides a single detailed message including the source and destination entry contents. • all-entries-summary. Provides multiple summary messages, which are logged for every entry that is loaded or compared. The contents of the entries are not included. • all-entries-detailed. Provides multiple detailed messages, which are logged for every entry that is loaded or compared. This option can impact performance as it generates a large output file. • debug. Provides multiple verbose messages, which are logged for every entry that is loaded or compared. This option should only be used to diagnose or troubleshoot a problem as its potential size could impact performance. The contents of the entries are included during processing.

TABLE 3-1. Useful Resync Command Options

Resync Options	Description
<code>--ratePerSecondFile</code>	Specifies a specific synchronization rate (synchronizing changes per second). The option allows you to adjust the rate during off-peak hours or adjust the rate based on measured loads for very long running resync operations. The file must contain a single positive integer number surrounded by white space (for example, 1) to start with. If the file is updated with an invalid number (for example, changing it to zero, a negative number, or something other than an integer number), the rate is not updated. To use this feature, run resync first at 100 operations/sec, measure the impact on the source servers, then adjust as desired.
<code>--secondsBetweenPass</code>	Specifies the number of seconds to wait between each pass to recheck entries that were out-of-sync. This option is used when entries are out-of-sync due to synchronization delays.
<code>--sourceDNsFile</code>	Specifies a file containing a list of DNs to be retrieved from the Sync Source and processed. The option allows for faster processing of very large data sets by targeting individual base-level searches for each source DN in the file.

The **resync** tool provides a number of other useful functions, including the ability to schedule a nightly synchronization if real-time synchronization is not necessary (for example, the creation of new entries during a specific time period can be resynced at a designated nightly time). The tool also provides explicit control over which attributes are included or excluded during the synchronization process if fine-grained synchronization is required by the Attribute or DN maps. For more information, type **bin/resync --help** for information and examples.

Testing Attribute and DN Maps using Resync

You can use the **resync** tool to test how attribute maps and DN maps are configured by synchronizing a single entry. If the `--logFilePath` and `--logLevel` options are specified, the **resync** tool generates a log file with varying degrees of details to show any synchronization messages. You can specify the log file and the level of detail of processing messages.

To Test Attribute and DN Maps using Resync

Use the **resync** tool in "dry run" mode by specifying a single entry. Assume that the Sync Source topology contains an entry, `uid=user.0`.

```
$ bin/resync --pipe-name sun-to-UnboundID-sync-pipe \
  --sourceSearchFilter "(uid=user.0)" --dry-run \
  --logFilePath logs/resync.log --logLevel debug
```

Verifying the Synchronization Configuration using Resync

The most common example for resync is to test that the Sync Pipe configuration has been set up correctly. For example, the following procedure assumes that the configuration was set up with the Sync Source topology (two replicated Sun Directory Server 5.x servers) with 2003 entries; the Sync Destination topology (two replicated UnboundID Directory Servers) has only the base entry and the `cn=Sync User` entry. Both Source and Destination topologies have their LDAP Changelogs enabled. Because both topologies are not actively being updated, the **resync** tool can be run with one pass through the entries.

To Verify the Synchronization Configuration Using resync

Use **resync** with the **--dry-run** option to check the synchronization configuration. The following example does a dry-run process to verify the Sync configuration and creates entries that are not present in the Source Destination. The output also displays a timestamp that can be tracked in the logs.

```
$ bin/resync --pipe-name sun-to-UnboundID-sync-pipe \  
  --sourceSearchFilter "(objectclass=*)" --numPasses 1 --dry-run  
  
Starting Pass 1  
  
Status after completing all passes[20/Mar/2010:10:20:07 -0500]  
-----  
Source entries retrieved      2003  
Entries missing              2002  
Entries out-of-sync          1  
Duration (seconds)           4  
  
Resync completed in 4 s.  
  
0 entries were in-sync, 0 entries were modified, 0 entries were created, 1 entries are  
still out-of-sync, 2002 entries are still missing, and 0 entries could not be processed  
due to an error
```

Populating an Empty Sync Destination Topology Using Resync

The **resync** tool can populate an empty Sync Destination with the Sync Source entries prior to real-time synchronization. If you already have data from the Sync Source in the Sync Destination, you can use the **resync** tool to synchronize entries with the Sync Source.

The following procedures shows how you can use **resync** to populate an empty Sync Destination topology for small datasets. For large deployments, see “To Populate an Empty Sync Destination Topology Using translate-ldif” on page 82.

To Populate an Empty Sync Destination Topology using resync

1. Run **resync** in a dry-run (see the previous example). Assume an error was generated during the process.
2. Rerun the **resync** command with the log file path and with the log level **debug**. Do not include the **--dry-run** option.

```
$ bin/resync --pipe-name sun-to-UnboundID-sync-pipe \  
  --sourceSearchFilter "(objectclass=*)" \  
  --numPasses 1 \  
  --logFilePath logs/resync.log \  
  --logLevel debug
```

3. Open the **logs/resync-failed-DNs.log** file in a text editor to locate the error and fix it. Sometimes an entry could not be created because the parent entry was not present as seen below. You can rerun the **resync** command to create the missing entries.

```
# Entry '(see below)' was dropped because there was a failure at the resource:  
Failed to create entry uid=mlott,ou=People,dc=example,dc=com. Cause: LDAPException(resultCode=no such object, errorMessage='Entry uid=user.38,ou=People,dc=exam-
```

```
ple,dc=com cannot be added because its parent entry ou=People,dc=example,dc=com does
not exist in the server', matchedDN='dc=example,dc=com')
(id=1893859385ResourceOperationFailedException.java:126 Build revision=4881)
dn: uid=user.38,ou=People,dc=example,dc=com
```

4. Rerun the **resync** command. The command creates the entries in the Sync Destination topology that are present in the Sync Source topology.

```
$ bin/resync --pipe-name sun-to-UnboundID-sync-pipe --numPasses 3 \
  --sourceSearchFilter "(objectclass=*)" "
```

...(output from each pass)...

```
Status after completing all passes[20/Mar/2010:10:23:33 -0500]
```

```
-----
Source entries retrieved 160
Entries in-sync          156
Entries created          4
Duration (seconds)      11
```

Resync completed in 12s.

156 entries were in-sync, 0 entries were modified, 4 entries were created, 0 entries are still out-of-sync, 0 entries are still missing, and 0 entries could not be processed due to an error

Populating an Empty Sync Destination Topology Using Translate-LDIF

If you populate a Sync Destination using the **resync** tool, it could take some time to load a large dataset. For a faster method, you can use the **translate-ldif** tool to populate an empty Sync Destination topology for a very large number of entries. The **translate-ldif** tool translates the contents of an LDIF file in Sync Source format to Sync Destination format using the filtering and mapping criteria defined for the Sync Pipe's Sync Classes.

To Populate an Empty Sync Destination Topology Using translate-Ldif

1. On a Sync Source Server, export the data to an LDIF file.
2. On the Synchronization Server, run the **translate-ldif** tool to translate or filter the entries into the Sync Destination format, if necessary. Make sure to specify the path to the LDIF file on the Sync Source server and the path to the output file.

```
$ bin/translate-ldif --pipe-name sun-to-UnboundID-sync-pipe \
  --sourceLDIF /path/to/sync-source-data.ldif \
  --destinationLDIF /path/to/sync-dest-data.ldif
```

3. On a Sync Destination Server, import the data using the path to the translated LDIF file.

Setting the Synchronization Rate Using Resync

The **resync** command has a **--ratePerSecondFile** option that allows you to set a specific synchronization rate (sync changes per second). The option allows you to adjust the rate during off-peak hours or adjust the rate based on measured loads for very long running **resync** operations by simply changing the rate in the file.

To use this feature, run **resync** first at 100 operations/sec, measure the impact on the source servers, then adjust as desired. The file must contain a single positive integer number surrounded by white space (for example, 1) to start with. If the file is updated with an invalid number (for example, changing it to zero, a negative number, or something other than an integer number), the rate is not updated.

To Set the Synchronization Rate Using Resync

1. Create a text file containing the resync rate. The number must be a positive integer surrounded by white space.

```
$ echo ' 100 ' > rate.txt
```

2. Run the **resync** command with the **--ratePerSecondFile** option.

```
$ bin/resync --pipe-name "sun-to-UnboundID-sync-pipe" \  
  --ratePerSecondPath rate.txt
```

Note

The **resync** command also has a **--ratePerSecond** option that allows you to set the sync rates per second by specifying the target rate. The option allows you to throttle resync and reduce its load on the end servers. If this option is not provided, then the tool resyncs at the maximum rate.

3. Check the rate on your system, and then update the rate file again to change the **resync** rate.

```
$ echo ' 150 ' > rate.txt
```

Synchronizing a Specific List of DNs

The **resync** command allows you to synchronize a specific set of DNs that are read from a file using the **--sourceDNsFile** option. The option is best used for very large datasets that require faster processing by targeting individual base-level searches for each source DN in the file. If any DN fails for any reason (parsing, search, or process errors), the command creates an output file of the skipped entries (**resync-failed-DNs.log**), which can be rerun again.

The file must contain only a list of DNs in LDIF format with "dn:" or "dn:". The file can include comment lines by starting each line with a pound sign (#). All DNs can be wrapped and are assumed to be wrapped on any lines that begin with a space followed by text. Empty lines are ignored.

For small files, you can create a file manually. For large files, you can use **ldapsearch** to create an LDIF file, as seen below.

To Synchronize a Specific List of DNs

1. To create a file of DNs, you can enter each manually for small files, or you can run an **ldapsearch** command using the special OID "1.1" extension, which only returns the DNs

in your DIT. For example, on the Sync Source directory server, run the following command:

```
$ bin/ldapsearch --port 1389 --bindDN "uid=admin,dc=example,dc=com \  
--baseDN dc=example,dc=com --searchScope sub "(objectclass=*)" "1.1" >  
dn.ldif
```

2. Run the **resync** command with the **--sourceDNsFile** option to run on individual DN's in the file. For this example,

```
$ bin/resync --pipe-name "sun-to-UnboundID-pipe" --sourceDNsFile dn.ldif
```

```
Starting pass 1
```

```
[20/Mar/2010:10:32:11 -0500]
```

```
-----  
Resync pass                1  
Source entries retrieved    1999  
Entries created             981  
Current pass, entries processed 981  
Duration (seconds)         10  
Average ops/second         98
```

```
Status after completing all passes[20/Mar/2010:10:32:18 -0500]
```

```
-----  
Source entries retrieved    2003  
Entries created             2003  
Duration (seconds)         16  
Average ops/second         98
```

```
Resync completed in 16 s.
```

```
0 entries were in-sync, 0 entries were modified, 2003 entries were created, 0  
entries are still out-of-sync, 0 entries are still missing, and 0 entries could not  
be processed due to an error
```

3. If any errors occurred, view the `<UnboundID-Sync>/resync-failed-DNs.log` to see the skipped DN's. Then, correct the source DN's file, and rerun the **resync** command.

Controlling Real Time Synchronization

The UnboundID Synchronization Server provides a utility to control real-time synchronization including starting and stopping synchronization globally or for individual Sync Pipes. The tool also provides features to set a specific starting point for real-time synchronization, so that changes made before the current time are ignored, and to schedule a stop or start at a future date.

TABLE 3-2. Useful Realtime-Sync Command Options

Realtime-Sync Options	Description
start	Start synchronization globally or for a specific Sync Pipe.

TABLE 3-2. Useful Realtime-Sync Command Options

Realtime-Sync Options	Description
stop	Stop synchronization globally or for a specific Sync Pipe.
set-startpoint	<p>Start synchronization for a specific Sync Pipe at a specified time. When specified, all changes made prior to the current time the command is invoked will be ignored by the Sync Pipe. Additional options include:</p> <ul style="list-style-type: none"> • --change-number {change number}. Begin synchronization at a specific change number in the changelog. This feature cannot be used if the endpoint server is the UnboundID Directory Proxy Server. See “Syncing Through Proxy Servers” on page 145 for more information. • --startpoint-rewind {duration}. Begin synchronization by "rewinding" or starting the synchronization back at a specified duration from the current time. The duration string has the format: d (days), h (hours), m(minutes), s (seconds), ms (milliseconds). For example, to start the synchronization state that occurred 1 day, 2 hours, 12 minutes, and 30 seconds, use "1d2h12m30s". You can also specify milliseconds, for example, "300ms". <p>The set-startpoint option cannot be run on a Sync Pipe that has already started.</p>
Note	To get an accurate picture of the current status of real-time synchronization, view the monitor properties: num-sync-ops-in-flight , num-ops-in-queue , and source-unretrieved-changes .

Starting Real Time Synchronization Globally

You can start real time synchronization globally for all Sync Pipes using the **realtime-sync** tool in the **bin** directory (or **bat** directory for Microsoft Windows systems). The command assumes that you have properly configured your Synchronization topology.

To Start Real Time Synchronization Globally

Use **realtime-sync** to start a synchronization topology globally. Assume that a single Sync Pipe called "dsee-to-UnboundID-sync-pipe" exists.

```
$ bin/realtime-sync start --pipe-name "dsee-to-UnboundID-sync-pipe" \
  --port 389 --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret
```

If you have more than one Sync Pipe configured, specify each Sync Pipe using the **--pipe-name** option. The following example starts **realtime-sync** for a bidirectional synchronization topology.

```
$ bin/realtime-sync start --pipe-name "Sun DS to UnboundID DS" \
  --pipe-name "UnboundID DS to Sun DS" --port 389 \
  --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret
```

Pausing Synchronization

You can pause or unpause synchronization using the 'start' and 'stop' subcommands. If synchronization is stopped and then restarted, then changes made at the Sync Source while synchronization was stopped will still be detected and applied.

Synchronization for individual Sync Pipes can be started or stopped using the `--pipe-name` argument. If the `--pipe-name` argument is omitted, then synchronization is started or stopped globally.

To Stop Real Time Synchronization Globally

Use `realtime-sync` to stop a synchronization topology globally. This command will stop all Sync Pipes started.

```
$ bin/realtime-sync stop --port 389 --bindDN "uid=admin,dc=example,dc=com" \
  --bindPassword secret --no-prompt
```

To Stop an Individual Sync Pipe

Use `realtime-sync` to stop an individual Sync Pipe. Assume the topology has two Sync Pipes, Sync Pipe1 and Sync Pipe2. This command stops Sync Pipe1.

```
$ bin/realtime-sync stop --pipe-name "Sync Pipe1" --port 389 \
  --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret --no-prompt
```

Setting Startpoints

You can set startpoints that instructs the Sync Pipe to ignore all changes made prior to the current time using the `set-startpoint` subcommand with the `realtime-sync` command. Once synchronization is started, only changes made after this command is run will be detected at the Sync Source and applied at the Sync Destination.

The `set-startpoint` subcommand is often run during the initial setup prior to starting real-time synchronization for the first time. It should be run prior to initializing the data in the Sync Destination, which is usually done either by using the `resync` command or by exporting data from the Sync Source, running `translate-ldif`, and then importing the data into the Sync Destination.

The `set-startpoint` subcommand also has two convenient options that can start synchronization at a specific change log number or back at a sync state that occurred at a specific time duration ago (for example, you can start synchronizing at a sync state that occurred 10 minutes ago from the current time).

To Set a Synchronization Startpoint

1. Stop the synchronization topology globally (if it had been started previously) using the `realtime-sync` command with the `stop` subcommand.


```
$ bin/realtime-sync stop --pipe-name "Sync Pipe1" \  
--port 389 --bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

2. Set the startpoint for the synchronization topology. Any changes made before setting this command will be ignored.

```
$ bin/realtime-sync set-startpoint --pipe-name "Sync Pipe1" \  
--port 389 --bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

```
Set StartPoint task 2009072109564107 scheduled to start immediately  
[21/Jul/2009:09:56:41 -0500] severity="INFORMATION" msgCount=0 msgID=1889535170  
message="The startpoint has been set for Sync Pipe 'sun-to-UnboundID-sync-pipe'.  
Synchronization will resume from the last change number in the Sync Source"  
Set StartPoint task 2009072109564107 has been successfully completed
```

To Restart the Sync at a Specific Changelog Event

1. First, search for a specific changelog event from which you want to restart the synchronization state. On one of the endpoint servers, run `ldapsearch` to search the changelog.

```
$ bin/ldapsearch -p 1389 --bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --baseDN cn=changelog --dontWrap  
"(objectclass=*)"   
  
dn: cn=changelog  
objectClass: top  
objectClass: untypedObject  
cn: changelog  
  
dn: changeNumber=1,cn=changelog  
objectClass: changeLogEntry  
objectClass: top  
targetDN: uid=user.13,ou=People,dc=example,dc=com  
changeType: modify  
changes: :  
cmVwbGFjZTogcm9vbU51bWJlcgpyb29tTnVtYmVyOiAwMTM4Ci0KcmVwbGFjZTogbW9kaWZpZXJzTmFtZQptb2RpZmllcnN  
OYWl1OiBjb1EaXJlY3RvcnkgTWFuYWdlcixjb1Sb290IEROcycxbj1jb25maWcKLQpyZXBsYWN1OiBkcy1lcGRhdGUtdG  
ltZQpkcy1lcGRhdGUtdGltZTo6IEFBQUJKZ250W1UwPQotCgA=  
changenumber: 1  
... (more output)  
dn: changeNumber=2329,cn=changelog  
objectClass: changeLogEntry  
objectClass: top  
targetDN: uid=user.49,ou=People,dc=example,dc=com  
changeType: modify  
changes: :  
cmVwbGFjZTogcm9vbU51bWJlcgpyb29tTnVtYmVyOiAwNDMzCi0KcmVwbGFjZTogbW9kaWZpZXJzTmFtZQptb2RpZmllcnN  
OYWl1OiBjb1EaXJlY3RvcnkgTWFuYWdlcixjb1Sb290IEROcycxbj1jb25maWcKLQpyZXBsYWN1OiBkcy1lcGRhdGUtdG  
ltZQpkcy1lcGRhdGUtdGltZTo6IEFBQUJKZ250MC84PQotCgA=  
changenumber: 2329
```

2. Restart synchronization from change number 2329 using the `realtime-sync` tool. Any event before this change number will not be synchronized to the target endpoint.

```
$ bin/realtime-sync set-startpoint --change-number 2329 \  
--server-name "External Server 1" --pipe-name "Sync Pipe 1" \  
--bindPassword secret --no-prompt
```

To Rewind the Sync State by a Specific Time Duration

Use `realtime-sync` with the `--startpoint-rewind` option to "rewind" the synchronization state and begin synchronizing at the specified time duration ago. The following command will start begin synchronizing data at the state that occurred 2 hours and 30 minutes ago from the

current time on External Server 1 for Sync Pipe "Sync Pipe 1". Any changes made before this time will not be synchronized to the target servers. You can specify days (d), hours (h), minutes (m), seconds (s), or milliseconds (ms).

```
$ bin/realtime-sync set-startpoint --startpoint-rewind 2h30m \  
--server-name "External Server 1" --pipe-name "Sync Pipe 1" \  
--bindPassword secret --no-prompt
```

Scheduling a Realtime Sync as a Task

The `realtime-sync` tool features both an offline mode of operation as well as the ability to schedule an operation to run within the Synchronization Server's process. To schedule an operation, supply LDAP connection options that allow this tool to communicate with the server through its task interface. Tasks can be scheduled to run immediately or at a later time. Once scheduled, tasks can be managed using the `manage-tasks` tool.

To Schedule a Realtime Sync Task

1. Use the `--start` option with the `realtime-sync` command to schedule a start for the synchronization topology. The following command will set the start time at July 21, 2009 at 12:01:00 AM. You can also schedule a stop using the `stop` subcommand.

```
$ bin/realtime-sync set-startpoint \  
--pipe-name "sun-to-UnboundID-sync-pipe" \  
--port 389 --bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --start 20090721000100 --no-prompt  
Set StartPoint task 2009072016103807 scheduled to start Jul 21, 2009 12:01:00 AM CDT
```

2. Run the `manage-tasks` tool to manage or cancel the schedule task.

```
$ bin/manage-tasks --port 7389 \  
--bindDN "uid=admin,dc=example,dc=com" --bindPassword secret
```

Configuring Attribute Maps

Attribute Maps are collections of Attribute Mappings, where each mapping defines those destination attributes and value that differ from that of source attributes and how the system will translate the data from one system to another. There are three types of Attribute mappings that can be defined:

- **Direct mapping.** Attributes are directly mapped to another attribute. For example, `employeenumber->employeeid`
- **Constructed Mapping.** Destination attribute values are derived from source attribute values and static text. For example: `{givenname}.{sn}@example.com->mail`

- **DN Mapping.** Attributes are mapped for attributes that store DNs. You can reference the same DN maps that map entry DNs. For example, an attribute called **manager**.

Note	The Synchronization Server automatically validates any attribute mapping prior to applying the configuration.
------	---

Configuring an Attribute Map Using dsconfig Interactive

You can use the **dsconfig** tool in interactive mode to create an attribute map. A Sync Class can reference multiple Attribute Maps. Multiple Sync Classes can share the same Attribute Map.

To Configure an Attribute Map Using dsconfig Interactive

1. On the Synchronization Server Configuration console main menu, type **1** to start up the Attribute Map.
2. On the Attribute Map management menu, type **2** to create a new attribute map.
3. Next, enter a name for the Attribute Map.
4. On the Attribute Map Property menu, type **1** to enter a general description for the Attribute Map. This step is optional. Follow the prompts to enter a description for the Attribute Map. When completed, type **¶** to save the changes and apply.

You have successfully created an attribute map. Next, you must create specific add attribute mappings to your map.

	You can use dsconfig in non-interactive mode to create an attribute for easy scripting.
Note	<pre>\$ bin/dsconfig create-attribute-map \ --map-name test-attribute-map \ --set "description:Test Attribute Map" \ --port 389 --bindDN "uid=admin,dc=example,dc=com" \ --bindPassword secret --no-prompt</pre>

Configuring an Attribute Mapping Using dsconfig Interactive

You can use the **dsconfig** tool in non-interactive mode to create one or more attribute mappings. In this example, the Attribute Mapping sets up a Direct Mapping for one attribute:

employeeNumber -> employeeID.

Note

The Synchronization Server provides a **scramble-value** advanced property that can be configured with each Attribute Mapping. The scramble feature allows you to load a Sync Destination topology with the scrambled values of real production data attributes. Obfuscating production data is convenient in testing environments.

To Configure an Attribute Mapping Using dsconfig Interactive

1. On the Synchronization Server Configuration console main menu, type **2** to start up the Attribute Mapping.
2. On the Attribute Mapping management menu, type **2** to create a new mapping.
3. Select the attribute map that you want to configure. If there is only one Attribute Map, press Enter or Return to accept the default.
4. Select the type of Attribute Mapping that you want to create: 1 for Constructed, 2 for Direct Attribute, 3 for DN Attribute. In this example, type **2** for a Direct Attribute Mapping.
5. Enter a name for the **to-attribute** for the Direct Attribute Mapping. For this example, type **employeeID**.
6. Enter a name for the **from-attribute** for the Direct Mapping. For this example, type **employeeNumber**.
7. On the Directory Attribute Mapping menu, type **f** to save and apply the changes.
8. After you have configured your Attribute Mappings, remember to add the new Attribute Map to a new Sync Class or modify an existing Sync Class.

Configuring an Attribute Mapping Using dsconfig Non-Interactive

You can use the **dsconfig** tool in non-interactive mode to create an attribute mapping. You can view the log of all configuration changes in the **logs/config-audit.log** as well as view the analogous commands to back out of each change.

To Configure an Attribute Mapping Using dsconfig Non-Interactive

1. On the Synchronization Server, use **dsconfig** in non-interactive mode to create an Attribute Mapping.

```
$ bin/dsconfig create-attribute-mapping \  
  --map-name test-attribute-map \  
  --to-attribute employeeID \  
  --from-attribute employeeNumber
```

```
--mapping-name employeeID \  
--type direct \  
--set from-attribute:employeeNumber \  
--port 389 --bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

2. After you have configured your Attribute Mappings, remember to add the new Attribute Map to a new Sync Class or modify an existing Sync Class.

```
$ bin/dsconfig set-sync-class-prop \  
--pipe-name test-sync-pipe \  
--class-name test-sync-class \  
--set attribute-map:test-attribute-map \  
--port 389 --bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --no-prompt
```

Configuring the Directory Server Backend for Synchronizing Deletes

One important attribute that must be configured on the directory server's changelog backend is the **changelog-deleted-entry-include-attribute** property. The property specifies which attributes should be recorded in the changelog entry during a DELETE operation. Normally, the Sync Server cannot correlate a deleted entry to the entry on the destination as there is not enough information to figure out what was deleted. If you have a Sync Class configured with a filter, such as, "**include-filter: objectClass=person**," then you need the **object-class** attribute to be recorded in the changelog entry. Likewise, if you have special correlation attributes (other than DN), you will need those attributes recorded on the changelog entry to be properly synchronized at the endpoint server.

To Configure the Changelog-Deleted-Entry-Include-Attribute Property

- On each directory server backend (UnboundID Directory Server), use the **dsconfig** command to set the property. Remember to add the connection parameters specific to your server (hostname, port, bind DN, and bind DN password).

```
$ bin/dsconfig set-backend-prop --backend-name changelog --set changelog-deleted-  
entry-include-attribute:objectClass
```

Configuring DN Maps

Similar to Attribute Maps, DN Maps define mappings when destination DNs differ from source DNs. These differences must be resolved using DN Maps in order for synchronization to successfully take place. For example, the Sync Source could have a DN in the following format:

```
uid=jdoe,ou=People,dc=example,dc=com
```

While the Sync Destination could have the standard X.500 DN format:

```
cn=John Doe,ou=Engineering,o=example.com
```

A DN Map can hold one or more DN mappings and can be defined as follows:

- **Wildcards.** DN Mappings allow the use of wild cards for DN transformations. A single wild card ("*") matches a single RDN component and can be used any number of times. The double wild card ("**") matches zero or more RDN components and can be used only once. The wild card values can be used in the `to-dn-pattern` attribute using "{1}" to replace their original index position in the pattern, or "{attr}" to match an attribute value. For example:

```
** ,dc=myexample,dc=com->{1},o=example.com
```

- **Regular Expressions.** You can also use regular expressions and attributes from the user entry. For example, the following expression constructs a value for the `uid` attribute, which is the RDN, out of the initials (first letter of `givenname` and `sn`) and the employee ID (the `eid` attribute) of a user.

```
uid={givenname:/^(.)(.*)/$1/s}{sn:/^(.)(.*)/$1/s}{eid},{2},o=example
```

Note

The Synchronization Server automatically validates any DN mapping prior to applying the configuration.

Configuring a DN Map Using dsconfig Interactive

You can use the `dsconfig` tool in interactive mode to create a DN Map. A Sync Class can reference multiple DN Maps. Multiple Sync Classes can share the same DN Map.

To Configure a DN Map Using dsconfig Interactive

1. On the Synchronization Server Configuration console main menu, type **4** to start up the DN Map.
2. On the DN Map management menu, type **2** to create a new DN map.
3. Enter a unique name for the DN Map.
4. For the `from-dn-pattern` property, enter a value. For example, type `** ,dc=myexample,dc=com`.
5. For the `to-dn-pattern` property, enter a value. For example, type `{1},o=example.com`.
6. On the DN Map Properties menu, type **1** to enter a general description for the DN Map. This step is optional. Follow the prompts to enter a description for the DN Map. When completed, type **£** to save the changes and apply.
7. After you have configured your DN Mappings, remember to add the new DN Map to a new Sync Class or modify an existing Sync Class.

Configuring a DN Map Using dsconfig Non-Interactive Mode

You can configure a DN Map using the **dsconfig** tool in non-interactive mode that can be included in a setup script in another Synchronization Server installation. Make sure that you understand the mapping process. If you need any assistance, contact UnboundID.

To Configure a DN Map Using dsconfig Non-Interactive Mode

1. Use **dsconfig** to create a DN Map for the Synchronization Server.

```
$ bin/dsconfig create-dn-map \  
  --map-name nested-to-flattened \  
  --set "from-dn-pattern:*,*,dc=example,dc=com" \  
  --set "to-dn-pattern:uid={givenname:/^(.) (.*)/\$1/s}{sn:/^(.) (.*)/\$1/  
s}{eid},{2},o=example" \  
  --port 1389 --bindDN "uid=admin,dc=example,dc=com" \  
  --bindPassword secret --no-prompt
```

2. After you have configured your DN Mappings, remember to add the new DN Map to a new Sync Class or modify an existing Sync Class.

```
$ bin/dsconfig set-sync-class-prop \  
  --pipe-name test-sync-pipe \  
  --class-name test-sync-class \  
  --set dn-map:test-dn-map \  
  --port 389 --bindDN "uid=admin,dc=example,dc=com" \  
  --bindPassword secret --no-prompt
```

Configuring Fractional Replication

The UnboundID Synchronization Server supports fractional replication to any type of server. For example, if a replica only performs user authentications, then the Synchronization Server can be configured to propagate, for example, only the **uid** and **userpassword** password policy attributes, reducing the database size at the replica and the network traffic needed to keep this server in sync to this server.

The following example presents a fractional replication use case, where the **uid** and **userPassword** attributes of all entries in the Source topology are synchronized to the Destination topology. Because the **uid** and **userPassword** attributes are present, you also need to synchronize the **objectclass** attribute. The example assumes that you have already configured a Synchronization Server and defined the sync pipe, sync class, and external servers but have not run realtime synchronization or bulk resync.

To Configure Fractional Replication

1. On the Synchronization Server Configuration console main menu, type the number corresponding to Sync Classes.

2. On the Sync Class management menu, type **3** to view and edit an existing Sync Class. Assume that only one Sync Class has been defined thus far.
3. Verify that the Sync Pipe and Sync Class exist.
4. On the Sync Class Properties menu, type **4** to specify the source LDAP filter (include-filter property) that defines which source entries be included in the Sync Class.
5. On the Include-Filter Property menu, type **2** to add a filter value. For this example, type (**objectclass=person**). You will be prompted to enter another filter, press **Enter** or **Return** to continue. On the menu, enter **1** to use the value when specifying it.
6. On the Sync Class Properties menu, type **7** for the **auto-mapped-source-attribute** property. When you change the value from **"-all-"** to a specific attribute, then only the specified attribute is automatically mapped from the Source topology to the Destination topology.
7. On the Auto-Mapped-Source-Attribute Property menu, type **2** to add the source attributes that will be automatically mapped to the Destination attributes of the same name. When prompted, enter each attribute, and then press **Enter**.

Continuing from the previous step, type the attributes that will be auto-mapped from the Source to the Destination topology, and then press **Enter**. Repeat the process for each attribute that you want to define.

```
Enter another value for the 'auto-mapped-source-attribute' property [continue]: uid
Enter another value for the 'auto-mapped-source-attribute' property [continue]:
userPassword
Enter another value for the 'auto-mapped-source-attribute' property [continue]:
objectclass
Enter another value for the 'auto-mapped-source-attribute' property [continue]:
```

8. On the Auto-Mapped-Source-Attribute Property menu, type **3** to remove one or more values. In this example, we want to remove the **"-all-"** value, so that only the **objectclass**, **uid**, and **userPassword** attributes are only synchronized.
9. On the Auto-Mapped-Source-Attribute Property menu, press **Enter** to accept the values.
10. On the Sync Class Properties menu, type **8** to exclude some attributes from the synchronization process. Because we are using the **objectclass=person** filter, we must exclude the **cn**, **givenName**, and **sn** attributes.

Continuing from the previous step, type **2** to add one or more attributes, and then, add each attribute that you want to exclude on the Excluded-Auto-Mapped-Source-Attributes Property menu. Here, we want to exclude the **cn**, and **sn** attributes, which are required attributes of the **Person** objectclass. We also exclude the **givenName** attribute, which is an optional attribute of the **inetOrgPerson** objectclass.

```
Enter another value for the 'excluded-auto-mapped-source-attributes' property [continue]: givenName
Enter another value for the 'excluded-auto-mapped-source-attributes' property [continue]: sn
Enter another value for the 'excluded-auto-mapped-source-attributes' property [continue]:
```


11. On the Excluded-Auto-Mapped-Source-Attributes Property menu, confirm your selections, and then press **Enter** to accept the changes.
12. On the Sync Class Properties menu, review the configuration, and then type **Y** to accept the changes.
13. On the server instances in the Destination topology, you must turn off schema checking due to a schema error that occurs when the required attributes in the **Person** objectclass are not present. The command assumes that you have already set the global configuration property for the server-group to "all-servers". You can use **bin/dsconfig** with the **--applyChangeTo server-group** in non-interactive mode to turn off schema checking on all of the servers in the group.

```
$ bin/dsconfig set-global-configuration-prop \  
  --set check-schema:false --applyChangeTo server-group \  
  --port 3389 --bindDN "uid=admin,dc=example,dc=com" \  
  --bindPassword secret --no-prompt
```

14. Run **bin/resync** to load the filtered data from the source endpoint to the target endpoint.

```
$ bin/resync --pipe-name "test-sync-pipe" --numPasses 3
```

15. Run **bin/realtime-sync** to start synchronization.

```
$ bin/realtime-sync start --pipe-name "test-sync-pipe" \  
  --port 7389 --bindDN "uid=admin,dc=example,dc=com" \  
  --bindPassword secret --no-prompt
```

You have successfully configured a fractional replication example.

Managing Failover Behavior

The Synchronization Server delivers high availability in production environments using robust failover mechanisms. To illustrate the generalized failover behavior of the Synchronization Server, Figure 3-1 shows a simplified synchronization topology with a single failover server on the source, destination, and Synchronization Server, respectively. The gray lines represent possible failover connections in the event the server is down. It is assumed that the external servers are prioritized so that src1 has higher priority than src2; dest1 has higher priority than dest2.

The main Synchronization Server and its redundant failover instance communicate with each other over LDAP and bind using "cn=IntraSync User,cn=Root DNs,cn=config". The servers run periodic health checks on each other and share information on all changes that have been processed. Whenever the failover server loses connection to the main Sync Server, for example, during a ping or an LDAP search request, it assumes that the main server is down and begins processing changes from the last known change. Control reverts back to the main server once it is back online.

Unlike the Synchronization Servers, the external servers and their corresponding failover server(s) do not run periodic health checks. If an external server goes offline (e.g., dest1), the failover server (e.g., dest2) will receive transactions and remained connected to the Synchroni-

zation Server until the Sync Pipe is restarted, regardless if the main external server goes back online.

FIGURE 3-1. The Synchronization Server in a Simplified Setup





FIGURE 3-2. The Synchronization Server during Failover



Conditions that Trigger Immediate Failover

Immediate failover occurs when the Synchronization Server encounters the following error code that are returned from an external server. (The error code numbers are presented in parentheses.):

- BUSY (51)
- UNAVAILABLE (52)
- SERVER CONNECTION CLOSED (81)
- CONNECT ERROR (91)

For example, if the Synchronization Server attempts a write operation to a target server (e.g., src1 or dest1) that is in lock-down mode, the Sync Server will see a returned **UNAVAILABLE** error code. The Synchronization Server will then automatically fail over to the next highest prioritized redundant server instance in the target topology (e.g., src2 or dest2), issue an alert, and then reissue the retry attempt. If the operation is unsuccessful for any reason, the server logs the error.

Configuration Properties that Control Failover Behavior

The Synchronization Server's out-of-the-box configuration settings should meet the requirements for most applications. Administrators should be aware of four important advanced properties to fine tune the failover mechanism (each property presented in the next section):

- max-operation-attempts (sync pipe)
- response-timeout (source and destination endpoints)
- max-failover-error-code-frequency (source and destination endpoints)
- max-backtrack-replication-latency (source endpoints only)

These properties apply to the following LDAP error codes:

TABLE 3-3. LDAP Error Codes

Error Codes	Description
ADMIN_LIMIT_EXCEEDED (11)	Indicates that processing on the requested operation could not continue, because an administrative limit was exceeded.
ALIAS_DEREFERENCING_PROBLEM (36)	Indicates that a problem was encountered while attempting to dereference an alias for a search operation.
CANCELED (118)	Indicates that a cancel request was successful, or that the specified operation was canceled.
CLIENT_SIDE_LOCAL_ERROR (82)	Indicates that a local (client-side) error occurred.
CLIENT_SIDE_ENCODING_ERROR (83)	Indicates that an error occurred while encoding a request.
CLIENT_SIDE_DECODING_ERROR (84)	Indicates that an error occurred while decoding a request.
CLIENT_SIDE_TIMEOUT (85)	Indicates that a client-side timeout occurred.
CLIENT_SIDE_USER_CANCELLED (88)	Indicates that a user cancelled a client-side operation.
CLIENT_SIDE_NO_MEMORY (90)	Indicates that the client could not obtain enough memory to perform the requested operation.
CLIENT_SIDE_CLIENT_LOOP (96)	Indicates that a referral loop is detected.
CLIENT_SIDE_REFERRAL_LIMIT_EXCEEDED (97)	Indicates that the referral hop limit was exceeded.
DECODING_ERROR (84)	Indicates that an error occurred while decoding a response.
ENCODING_ERROR (83)	Indicates that an error occurred while encoding a response.
INTERACTIVE_TRANSACTION_ABORTED (30221001)	Indicates that an interactive transaction was aborted.
LOCAL_ERROR (82)	Indicates that a local error occurred.
LOOP_DETECT (54)	Indicates that a referral or chaining loop was detected while processing a request.
NO_MEMORY (90)	Indicates that not enough memory could be obtained to perform the requested operation.
OPERATIONS_ERROR (1)	Indicates that an internal error prevented the operation from being processed properly.
OTHER (80)	Indicates that an error occurred that does not fall into any of the other categories.
PROTOCOL_ERROR (2)	Indicates that the client sent a malformed or illegal request to the server.

TABLE 3-3. LDAP Error Codes

Error Codes	Description
TIME_LIMIT_EXCEEDED (3)	Indicates that a time limit was exceeded while attempting to process the request.
TIMEOUT (85)	Indicates that a timeout occurred.
UNWILLING_TO_PERFORM (53)	Indicates that the server is unwilling to perform the requested operation.

max-operation-attempts

The **max-operation-attempts** property (part of the Sync Pipe configuration) specifies the maximum number of times to retry a synchronization operation that fails for reasons other than the Sync Destination being busy, unavailable, server connection closed, or connect error.

To Change the max-operation-attempts Property

To change the default number of retries, use **dsconfig** in non-interactive mode to change the **max-operation-attempts** value on the Sync Pipe object. The following command changes the number of maximum attempts from 5 (default) to 4. Remember to include the LDAP or LDAPS connection parameter (hostname, port, bindDN, bindDN password).

```
$ bin/dsconfig set-sync-pipe-prop --pipe-name "Test Sync Pipe" \
--set max-operation-attempts:4
```

response-timeout

The **response-timeout** property (part of the Sync Source and Sync Destination configuration) specifies how long the Synchronization Server should wait for a response from a search request to a source server before failing with LDAP result code 85 (client-side timeout). When a client-side timeout occurs, the Sync Source will retry the request according to the **max-failover-error-code-frequency** property before failing over to a different source server and performing the retry. The total number of retries will not exceed the **max-operation-attempts** property defined in the Sync Pipe configuration. A value of zero indicates that there should be no client-side timeout. The default value is one minute.

To Change the response-timeout Property

To set the **response-timeout** property, use the **dsconfig** tool to set it. Assuming a bidirectional topology, you can set the property on the Sync Source and Sync Destination, respectively. Remember to include the LDAP or LDAPS connection parameter (hostname, port, bindDN, bindDN password).

```
$ bin/dsconfig set-sync-source-prop --source-name src --set "response-timeout:8 s"
$ bin/dsconfig set-sync-destination-prop --destination-name U4389 --set "response-timeout:9 s"
```

max-failover-error-code-frequency

The **max-failover-error-code-frequency** property (part of the Sync Source configuration) specifies the maximum time period that an error code can re-appear until it fails over to another server instance. This property allows the retry logic to be tuned, so that retries can be performed once on the same server before giving up and trying another server. The value can be set to zero if there is no acceptable error code frequency and failover should happen immediately. It can also be set to a very small value (such as, 10 ms) if a high frequency of error codes is tolerable. The default value is 3 minutes.

To Change the max-failover-error-code-frequency Property

To change the maximum failover error code frequency, use **dsconfig** in non-interactive mode to change the property on the Sync Source object. The following command changes the frequency from 3 minutes to 2 minutes. Remember to include the LDAP or LDAPS connection parameter (hostname, port, bindDN, bindDN password) with the **dsconfig** command.

```
$ bin/dsconfig set-sync-source-prop --source-name source1 \  
--set "max-failover-error-code-frequency:2 m"
```

max-backtrack-replication-latency

The **max-backtrack-replication-latency** property (part of the Sync Source configuration) sets the time period that a new Synchronization Server will look for any missed changes in the changelog to account for any changes that come in due to replication delays. The property should be set to a conservative upper-bound of the maximum replication delay between two servers in the topology. A value of zero implies that there is no limit on the replication latency. The default value is 2 hours. The Synchronization Server stops looking in the changelog once it finds a change that is older than the maximum replication latency than the last change it processed on the failed server.

For example, after failing over to another server, the Synchronization Server must look through the new server's changelog to find the equivalent place to begin synchronizing any changes. Normally, the Synchronization Server can successfully backtrack with only a few queries of the directory, but in some situations, it might have to look further back through the changelog to make sure that no changes were missed. Because the changes can come from a variety of sources (replication, synchronization, and over LDAP), the replicated changes between directory servers are interleaved in each changelog. When the Synchronization Server fails over between servers, it has to backtrack to figure out where synchronization can safely pick up the latest changes.

Backtracking occurs until the following:

- It determines that there is no previous changelog state available for any source servers, so it must start at the beginning of the change log.
- It finds the last processed replication change sequence number (CSN) from the last time it was connected to that replica, if at all. This process is similar to the "set-startpoint" functionality on the **realtime-sync** tool.

- It finds the last processed replication CSN from every replica that has produced a change so far, and it determines that each change entry in the next oldest batch of changes has already been processed.
- It finds a change that is separated by more than a certain duration (specified by the `max-backtrack-replication-latency` property) from the most recently processed change.

Changing the max-backtrack-replication-latency Property

To change the maximum backtrack replication, use `dsconfig` in non-interactive mode to change the `max-backtrack-replication-latency` value to some time period. The following command changes the maximum backtracking from two hours to three hours. Remember to include the LDAP or LDAPS connection parameter (hostname, port, bindDN, bindDN password) with this command.

```
$ bin/dsconfig set-sync-source-prop --source-name source1 \  
--set "max-backtrack-replication-latency:3 h"
```

4 Configuring Synchronization with Microsoft Active Directory Systems

Overview

The UnboundID Synchronization Server supports full synchronization for newly created or modified accounts with native password changes between source and destination topologies comprising of UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Servers (3.x), Sun Directory Server Enterprise Edition (DSEE 6.x, 7.x), Sun Directory Server (5.2 patch 3 or higher), Oracle (10g, 11g), Microsoft SQL Server (2005, 2008) with Microsoft Active Directory systems. Synchronization with Active Directory systems provides a robust and scalable solution for large multi-directory and multi-national networks. The Synchronization Server also delivers immediate failover capabilities to source and destination instances without data loss in case the target systems go down.

This chapter presents the configuration procedures needed to set up synchronization between UnboundID Directory Server, Alcatel-Lucent 8661 Directory Server, Sun DSEE, or Sun Directory Server source or targets with Microsoft Active Directory systems:

- [Before You Begin](#)
- [Configuring Active Directory Synchronization](#)
- [Installing the UnboundID Password Sync Agent](#)

Before You Begin

Synchronization with Microsoft Active Directory systems requires that SSL be enabled on the Active Directory domain controller, so that the UnboundID Synchronization Server can securely propagate the `cn=Sync User` account password and other user passwords to the Active Directory target. Likewise, the UnboundID Synchronization must be configured to accept SSL connections.

- For information on setting up an SSL connection on the Synchronization Server, see “Installing the Synchronization Server” on page 23.

Configuring Active Directory Synchronization

To install and configure synchronization with Active Directory systems, you must do the following steps depending on the type of synchronization data:

- On the Synchronization Server, use the **create-sync-pipe-config** tool to configure the Sync Pipes to communicate with the Active Directory source or target.
- If you plan to provide outbound password synchronization from the UnboundID Directory Server, enable the Password Encryption component on all UnboundID Directory Server sources that receive password modifications. The UnboundID Directory Server uses the Password Encryption component, analogous to the Password Sync Agent component, to intercept password modifications and add an encrypted attribute, **ds-changelog-encrypted-password**, to the changelog entry. The component allows passwords to be synced securely to the Active Directory system, which uses a different password storage scheme. The encrypted attribute appears in the changelog and gets synchronized to the other servers but does not appear in the entries. If you do not plan to synchronize passwords, you can skip this step.
- If you plan to provide outbound password synchronization from the Active Directory system, install the Password Sync Agent (PSA) presented in “Installing the UnboundID Password Sync Agent” on page 108 after configuring the Synchronization Server. If you do not plan to synchronize passwords, you can skip this step.

To Configure Active Directory Synchronization

The following procedure configures a one-way sync pipe with the Active Directory topology as the Sync Source and the UnboundID Directory Server topology as the Sync Destination.

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

2. From the server-root directory, start the Synchronization Server.

```
$ bin/start-sync-server
```

3. Run the **create-sync-pipe-config** tool to set up the initial Synchronization topology.

```
$ bin/create-sync-pipe-config
```

4. On the Initial Synchronization Configuration Tool menu, press Enter or Return (yes) to continue the configuration.
5. On the Synchronization Directory menu, select if the Synchronization topology will be one-way (1) or bidirectional (2). In this example, enter "2" for bidirectional.

6. On the Source Endpoint Type menu, enter 6 for Microsoft Active Directory.

```
>>>> Source Endpoint Type
```

```
Enter the type of data store for the source endpoint:
```

- 1) UnboundID Directory Server
- 2) UnboundID Proxy Server
- 3) Alcatel-Lucent Directory Server
- 4) Alcatel-Lucent Proxy Server
- 5) Sun Directory Server
- 6) Microsoft Active Directory
- 7) Microsoft SQL Server
- 8) Oracle Database
- 9) Generic JDBC

- b) back
- q) quit

```
Enter choice [2]:
```

7. On the Source Endpoint Name menu, type a name for the Source Server, or accept the default ("Microsoft Active Directory Source"). For this example, use the default "ADSync".
8. On the Base DN menu, type the base DN on the Source Server where the entries will be searched. In this example, type `dc=adsync,dc=UnboundID,dc=com`. If you have other base DNs, type the DN or press Enter when finished. You also have the option to remove the base DN if desired.
9. On the Server Security menu, select the security connection type for the source server, which will be SSL by default for Active Directory configurations. Note that any connection with the Active Directory topology **requires** an SSL connection, while connections with the UnboundID Directory Server, Sun DSEE, or Sun Directory Server can use a standard LDAP or SSL connection.
10. On the Servers menu, enter the host name and listener port number for the Source Server, or accept the default (port 389). The server will attempt a connection to the server. If the server is unresponsive, you will be asked to retry `<hostname>:636`, contact, discard, or keep the server.

After entering the first server, enter the additional servers (hostname:port) for the source endpoints, which will be prioritized below the first server. You also have the option to remove any existing servers.
11. On the Synchronization User Account DN menu, enter the User Account DN for the source servers. The account will be used exclusively by the Synchronization Server to communicate with the source external servers. This step will ask you to enter a User Account DN and password, or accept the default account DN (`cn=SyncUser,cn=Users,DC=adsync,DC=UnboundID,DC=com`). The User Account DN password must meet the minimum password requirements for Active Directory domains.
12. At this point, you must set up the Destination Endpoint servers. The setup steps are similar to steps 6–11. Select the option for UnboundID and then set up an external destination server and User Account DN.

Preparing the External Servers

1. After you have configured the Source and Destination Endpoints, the Synchronization Server will prompt you to "prepare" each external server. This step entails asking you if you trust the certificate presented to it, and then testing the connection. The following example shows the user interaction involved in preparing one external server. If you do not prepare the external servers, you can do so after configuring the Sync Pipes using the `prepare-endpoint-server` tool.

The following shows example a snippet of a user session:

```
>>>> Prepare Server '10.8.1.163:636'
```

```
Servers in a synchronization topology must be 'prepared' for synchronization. This involves making sure the synchronization user account exists and has the proper privileges.
```

```
Would you like to prepare server '10.8.1.163:636' for synchronization?
```

- 1) Yes
- 2) No

- b) back
- q) quit

```
Enter choice [1]:
```

```
Testing connection to 10.8.1.163:636
```

```
Do you wish to trust the following certificate?
```

```
Certificate Subject: CN=WIN-G2R2NXV87VX.adsync.UnboundID.com
```

```
Issuer Subject: CN=adsync-WIN-G2R2NXV87VX-CA,DC=adsync,DC=UnboundID,DC=com
```

```
Validity: Thus Nov 12 11:39:52 CST 2009 to Fri Nov 12 11:39:52 CST 2010
```

```
Enter 'y' to trust the certificate or 'n' to reject it.
```

```
y
```

```
Testing connection to 10.8.1.163:636 ..... Done
```

```
Testing 'cn=Sync User,cn=Users,DC=adsync,DC=UnboundID,DC=com' access ..... Done
```

```
Configuring this server for synchronization requires manager access. Enter the DN of an account capable of managing the external directory server [cn=Administrator,cn=Users,DC=adsync,DC=UnboundID,DC=com]:
```

```
Enter the password for 'cn=Administrator,cn=Users,DC=adsync,DC=unbound,DC=com':
```

```
Verifying base DN 'dc=adsync,dc=UnboundID,dc=com' ..... Done
```

2. Next, you will be prompted if you want to prepare another server in the topology. Repeat the process for each server that you have configured in the system.

Configuring the Sync Pipes and its Sync Classes

1. Next, on the Sync Pipe Name menu, you will be prompted to set up the Sync Pipe name. Type a unique name to identify the Sync Pipe or accept the default.
2. On the Pre-Configured Sync Class Configuration for Active Directory Sync Source menu, enter "yes" if you want to synchronize user **CREATE** operations, and then enter the object class for the user entries at the destination server (default is `inetOrgPerson`). Next, you

will be prompted if you want to synchronize user **MODIFY** and **DELETE** operations from Active Directory. Enter "yes" if you want to do so.

3. Next, you will be asked if you want to synchronize user passwords from Active Directory. Press Enter or Return to accept the default (yes). If you plan to synchronize passwords from Active Directory, you must also install the UnboundID Password Sync Agent component on each domain controller. See “Installing the UnboundID Password Sync Agent” on page 108 for more information.
4. Next, you will be asked if you want to create a DN map for the user entries in the Sync Pipe. Enter the base DN for the user entries at the Microsoft Active Directory Sync Source (for example, **CN=Users,DC=adsync,DC=UnboundID,DC=com**), and then enter the base DN for the user entries at the UnboundID Directory Server Sync Destination (for example, **OU=users,DC=adsync,DC=UnboundID,DC=com**).
5. At this stage, you will see a list of basic attribute mappings from the Microsoft Active Directory Source to the UnboundID Directory Server destination. If you want to add more complex attribute mappings involving constructed or DN attribute mappings, you must quit the command and use the **dsconfig** tool. The following example shows a sample user session.

Below is a list of the basic mappings that have been set up for user entries synchronized from Microsoft Active Directory -> UnboundID Directory Server. You can add to or modify this list with any direct attribute mappings. To set up more complex mappings (such as constructed or DN attribute mappings), use the 'dsconfig' tool.

```
1)  cn -> cn
2)  sn -> sn
3)  givenName -> givenName
4)  description -> description
5)  sAMAccountName -> uid
6)  unicodePwd -> userPassword

b)  back
q)  quit
n)  Add a new attribute mapping
```

Enter "n" to add a new attribute mapping. First, enter the source attribute, and then enter the destination attribute. The following example shows a sample user session, where a mapping is set up for the **telephoneNumber** attribute (Active Directory) to the **otherTelephoneNumber** attribute (UnboundID).

```
Select an attribute mapping to remove, or choose 'n' to add a new one [Press ENTER to continue]: n
```

```
Enter the name of the source attribute: telephoneNumber
```

```
Enter the name of the destination attribute: otherTelephoneNumber
```

6. Next, enter **yes** if you want to synchronize group **CREATE**, **MODIFY**, and **DELETE** operations from Active Directory.
7. Next, type "yes" if you want to synchronize group entries from Active Directory. Then, review the basic user group mappings. If you want to use more complex mappings, such as constructed or DN attribute mappings, use the **dsconfig** tool. If you want to add new group

attribute mappings, type "n". In this example, press Enter to continue as no new group mappings will be created. The following example shows a portion of a user session.

```
Would you like to sync group entries from Active Directory? (yes / no) [no]: yes
```

```
Below is a list of the basic mappings that have been set up for user entries in the
Sync Class: 'AD Groups Sync Class'. You can add to or modify this list with any
direct attribute mappings. To set up more complex mappings (such as constructed or
DN attribute mappings), use the 'dsconfig' tool.
```

- 1) {cn} -> {cn} (Direct Mapping)
- 2) 'groupOfUniqueNames' -> {objectClass} (Constructed Mapping)
- 3) {member} -> {uniqueMember} (Direct Mapping)

- b) back
- q) quit
- n) Add a new direct attribute mapping

```
Select an attribute mapping to remove, or choose 'n' to add a new one [Press ENTER
to continue]:
```

8. On the Sync Pipe Sync Class Definitions menu, enter another name for a new Sync class if required. You will essentially repeat the steps in 2–7 to define this new Sync Class. You also have the option to remove any existing sync classes already defined. If you do not require any additional sync class definitions, press Enter to continue.

9. Review the Sync Pipe Configuration Summary, and then, press Enter to accept the default ("write configuration"), which records the commands in a batch file (`sync-pipe-cfg.txt`). The batch file can be re-used to set up other Sync topologies. The following summary shows two Sync Pipes (from Active Directory to UnboundID; the other, from UnboundID to Active Directory) and its associated Sync Classes.

```
>>>> Configuration Summary
```

```
Sync Pipe: AD to UnboundID
Source: Microsoft Active Directory
Type: Microsoft Active Directory
Access Account: cn=Sync User,cn=Users,DC=adsync,DC=UnboundID,DC=com
Base DN: DC=adsync,DC=UnboundID,DC=com
Servers: 10.5.1.149:636

Destination: UnboundID Directory Server
Type: UnboundID Directory Server
Access Account: cn=Sync User,cn=Root DNs,cn=config
Base DN: dc=example,dc=com
Servers: localhost:389

Sync Classes:
Microsoft Active Directory Users Sync Class
Base DN: DC=adsync,DC=UnboundID,DC=com
Filters: (objectClass=user)
DN Map: **,CN=Users,DC=adsync,DC=UnboundID,DC=com ->{1},ou=users,dc=example,dc=com
Synchronized Attributes: Custom set of mappings are defined
Operations: Creates,Deletes,Modifies

Sync Pipe: UnboundID to AD
Source: UnboundID Directory Server
Type: UnboundID Directory Server
Access Account: cn=Sync User,cn=Root DNs,cn=config
Base DN: dc=example,dc=com
Servers: localhost:389
```

```
(more output on the next page)
```

```
Destination: Microsoft Active Directory
Type: Microsoft Active Directory
Access Account: cn=Sync User,cn=Users,DC=adsync,DC=UnboundID,DC=com
Base DN: DC=adsync,DC=UnboundID,DC=com
Servers: 10.5.1.149:636
```

```
Sync Classes:
UnboundID Directory Server Users Sync Class
Base DN: dc=example,dc=com
Filters: (objectClass=inetOrgPerson)
DN Map: **,ou=users,dc=example,dc=com ->{1},CN=Users,DC=adsync,DC=UnboundID,DC=com
Synchronized Attributes: Custom set of mappings are defined
Operations: Creates,Deletes,Modifies

w) write configuration
b) back
q) quit

Enter choice [w]:
```

10. Next, you will be prompted as to whether you want to apply the configuration to the local Synchronization Server instance. Type "yes" to apply the configuration changes.
11. Connect to the Synchronization Server by entering the LDAP Connection Parameters: connection type (LDAP, SSL, or StartTLS), **host name**, **port**, **user bind DN** and **bind DN password**. The configuration is also recorded at `<server-root>/logs/tools/create-sync-pipe-config.log`.

Configuring the Password Encryption Component

1. The next two steps are only required if you are syncing passwords **from UnboundID Directory Server to Active Directory**. It is **not** required if you are syncing from Active Directory to UnboundID Directory Server, or have no plans to sync passwords.

On the UnboundID Directory Server that will receive the password modifications, enable the "Changelog Password Encryption" component on the directory server. The component intercepts password modifications, encrypts the password and adds an encrypted attribute, **ds-changelog-encrypted-password**, to the changelog entry. You can copy and paste the encryption key from the output if displayed, or you can access it from the `<server-root>/sync-pipe-cfg.txt`.

```
$ bin/dsconfig set-plugin-prop --plugin-name "Changelog Password Encryption" --set
enabled:true --set changelog-password-encryption-key:ej5u9e39pqo68"
```

2. This step is only required if you are syncing passwords **from UnboundID Directory Server to Active Directory**. It is **not** required if you are syncing from Active Directory to UnboundID Directory Server, or have no plans to sync passwords.

On the Synchronization Server, set the decryption key used to decrypt the user password value in the changelog entries. The key allows the user password to be synchronized to other servers that do not use the same password storage scheme.

```
$ bin/dsconfig set-global-sync-configuration-prop \
--set changelog-password-decryption-key:ej5u9e39pqo68
```

You have successfully configured the Active Directory Synchronization topology.

You can test the configuration or populate data in the Destination Servers using bulk resync mode. See "Using Bulk Resync on the Synchronization Server" on page 79. Then, you can use **realtime-sync** to start synchronizing the data. See "Controlling Real Time Synchronization" on page 84 for more information. Finally, if you are planning on synchronizing

passwords, you must install the Password Sync Agent (PSA) on all of the domain controllers in the topology. See the next chapter on how to install the PSA.

Installing the UnboundID Password Sync Agent

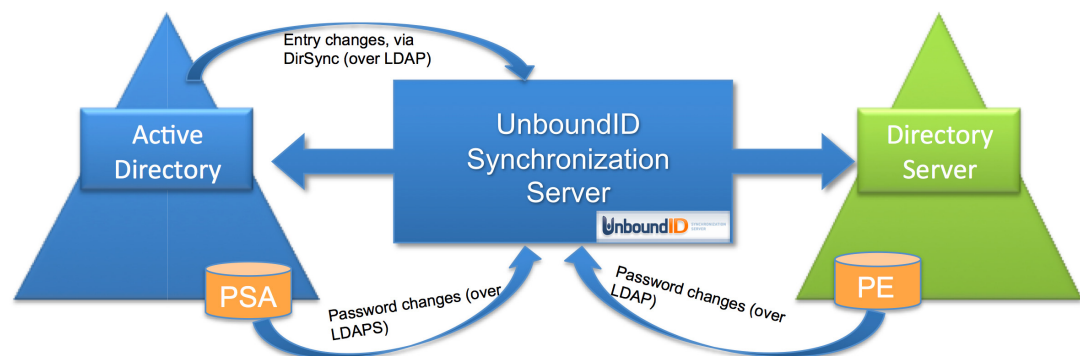
When synchronizing passwords with Active Directory systems, the UnboundID Synchronization Server requires that an additional software component, the UnboundID Password Sync Agent (PSA), be installed on all domain controllers in the synchronization topology. This component provides real-time outbound password synchronization from Microsoft Active Directory to any of the Sync Destinations supported by the UnboundID Synchronization Server. Currently these includes the UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun Directory Server 5.x, Sun DSEE (6.x, 7.x), Oracle (10g, 11g), and Microsoft SQL Server (2005, 2008).

The PSA component was designed to provide realtime password synchronization between directories that support differing password storage schemes. The PSA component immediately hashes the password with a 160-bit salted secure hash algorithm and erases the memory where the clear-text password was stored. The component only transmits data over a secure (SSL) connection. The PSA follows Microsoft's security guidelines when handling clear-text passwords (see [http://msdn.microsoft.com/en-us/library/ms721884\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms721884(VS.85).aspx)).

Note

For outbound password synchronization **from** UnboundID Directory Server **to** Active Directory, you can enable the Password Encryption component, which has similar functionality to that of the PSA component. See “Configuring the Password Encryption Component” on page 107 for more information.

FIGURE 4-1. Password Synchronization with Microsoft Active Directory Systems



The PSA also utilizes Microsoft Windows password filters, which are part of the local security authority (LSA) process. The password filters allow you to implement password policy validation and change notification mechanisms for your system. For more information on this topic, see [http://msdn.microsoft.com/en-us/library/ms721882\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms721882(VS.85).aspx).

The PSA supports failover between sync servers. It caches the hashed password changes in a local database until it can be guaranteed that all sync servers in the topology have received them. The failover features provide added flexibility as any or all of the sync servers can be taken offline or reconfigured in real-time without losing any password changes from Active Directory.

The UnboundID Password Sync Agent is safe to leave running on a domain controller indefinitely. If you want to temporarily (or permanently) stop synchronizing passwords, simply remove the `userPassword` attribute mapping in the sync server, or just stop the sync server. The PSA will not allow its local cache of password changes to grow out of control; it automatically cleans out records from its local database as soon as they have been acknowledged at a sync server, and furthermore, it also purges changes that have been in the database for over a week. This is both a safety feature as well as a maintenance convenience, so that the PSA will not require any manual updates if the sync configuration changes.

Supported Platforms

The Password Sync Agent (PSA) software has been tested and is supported for the following platforms:

- Windows Server 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2

Before You Begin

- If you have no plans to synchronize passwords, you do not need to install the PSA.
- Make sure that the Active Directory domain controller has SSL enabled and running on the Windows host machine.
- The UnboundID Synchronization Servers must be configured to accept SSL connections when communicating with the Active Directory host.
- At least one ADSyncSource needs to be configured on the UnboundID Synchronization Server and should point to the domain controller(s) on which the PSA is being installed.
- At the time of installation, all UnboundID Synchronization Servers in the sync topology must be online and available.
- The PSA component is for outbound-only password synchronization **from** the Active Directory Systems. The PSA component is not necessary if you are performing a one-way password sync **from** the UnboundID Directory Server **to** the Active Directory server.

Installing the Password Sync Agent

UnboundID distributes the Password Sync Agent (PSA) in zip file format. The PSA will be available together with each UnboundID Synchronization Server build. Before you install the PSA, ensure that your system meets the conditions presented in the previous section.

The initial (i.e., first time) installation of the PSA requires a system restart for the new PSA DLL (Microsoft requirement).

To Install the Password Sync Agent (restart needed)

1. On the domain controller, run the installer by double-clicking the "setup.exe" program.
2. Select an installation folder for the service files. This is where the PSA will store its binaries, local database, and log files.
3. Enter the host names (or IP addresses) and SSL ports of the UnboundID Synchronization Servers. They should be separated by a colon, for example, "sync.host.com:636". Do not add any prefixes to the hostnames.
4. Enter the Directory Manager DN and password. This is only used to set up a special "ADSync" user on the sync servers. You must also enter a password for this user. Since this is a first-time installation, the ADSync User password will be set to the password you supply. However, if it has been set before (by a previous installation), the password you supply will be verified against the existing password. In this case, make sure you use the same password that was used previously.
5. Enter the maximum size that you would like each log file to grow to in kilobytes. The PSA uses a simple two-file log system. When one reaches the maximum size, it is copied to the backup log file location and a new one is started.
6. Click "Next" to begin the installation. All of the specified sync servers will be contacted, and any failures will roll back the installation. If everything succeeds, you will see an information message indicating that a restart is required. At this point the PSA is installed but not running. It will not begin until the computer restarts, and the LSA process loads it into memory. Unfortunately, the LSA process cannot be restarted at runtime.
7. If you are syncing pre-encoded passwords from an Active Directory system to an UnboundID Directory system, you must allow pre-encoded passwords in the default password policy.

```
$ bin/dsconfig set-password-policy-prop --policy-name "Default Password Policy" \
--set allow-pre-encoded-passwords:true
```

Upgrading the Password Sync Agent

For software upgrades for the Password Sync Agent (PSA), the Synchronization Server provides the **update** tool that upgrades the server code to the latest version. New PSA builds are packaged with the Synchronization Server upgrade distributions. The upgrade does not require

a restart, because the core password filter is already running under LSA. The upgrade merely replaces the implementation binaries, which are encapsulated from the password filter DLL.

To upgrade the Password Sync Agent, see “Updating the Synchronization Server” on page 33.

Uninstalling the Password Sync Agent

If you are required to uninstall the PSA, you can simply use the Add/Remove Programs on the Windows Control Panel.

To Uninstall the Active Directory Password Synchronization Service (restart optional)

1. Go to Control Panel, select Add/Remove Programs. Find the Password Sync Agent and click "Remove". Click "Yes" on the warning that asks if you are sure.
2. At this stage, the PSA has been stopped, and passwords will no longer get synced to the synchronization servers or stored in the local database. The implementation DLL has been unloaded, and the database and log files are deleted. Only the binaries remain.
3. The core password filter, however, is still running under the LSA process (it cannot be unloaded at runtime). At this point, it imposes zero overhead on the domain controller, because its implementation DLL has been unloaded. If it is absolutely necessary to remove the password filter itself (located at `C:\WINDOWS\System32\ubidPassFilt.dll`), simply restart the computer. On restart, the password filter and implementation binaries (found in the install folder) can be deleted.

Note

After uninstalling the Password Sync Agent, it cannot be reinstalled without another reboot (i.e. it will revert back to a first-time installation state).
--

Manual Configuration for Advanced Users

All the configuration settings for the Password Sync Service are stored in the Windows registry under the key "`HKLM\SOFTWARE\UnboundID\PasswordSync`". If you want to manually change any of the configuration properties, you can do so at runtime by modifying the values under this registry key. The agent will automatically reload and refresh its settings from the registry. You can verify that the agent is working by checking the current log file, found in the installation directory under "`logs\password-sync-current.log`".

5

Configuring Synchronization with Relational Databases

Overview

The UnboundID Synchronization Server supports high-scale, highly-available data synchronization from one endpoint consisting of the UnboundID Directory Server, UnboundID Directory Proxy Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Server (3.x), Sun™ Directory Server Enterprise Edition (DSEE 6.x, 7.x), Sun™ Directory Server (5.2 p3 or higher), or Microsoft Active Directory systems to another endpoint consisting of a relational database management system (RDBMS). UnboundID officially supports synchronization with Oracle Database 10g and 11g as well as Microsoft SQL Server 2005 and 2008. The architecture, however, does not make any assumptions about the type of database or schema being managed; any database with a JDBC 3.0 or higher driver compatible with Java 1.6 can be used.

This chapter presents the following information:

- [About the DBSync Process](#)
- [About the Server SDK](#)
- [About the DBSync Example](#)
- [About the Overall DBSync Configuration Process](#)
- [Downloading the Software Packages](#)
- [Creating the JDBC Adapter Script](#)
- [Configuring the Database for Synchronization](#)
- [Pre-Configuration Checklist](#)
- [Configuring the Directory-to-Database Sync Pipe](#)
- [Configuring the Database-to-Directory Sync Pipe](#)

About the DBSync Process

The Synchronization Server is designed for high-scale, parallelized, point-to-point data synchronization between a directory server and a RDBMS system via a JDBC 3.0 or higher driver compatible with Java 1.6. The Synchronization Server provides multiple configuration options, such as, advanced filtering (fractional and subtree), attribute and DN mappings, trans-

formations, correlations, and configurable logging features for seamless one-way or bidirectional synchronization.

To support synchronizing changes out of a database, the database must be configured with a change tracking mechanism. We recommend a general approach involving triggers (one trigger per table) to record all changes to a changelog table. The database changelog table should record the type of change (INSERT, UPDATE, DELETE) that occurred, the specific table name, the unique identifier for the row that was changed, the database entry type, the changed columns, the modifier's name, and the timestamp of the change. An example is presented later in this chapter.

The Synchronization Server delegates the physical interaction with the database to a scripted layer, which has full control of the SQL queries. The scripted layer provides flexibility in how you define the mapping semantics between your LDAP environment and your relational database environment. The connection management, pooling, retry logic, and other boilerplate code are all handled internally by the Synchronization Server.

The RDBMS Synchronization (DBSync) implementation does not support failover between different physical database servers as is the case for directory servers. Most enterprise databases have a built-in failover layer (for example, Microsoft's node-based SQL Server Failover Cluster) from which the Sync Server can point to a single virtual address and port and still be highly available. Note that while you can have a single RDBMS node, you can scale to multiple directory server instances at the other endpoint.

About the Server SDK

To synchronize LDAP data to or from a relational database, you must first create a JDBC Adapter script to act as an interface between the UnboundID Synchronization Server and your database environment. You can create the adapter script using the UnboundID Server SDK, which provides APIs to develop plugins and third-party extensions to the server using Java or Groovy. The Server SDK's documentation (javadoc and examples) is delivered with the Server SDK build in zip format.

Note	Server SDK support is provided if you have purchased Premium Support for the product that you are developing extensions for. However, UnboundID does not provide support for the third party extensions developed using the Server SDK. You should contact your product level support group if you need assistance.
------	---

The Server SDK contains two abstract classes that will likely be required for your implementation:

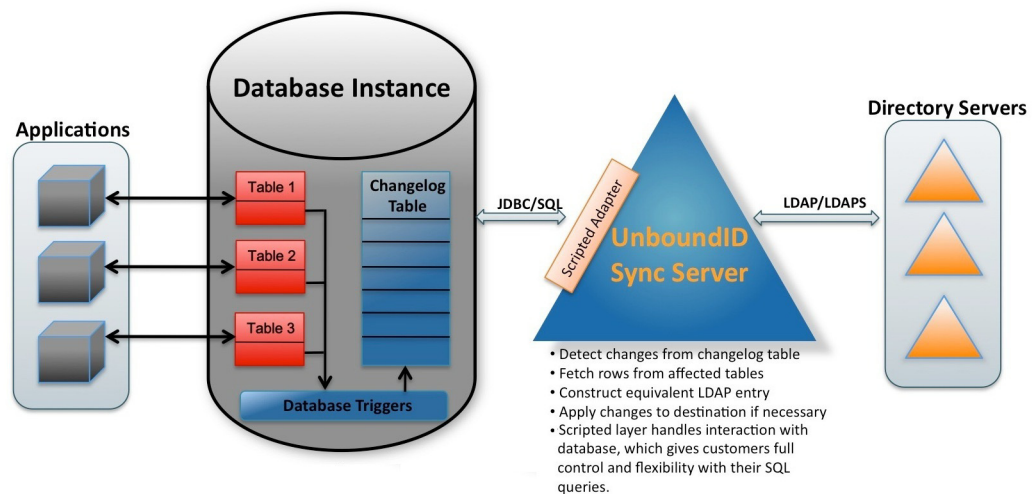
- `ScriptedJDBCSyncSource`
- `ScriptedJDBCSyncDestination`

The abstract classes correspond to how you use the database, either as a source of synchronization or as a target destination. The remainder of the SDK contains helper classes and utility functions to make the script implementation simpler.

The SDK allows you to use any change tracking mechanism to detect changes in the database. However, we provide a recommended generic approach using a simplified changelog table and triggers to record changes. Our solution is largely vendor and database version-independent and is configurable on any database that supports row-based trigger semantics. Examples are provided in the `config/jdbc/samples` directory for Oracle Database and Microsoft SQL Server.

The Synchronization Server uses the scripted adapter layer (shown in Figure 5-1) to convert any database change to an equivalent LDAP entry. From there, the Sync Pipe processes the data through inclusive (or exclusive) filtering together using attribute and DN maps defined in the Sync Classes to update the endpoint servers. For example, once you have implemented a script using Groovy, you can configure it for use by setting the `script-class` property on a `GroovyScriptedJDBCSyncSource` or `GroovyScriptedJDBCSyncDestination` configuration object within the Synchronization Server. The procedures to accomplish this are presented later in this chapter.

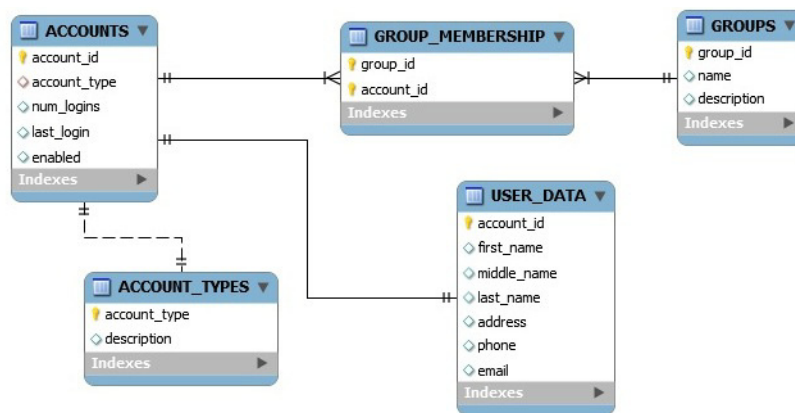
FIGURE 5-1. Architectural Overview



About the DBSync Example

The Synchronization Server provides a DBSync example between two endpoints consisting of an UnboundID Directory Server source and a RDBMS system, which will be used throughout this chapter. The Entity-Relational diagram for the normalized database schema is available in `/config/jdbc/samples/oracle-db/ComplexSchema.jpg` and is shown in Figure 5-2. Five tables are represented with their primary keys in bold. The entity relations and foreign keys are marked by the relationship lines.

FIGURE 5-2. ER Diagram for the Schema Tables



Example DS Entries

The synchronization example assumes that the directory server's schema has been configured to handle the mapped attributes. If you are configuring a database-to-directory sync pipe with a newly installed directory server, you must ensure that the schema has the correct `attributeType` and `objectClass` definitions in place. You can add the definitions in a custom `99-user.ldif` file in the `config/schema` folder of your directory server implementation, if necessary. The following snippet shows an example of the LDAP entries that are used in the synchronization example. The "accounts" table from Figure 5-2 maps to a custom object class on the directory server, while the "groups" table maps to a standard LDAP group entry with `objectclass:groupOfUniqueNames`. Example entries appear as follows:

```

dn: accountid=0,ou=People,dc=example,dc=com
objectClass: site-user
firstName: John
lastName: Smith
accountID: 1234
email: jsmith@example.com
phone: +1 556 805 4454
address: 17121 Green Street
numLogins: 4
lastLogin: 20070408182813.196Z
enabled: true

dn: cn=Group 1,ou=Groups,dc=example,dc=com
objectClass: groupOfUniqueNames
description: This is Group 1
uniqueMember: accountID=0,ou=People,dc=example,dc=com
uniqueMember: accountID=1,ou=People,dc=example,dc=com
  
```

About the Overall DBSync Configuration Process

The procedure to configure a DBSync system is slightly more complicated than that of a directory server-to-directory server synchronization configuration due to the extra tasks required to create the Adapter Scripts and to configure the database. The overall configuration process is as follows:

1. Download the appropriate JDBC 3.0 or higher driver. UnboundID cannot bundle any JDBC drivers with the Synchronization Server due to licensing restrictions, but most are freely available to end users. When ready, you can place it in the `lib` directory in the Synchronization Server's installation directory (`/UnboundID-Sync/lib`), and then restart the server for the driver to take effect.

For example, you should download the `ojdbc6.jar` file for Oracle systems or the `sqljdbc4.jar` file for MS SQL Server systems.

2. Create one or more Adapter scripts based on the Server SDK. If you are configuring for bidirectional synchronization, you will need two scripts: one for the JDBC Sync Source; the other for the JDBC Sync Destination.
3. Configure the database changelog table and triggers (presented later). While you can use the vendor's native change tracking mechanism, we recommend setting up a changelog table, shown later in the configuration procedures. Each table requires one database trigger to detect the changes and loads them into the changelog table.
4. Configure the Sync Pipes including the Sync Classes, external servers, DN and attribute maps for one direction (e.g., from directory server to database).
5. Run the `resync --dry-run` command to test the configuration settings.
6. Run `realtime-sync set-startpoint` to initialize the starting point for synchronization.
7. Run the `resync` command to populate data on the endpoint.
8. Start the Sync Pipes using the `realtime-sync start` command.
9. Monitor the synchronization server using the `status` commands and logs.
10. For bidirectional synchronization, configure another Sync Pipe in the other direction (e.g., from database to directory server), repeat steps 4–8 to test the complete synchronization system.

Downloading the Software Packages

You need to download your JDBC driver prior to setting up your synchronization environment with a RDBMS system.

To Download the Software Packages

1. Download the UnboundID Synchronization zip file. Unzip the server in a directory of your choice.
2. If you are configuring the Synchronization Server from scratch, you must ensure that you have JDK1.6 update 20. Set the `JAVA_HOME` environment variable and your `PATH` or `CLASSPATH` variables accordingly.

Download an appropriate JDBC 3.0 or higher driver for your system. Place it in the `lib` directory in the Synchronization Server's installation directory (`/UnboundID-Sync/lib`). For example, you should download the `ojdbc6.jar` file for Oracle systems or the `sqljdbc4.jar` file for MS SQL Server systems.

Note	You will need to restart the Sync Pipe to pick up changes to a script.
------	--

3. Download the Server SDK zip file and unzip it in a directory of your choice.

Creating the JDBC Adapter Script

The JDBC Adapter Script implementation must be written in Java or the Groovy scripting language (<http://groovy.codehaus.org/api/>). Consult the Server SDK documentation for details on Java implementations. The examples in this guide use Groovy to assist those unfamiliar with the language. Both languages have been tested and are fully supported by the Synchronization Server. UnboundID does not recommend one language over the other for your implementation requirements.

For those unfamiliar with Groovy, it is an open-source, dynamically-typed scripting language, similar to Java and provides quick adoption for those developers who already know the Java programming language. Groovy scripts can leverage existing Java classes and libraries to allow embedded applications within Java or as standalone scripts. The Scripted JDBC Adapter is loaded at Sync Pipe startup, which allows you to dynamically reload a script by restarting the Sync Pipe.

All scripts must live under the `/lib/groovy-scripted-extensions` directory (Java implementations using the Server SDK reside under `lib/extensions`), which may also contain other plugins built using the UnboundID Server SDK. If a script declares a package name, it must live under the corresponding folder hierarchy, just like a Java class. For example, to use a script class called `ComplexJDBCSyncSource` whose package is `com.unboundid.examples.oracle`, place it under the `/lib/groovy-scripted-extensions/com/unboundid/examples/oracle` and set the 'script-class' property on the Sync Source to `"com.unboundid.examples.oracle.Com-`

`plexJDBCSyncSource`". There are a few reference implementations provided in the `config/jdbc/samples` directory.

Note	Any changes to an existing script requires a manual Sync Pipe restart. Any configuration change automatically restarts the affected Sync Pipe.
------	--

The default libraries available on the classpath to your script implementation include:

- Groovy 1.7
(<http://groovy.codehaus.org/api/>)
- UnboundID LDAP SDK for Java 2.0.0
(<http://unboundid.com/products/ldapsdk/docs/javadoc/index.html>)
- JRE 1.6
(<http://download.oracle.com/javase/6/docs/api/>)

Logging from within a script can be done with the Server SDK's `ServerContext` abstract class. Some of `ServerContext`'s methods, such as `registerChangeListener()` or `getInternalConnection()` will not be available when running the Resync tool, because it runs independent of the running Sync Server. Any logging performed within a script during a Resync operation will appear in the `logs/tools/resync` file.

About Groovy

There are a few things to be aware of when using Groovy:

- Semicolons are optional.
- The 'return' keyword is optional.
- You can use the 'this' keyword inside static methods (which refers to this class).
- Methods and classes are public by default.
- The 'throws' clause in a method signature is not checked by the Groovy compiler, because there is no difference between checked and unchecked exceptions.
- You will not get compile errors like you would in Java for using undefined members or passing arguments of the wrong type.

Implementing a Scripted JDBC Sync Source

The `ScriptedJDBCSyncSource` abstract class must be implemented in order to synchronize data out of a relational database (e.g., for database to directory server synchronization). Since the UnboundID Synchronization Server is LDAP-centric, this abstract class allows you to take database content and convert it into LDAP entries. For more detailed information on the abstract class, consult the UnboundID Server SDK Javadoc.

The script imports classes from the Java API, UnboundID LDAP SDK for Java API, and the UnboundID Server SDK. Depending on the data, you will need to implement the following methods within your script:

- **initializeJDBCSyncSource**. Called when a Sync Pipe first starts up, or when the Resync process first starts up. Any initialization should be performed here, such as, creating internal data structures and setting up variables.
- **finalizeJDBCSyncSource**. Called when a Sync Pipe shuts down, or when the Resync process shuts down. Any clean up should be performed here, and all internal resources should be freed.
- **setStartpoint**. Sets the starting point for synchronization by identifying the starting point in the changelog. This method should cause all changes previous to the specified start point to be disregarded and only changes after that point to be returned by the `getNextBatchOfChanges` method. There are several different startpoint types (see `SetStartpointOptions` in the Server SDK), and this implementation is not required to support them all. If the specified startpoint type is unsupported, this method throws an exception (`IllegalArgumentException`). This method can be called from two different contexts:
 - When the `realtime-sync set-startpoint` command is used (the Sync Pipe is required to be stopped in this context).
 - Immediately after a connection is first established to the source server (e.g., before the first call to `getNextBatchOfChanges` method).

Important

The `RESUME_AT_SERIALIZABLE` startpoint type must be supported by your implementation, because this method is used when a Sync Pipe first starts up and loads its state from disk.

- **getStartpoint**. Gets the current value of the startpoint for change detection.
- **fetchEntry**. Returns a full source entry (in LDAP form) from the database, corresponding to the `DatabaseChangeRecord` object that is passed in. The `resync` command also uses this class to retrieve entries.
- **acknowledgeCompletedOps**. Provides a means for the Sync Server to acknowledge to the database which operations have completed processing.

Important

The internal value for the startpoint should only be updated after a sync operation is acknowledged back to this script (via this method). Otherwise it will be possible for changes to be missed when the Synchronization Server is restarted or a connection error occurs.

- **getNextBatchOfChanges**. Retrieves the next set of changes for processing. The method also provides a generic means to limit the size of the result set.
- **listAllEntries**. Used by the `resync` command to get a listing of all entries.
- **cleanupChangelog**. In general, we recommend implementing a `cleanupChangelog` method, so that the Synchronization Server can purge old records from the changelog table, based on a configurable age.

- **getResyncIncludeSourceAttrs.** Returns a set of source attribute names to synchronize when the resync tool is run. Corresponds with the `--includeSourceAttr` option. By default all source attributes are synchronized. If used as a method, it will only process MODIFY operations only and not CREATE operations.
- **getResyncExcludeSourceAttrs.** Returns a set of source attribute names to exclude from synchronization when the resync tool is run. Corresponds with the `--excludeSourceAttr` option. By default all source attributes are synchronized. If used as a method, it will only process MODIFY operations only and not CREATE operations.
- **getResyncIncludeDestinationAttrs.** Returns the set of destination attribute names to synchronize when the resync tool is run. Corresponds with the `--includeDestinationAttr` option. By default all destination attributes are synchronized. If used as a method, it will only process MODIFY operations only and not CREATE operations.
- **getResyncExcludeDestinationAttrs.** Returns the set of destination attribute names to exclude when the resync tool is run. Corresponds to the `--excludeDestinationAttr` option. By default all destination attributes are synchronized. If used as a method, it will only process MODIFY operations only and not CREATE operations.

See the `config/jdbc/samples` directory for example script implementations and the Server SDK javadoc for more detailed information on each method.

Implementing a Scripted JDBC Sync Destination

The `ScriptedJDBCSyncDestination` abstract class must be implemented in order to synchronize data into a relational database (e.g., for directory server to database synchronization). The abstract class allows you to take LDAP content and convert it to database content.

The script imports classes from the Java API, UnboundID LDAP SDK for Java API, and the UnboundID Server SDK, depending on your database configuration. You will need to implement the following methods within your script:

- **initializeJDBCSyncDestination.** Called when a Sync Pipe first starts up, or when the Resync process first starts up. Any initialization should be performed here, such as, creating internal data structures and setting up variables.
- **finalizeJDBCSyncDestination.** Called when a Sync Pipe shuts down, or when the Resync process shuts down. Any clean up should be performed here, and all internal resources should be freed.
- **createEntry.** Creates a full database entry (or row), corresponding to the the LDAP Entry that is passed in.
- **modifyEntry.** Modify a database entry, corresponding to the the LDAP Entry that is passed in.
- **fetchEntry.** Return a full destination database entry (in LDAP form), corresponding to the the source entry that is passed in.

- **deleteEntry**. Delete a full entry from the database, corresponding to the the LDAP Entry that is passed in.

For more detailed information on the abstract class, consult the Server SDK Javadoc.

Configuring the Database for Synchronization

To configure the database for synchronization, you must do three things: 1) set up a database **SyncUser** account; 2) set up the change tracking mechanism; and 3) set up the database triggers (one per table) for your application. To view an example, the full example setup script is available in `/config/jdbc/samples/oracle-db/OracleSyncSetup.sql`, where items in brackets (for example `[ubid_changelog]`) is a user-named label for the account, table or column.

Note	Database change tracking is only necessary if you are syncing FROM the database. If you are syncing TO a database, you only need to set up the SyncUser account and the correct privileges.
------	---

To Configure the Database for Synchronization

1. Create an Oracle login (SyncUser) for the Synchronization Server, so that the Sync Server can access the database server. Also make sure to grant sufficient privileges to the SyncUser for any tables to be synchronized. Make sure to change the default password on production systems.

```
CREATE USER SyncUser IDENTIFIED BY password DEFAULT TABLESPACE users TEMPORARY
TABLESPACE temp;
GRANT "RESOURCE" TO SyncUser;
GRANT "CONNECT" TO SyncUser;
```

2. Set up your changelog tables on the database. An example is presented as follows:

```
CREATE TABLE ubid_changelog (
  --This is the unique number for the change
  change_number    Number NOT NULL PRIMARY KEY,

  --This is the type of change (insert, update, delete). NOTE: This should
  --represent the actual type of change that needs to happen on the
  --destination(for example a database delete might translate to a LDAP
  --modify, etc.)
  change_type      VARCHAR2(10) NOT NULL,

  --This is the name of the table that was changed
  table_name       VARCHAR(50) NOT NULL,

  --This is the unique identifier for the row that was changed. It is up to
  --the trigger code to construct this, but it should follow a DN-like format
  --(e.g. accountID={accountID}) where at least the primary key(s) are
  --present. If multiple primary keys are required, they should be delimited
  --with a unique string, such as, '%%' (e.g. accountID={accoun-
  tID}%%groupID={groupID})
  identifier       VARCHAR2(100) NOT NULL,
```

```
--This is the database entry type. The allowable values for this must be
--set on the JDBC Sync Source configuration within the Synchronization
--Server.
entry_type          VARCHAR2(50) NOT NULL,

--This is a comma-separated list of columns that were updated as part of
--this change.
changed_columns     VARCHAR2(1000) NULL,

--This is the name of the database user who made the change
modifiers_name     VARCHAR2(50) NOT NULL,

--This is the timestamp of the change
change_time        TIMESTAMP(3) NOT NULL,

CONSTRAINT chk_change_type CHECK (change_type IN ('insert','update','delete'))
ORGANIZATION INDEX;
```

3. Create an Oracle function to get the SyncUser name. This is a convenience function for the triggers.

```
CREATE OR REPLACE FUNCTION get_sync_user RETURN VARCHAR2
IS
BEGIN
    RETURN 'SyncUser';
END get_sync_user;
```

4. Create an Oracle sequence object for the change-number column in the changelog table.

```
CREATE SEQUENCE ubid_changelog_seq MINVALUE 1 START WITH 1 NOMAXVALUE INCREMENT BY 1
CACHE 100 NOCYCLE;
```

5. Create a Database Trigger for each table that will participate in synchronization. An example is shown below and shows a trigger for the Accounts table that tracks all changed columns after any INSERT, UPDATE, and DELETE operation. The code generates a list of changed items and then inserts them into the changelog table. See the example in `/config/jdbc/samples/oracle-db/OracleSyncSetup.sql`.

```
CREATE OR REPLACE TRIGGER ubid_accounts_trg AFTER INSERT OR DELETE OR UPDATE ON
accounts
FOR EACH ROW
DECLARE
    my_identifier ubid_changelog.identifier%TYPE;
    my_changetype ubid_changelog.change_type%TYPE;
    my_changedcolumns ubid_changelog.changed_columns%TYPE := '';
    CURSOR column_cursor IS select COLUMN_NAME from USER_TAB_COLUMNS where
TABLE_NAME='ACCOUNTS';
BEGIN
    --Short circuit and do nothing if the change came from the Synchronization Server
    itself.
    --This prevents loopbacks when doing bidirectional synchronization.

    IF UPPER(USER) = UPPER(get_sync_user()) THEN
        RETURN;
    END IF;

    -- Figure out change type
    IF INSERTING THEN
        my_identifier := 'account_id' || :NEW.account_id;
        my_changetype := 'insert';
    ELSIF DELETING THEN
```

```
my_identifier := 'account_id=' || :OLD.account_id;
my_changetype := 'delete';
ELSIF UPDATING THEN
my_identifier := 'account_id=' || :NEW.account_id;
my_changetype := 'update';

--Figure out changed columns
FOR my_row IN column_cursor

LOOP
    IF UPDATING (my_row.COLUMN_NAME) THEN
        my_changedcolumns := my_changedcolumns || my_row.COLUMN_NAME || ',';
    END IF;
END LOOP;
END IF;

--Do the insert
INSERT INTO ubid_changelog (change_number, change_type, table_name, identifier,
entry_type,
changed_columns, modifiers_name, change_time) VALUES (ubid_changelog_seq.NEXTVAL,
my_changetype, 'ACCOUNTS', my_identifier, 'account', my_changedcolumns, USER, SYS-
TIMESTAMP);

--If changes to this table affect multiple LDAP entries, multiple records should
--be inserted into the changelog table. For example, if an update to an "account" in
--the database affected an "account" LDAP entry and a "groups" LDAP entry, then we
--would have another "INSERT INTO ubid_changelog..." here with a different entry
--type.

EXCEPTION
WHEN OTHERS THEN
    DBMS_OUTPUT.PUT_LINE('Changelog trigger exception:');
    DBMS_OUTPUT.PUT_LINE(DBMS_UTILITY.FORMAT_ERROR_BACKTRACE);
END;
```

Pre-Configuration Checklist

Before configuring the Synchronization Server, we assume that you have accomplished the following items:

- Create a Sync User account with the access privileges to the RDBMS server, so that the Sync Server can access the machine.
- Set up your changelog tables on the database.
- Set up the triggers in the database: one per table that will participate in synchronization.
- Create a JDBC Adapter Script using Java or Groovy to map the LDAP Entries to the RDBMS table rows or vice-versa, place it in the `/lib/groovy-scripted-extensions` directory (Java implementations using the Server SDK reside under `lib/extensions`). You must also create a script if configuring a Sync Pipe from database to directory server. See the example in the `config/jdbc/samples` directory.
- Make sure the database is configured to listen for external connections. If it is locked down for security, the Sync Server will register a connection error.

Configuring the Directory-to-Database Sync Pipe

The following procedure shows the interactive steps to set up a one-way Sync Pipe with an UnboundID Directory Server as the Sync Source and a RDBMS (Oracle) system as the Sync Destination. The procedure uses the `create-sync-pipe-config` tool in interactive command-line mode, which shows the configuration steps in a top-down flow from Sync Pipe to external servers.

The procedure is broken out into sections for easy access and is based on the interactive prompts that the `create-sync-pipe-config` tool will present. The instructions assume that the user has the proper root user or admin privileges to make configuration changes. Once you have configured the sync pipes, then you can fine-tune the configuration later using the `dsconfig` utility.

General Tips When Syncing to a Database Destination

When configuring a directory-to-database Sync Pipe, you should be aware of the following recommendations:

- **Identify the Object Classes.** It is advisable to identify the different object classes that will be synced. Create a Sync Class per object class, so that you can easily distinguish between them and have different mappings and sync rules set up for each one.
- **For each Sync Class**, make sure to set the following items:
 - **Set the Include-Filter Property.** Make sure the `include-filter` property is set to something that will uniquely identify the source entries, such as, "objectClass=<object class for this Sync Class>".
 - **Create Specific Attribute Mappings.** Create a specific Attribute Mapping for every LDAP attribute that you want to be synced to a database column(s); add all of these to a single Attribute Map and set it on the Sync Class. This way, the script will not have to know about the schema on the directory side. It may be desirable to add a Constructed Attribute Mapping that maps a literal value to the objectClass attribute, if needed by the script, to determine the "database entry type." For example, you could have something like "account" -> objectClass, which would result in the constructed destination LDAP entry always containing an objectClass of "account".
 - **Create Specific DN Maps (optional).** Create a DN Map that recognizes the DN's of the source entries and maps them to a desired destination DN. In most cases, this step is unnecessary, because the script will use the attributes rather than the DN to figure out which database entry needs changed.
 - **Set Auto-Mapped-Source-Attrs to "-none-".** Remove the default value of "-all-" from "auto-mapped-source-attrs" on the Sync Class, and replace it with the value "-none-". We do not want any values from the directory automatically mapped to an attribute with the same name when using explicit attribute mappings.

- **Avoid bidirectional Loopback.** Make sure to set the `ignore-changes-by-[user|dn]` property on both Sync Sources when configuring for bidirectional synchronization. This is important to make sure that changes are not looped back by the Sync Server.
- **Synchronizing DELETE Operations.** On UnboundID Directory Server and Alcatel-Lucent 8661 Directory Server systems, you must configure the `changelog-deleted-entry-include-attribute` property on the changelog backend. This property allows for the proper synchronization of DELETE operations that occur with this endpoint server. For the example presented in this section, you would set the `changelog-deleted-entry-include-attribute=accountid`. For more information, see “Configuring the Directory Server Backend for Synchronizing Deletes” on page 91.
- **Set the Attribute-Synchronization-Mode Appropriately for DB Sync.** For MODIFY operations, the Synchronization Server detects any change on the source changelog, fetches the source entry, applies mappings, computes the equivalent destination entry, fetches the actual destination entry, and then runs a diff between the two entries to determine the minimal set of changes to get the destination in sync with the source. By default, the Synchronization Server only makes changes on the destination entry for those attributes that were detected in the original changelog entry. However, this is configurable using the `attribute-synchronization-mode` property. The `attribute-synchronization-mode` property sets the type of diff operation that is performed between the source and destination entries on a MODIFY operation, which in turn determines the scope of attributes that are modified on the destination.

If the source endpoint is a database server (Oracle or MS SQL Server), we recommend setting the `attribute-synchronization-mode` property to `all-attributes`. In this way, the diff operation will consider all the source attributes and any that have changed will be updated on the destination, even if the change was not originally detected in the changelog. In some cases, you may not be able to get a list of changed columns in the database, in which case, you would have to use this mode, because `modified-attributes-only` will not change any destination attributes if it thinks that there are no source attributes changed. If both endpoints are directory servers, we recommend keeping the default configuration to `modified-attributes-only` to avoid any possible replication conflicts.

Step 1. Creating the Directory-to-Database Sync Pipe

The initial configuration steps show how to set up a single Sync Pipe from a directory server instance to a database using the `create-sync-pipe-config` tool in interactive mode. The `create-sync-pipe-config` tool prompts the user for input and leads you through the configuration steps in a wizard-like mode. The procedure will show how to set up and configure the Sync Pipe, external servers, and Sync Classes. The examples are based on the Complex JDBC sample in the `config/jdbc/samples/oracle-db` directory.

Optionally, you can run the `create-sync-pipe-config` tool with the server offline and import the configuration later.

To Create a Directory-to-Database Sync Pipe

1. Start the Synchronization Server.


```
$ bin/start-sync-server
```

2. Run the `create-sync-pipe-config` tool.

```
$ bin/create-sync-pipe-config
```

3. At the Initial Synchronization Configuration Tool prompt, press Enter to continue.
4. On the Synchronization Directory menu, enter 1 to create a one-way sync pipe from directory to database. If you are planning to deploy a bidirectional Sync configuration, enter 2 for bidirectional synchronization.

To Configure the Sync Source

1. On the Source Endpoint Type menu, enter the number for the sync source corresponding to the type of source external server. For this example, enter 1 to select the **UnboundID Directory Server**.
2. Next, you will be prompted to enter a name for the Source Endpoint. Enter a descriptive name for the Sync Source. For example, `dssync`.
3. Next, enter the base DN for the directory server, which is used as the base for LDAP searches. For example, enter `dc=example,dc=com`, and then press Enter again to return to the menu. If you enter more than one base DN, make sure the DNs do not overlap.
4. On the Server Security menu, select the type of secure communication that the Synchronization Server will use with the endpoint server instances. Select either 1) None; 2) SSL; or 3) StartTLS. For this example, select the default (None).
5. Next, enter the host and port of the first Source Endpoint server. The Sync Source can specify a single server or multiple servers in a replicated topology. The Synchronization Server will contact this first server if it is available, then contact the next highest priority server if the first server is unavailable, etc. After you have entered the host and port, the Synchronization Server tests that a connection can be established.
6. On the Synchronization Server User Account, enter the DN of the sync user account and create a password for this account. The Sync User account allows the Synchronization Server to access the source endpoint server. By default, the Sync User account is placed at `cn=Sync User,cn=Root DNs,cn=config`. Press Enter to accept the default configuration.

To Configure the Destination Endpoint Server

1. Next, on the Destination Endpoint Type menu, select the type of datastore on the endpoint server. In this example, enter 8 for Oracle Database.
2. Next, you will be prompted to enter a name for the Destination Endpoint. Enter a descriptive name for the Sync Destination. For example, `oraclesync`.

3. On the JDBC Endpoint Connection Parameters menu, enter the fully-qualified and resolvable host name or IP address for the Oracle database server. After you have entered the host name, the Synchronization Server checks if the hostname or IP address is resolvable.
4. Next, enter the listener port for the database server. For this example, press Enter to accept the default (1521).
5. You will be prompted to enter a database name. For this example, use `dbsync-test`.
6. Next, the Synchronization Server attempts to locate the JDBC driver in the `lib` directory. If the server found the file, it will generate a success message.

```
Successfully found and loaded JDBC driver for: jdbc:oracle:thin:@//dbsync-w2k8-vm-2:1521/dbsync-test
```

If the server cannot find the JDBC driver, you can add it later, or quit the `create-sync-pipe-config` tool and add the file to the `lib` directory. The following message is displayed to std-out.

```
Could not find an appropriate JDBC driver in the /UnboundID-Sync/lib directory for: jdbc:oracle:thin:@//dbsync-w2k8-vm-2:1521/dbsync-test
```

```
What do you want to do?
```

- 1) I will add the JDBC driver later
- 2) Quit this tool and add the JDBC driver now

- b) back
- q) quit

```
Choose an option [1]:
```

7. Next, you will be prompted if you want to add any additional JDBC connection properties for the database server. Please consult your JDBC driver's vendor documentation to see what properties are supported. For this example, press Enter to accept the default (no).
8. Next, you will be prompted to enter a name for the database user account with which the Synchronization Server will communicate. Press Enter to accept the default (SyncUser). Then, enter the password for the SyncUser account. For information on creating the SyncUser account on the Oracle Server, see step 1 in "To Configure the Database for Synchronization" on page 122.
9. At this stage, you will be prompted to enter the fully qualified name of the scripted JDBC adapter that uses the `ScriptedJDBCSyncDestination` abstract class. The script should reside in the `/lib/groovy-scripted-extensions` directory. For example, specify the fully-qualified class name for the script that resides in `/lib/groovy-scripted-extensions/com/unboundid/examples/oracle`.


```
Enter the fully qualified name of the Groovy class that will implement ScriptedJDBCSyncDestination: com.unboundid.examples.oracle.ComplexJDBCSyncDestination
```
10. Next, the Synchronization Server prompts if you want to prepare the Source Endpoint server, which tests the connection to the directory server and tests that the Sync User account is accessible. Press Enter to accept the default (**yes**). For the Sync User account, it

will return "Denied" as the account has not been written yet to the Directory Server at this time.

```
Testing connection to server1.example.com:1389 ..... Done
Testing 'cn=Sync User,cn=Root DNs,cn=config' access ..... Denied
```

11. Next, you will be prompted if you want to configure the Sync User account on the directory server. Press Enter to accept the default (**yes**). You will be prompted for the bind DN (e.g., `cn=Directory Manager`) and the bind DN password of the directory server so that you can configure the `cn=Sync User` account. The Sync Server creates the Sync User account, tests the base DN, and enables the changelog.

```
Created 'cn=Sync User,cn=Root DNs,cn=config'
Verifying base DN 'dc=example,dc=com' ..... Done
Enabling cn=changelog .....
```

12. Next, you will be prompted to enter the maximum age of the changelog entries. For this example, press Enter to accept the default (2h).

Step 2. Configuring the Sync Pipe and Sync Classes

In this section, we define the Sync Pipe and then create two Sync Classes. The first Sync Class is used to match the "accounts" objects. The second Sync Class is used to match the "group" objects. We'll set the basic Sync Class definitions and then add the attribute and DN maps in a later step.

To Configure the Sync Pipe and Sync Classes

1. Continuing from the previous session, enter a name for the Sync Pipe. Make sure the name is descriptive to identify it if you have more than one sync pipe configured. For example, enter `dssync to oraclesync`.
2. Next, you will be prompted if you would like to define one or more Sync Classes. Type **yes**. We'll define the Accounts Sync Class, and then the Groups Sync Class in the next sections.

To Configure the Accounts Sync Class

1. Next, enter a name for the Sync Class. Make sure the name is descriptive to identify the sync class. For example, type `accounts_sync_class`.
2. At this stage, if you plan to restrict entries to specific subtrees, then enter one or more base DNs. For this example, press Enter to accept the default (**no**).
3. Next, you will be prompted to set an LDAP search filter. For this example, type **yes** to set up a filter and enter the filter `"(accountid=*)"`. Press Enter again to continue. This property sets the LDAP filters and returns all entries that match the search criteria to be included in the Sync Class. In this example, we want to specify that any entry with an `accountID` attri-

bute be included in the Sync Class. If the entry does not contain any of these values, it will not be synchronized to the target server.

4. Continuing from the previous example, on the Sync Class menu, you will be prompted if you want to synchronize all attributes, specific attributes, or exclude specific attributes from synchronization. Press Enter to accept the default (all). We'll adjust these mappings in a later section.
5. Next, specify the operations that will be synchronized for the Sync Class. For this example, press Enter to accept the default (1, 2, 3) for creates, deletes, modifies.

To Configure the Groups Sync Class

For this current example, we need to configure another Sync Class to handle the Groups objectclass. The procedures are similar to that of the configuration steps for the `account_sync_class` Sync Class that were presented in the previous section.

1. On the Sync Class Management menu, enter a name for a new sync class. In this example, enter `groups_sync_class`.
2. At this stage, if you plan to restrict entries to specific subtrees, then enter one or more base DNs. Enter one or more base DNs. For this example, type `no`.
3. Next, you will be prompted to set an LDAP search filter. For this example, type `yes` to set up a filter and enter the filter "`(objectClass=groupOfUniqueNames)`". Press Enter again to continue. This property sets the LDAP filters and returns all entries that match the `groupOfUniqueNames` attribute to be included in the Sync Class. If the entry does not contain any of these values, it will not be synchronized to the target server.
4. Continuing from the previous example, on the Sync Class menu, you will be prompted if you want to synchronize all attributes, specific attributes, or exclude specific attributes from synchronization. Press Enter to accept the default (all). We'll adjust these mappings in a later section.
5. Next, specify the operations that will be synchronized for the Sync Class. For this example, press Enter to accept the default (1, 2, 3) for creates, deletes, modifies.
6. At this point, you will see the Sync Class menu again asking you to enter the name of another Sync Class. Press Enter to continue.
7. Next, on the Default Sync Class Operations menu, press Enter to accept the default (1,2,3) for creates, deletes, and modifies. The Default Sync Class determines how all entries that do not match any other Sync Class are handled, including whether create, delete, and/or modify operations are synchronized.
8. Review the configuration, and then press Enter to write the configuration to the Synchronization Server. If you want to change any property, you can go back to the particular menu, or make the adjustments later using the `dsconfig` tool. If you decide to write the configura-

tion to the Synchronization Server, press Enter, and then enter the connection properties for your Synchronization Server (bind DN, bind DN password).

>>>> Configuration Summary

Sync Pipe: dssync to oraclesync

Source: dssync
 Type: UnboundID Directory Server
 Access Account: cn=Sync User,cn=Root DNs,cn=config
 Base DN: dc=example,dc=com
 Servers: server1.example.com:1389

Destination: oraclesync
 Type: Oracle Database
 Access Account: SyncUser
 Servers: dbsync-w2k8-vm-2.unboundid.lab:1521

Sync Classes:
 accounts_sync_class
 Base DN:
 Filters: (accountID=*)
 DN Map: None
 Synchronized Attributes: -none-
 Operations: Creates,Deletes,Modifies

 groups_sync_class
 Base DN:
 Filters: (objectClass=groupOfUniqueNames)
 DN Map: None
 Synchronized Attributes: -none-
 Operations: Creates,Deletes,Modifies

DEFAULT
 Operations: Creates,Deletes,Modifies

w) write configuration
 b) back
 q) quit

Enter choice [w]:

9. The create-sync-pipe-config tool outputs the following messages. If you have to make any manual changes to the external servers, it will present them.

```
Creating External Servers ..... Done
Creating Endpoints ..... Done
Creating Sync Pipes ..... Done
Creating Attribute and DN Mappings ..... Done
Creating Sync Classes ..... Done
```

The following issues should be resolved before starting synchronization:

Server 'dbsync-w2k8-vm-2:1521' needs manual preparation before starting synchronization.

* You need to manually create the 'SyncUser' user account on this server and grant the proper privileges.

You need to implement the following scripted adapter(s): com.unboundid.examples.samples.ComplexJDBCSyncDestination.

Refer to the product documentation for a recommended approach for initially bringing the two ends points into sync. Once this is done, you can enable real-time synchro-

```
nization using the 'realtime-sync' tool.
```

```
Press RETURN to continue
```

```
See /UnboundID-Sync/logs/tools/create-sync-pipe-config.log for a detailed log of  
this operation
```

Step 3. Fine-tuning the Sync Classes

The Accounts and Groups Sync Classes require more fine-tuning as the DN and attributes maps need to be configured. Some additional properties are required for the example presented in this chapter.

To Configure the Accounts Sync Class

1. Start the `dsconfig` tool. Then, enter or select the LDAP (or LDAPS) connection parameters for the Synchronization Server.

```
$ bin/dsconfig
```

2. On the UnboundID Synchronization Server configuration console main menu, enter the number corresponding to Sync Class. On the Standard Objects menu, enter **14** for Sync Class.
3. On the Sync Class Management menu, type **3** to view and edit an existing Sync Class.
4. Select or confirm that you are configuring a given Sync Pipe. Press Enter to continue.
5. Next, select the specific Sync Class that you want to modify. For this example, enter **1** for the `accounts` sync class.

```
>>>> Select the Sync Class from the following list:
```

- ```
1) accounts_sync_class
2) DEFAULT
3) groups_sync_class

b) back
q) quit
```

```
Enter choice [b]: 1
```

6. On the Sync Class Properties menu, enter **1** to change the description property. For this example, enter "This Sync Class matches the site-user, guest, and administrator object-Classes." This step is optional but if you configure more than one Sync Class, you should add a general description describing the sync class's purpose.

```
>>>> Configure the properties of the Sync Class
```

```
>>>> via creating 'account_sync_class' Sync Class
```

- ```
1) description
2) evaluation-order-index
3) include-base-dn
```

```
"This Sync Class matches the site-user,  
guest, and administrator objectclasses."  
10  
The location of the entry in the Sync  
Source is not taken into account when  
determining whether an entry is  
part of this Sync Class.
```

```

4) include-filter (accountID=*)
5) attribute-map No attribute map is used.
6) dn-map No dn map is used.
7) auto-mapped-source-attribute all
8) excluded-auto-mapped-source-attributes No source attributes are excluded from
synchronization.
9) destination-correlation-attributes dn
10) synchronize-creates true
11) synchronize-modifies true
12) synchronize-deletes true

?) help
f) finish - create the new Sync Class
a) show advanced properties of the Sync Class
d) display the equivalent dsconfig arguments to create this object
b) back
q) quit

Enter choice [b]:

```

To Configure an Attribute Map

1. On the Sync Class Property menu, enter 5 to set the attribute map. On the Attribute Map Property menu, enter 2 to add one or more values, and then, enter 1 to create a new attribute map.
2. Next, enter a name for the Attribute Map. Make sure the name is descriptive as you can typically have more than one attribute map in a Sync Class. For this example, enter **Directory to DB Attr Map**. Review the configuration on the Attribute Map Properties menu, and then enter **f** to save the configuration. We'll add the attribute mappings in a later section.

```

>>>> Configure the properties of the Attribute Map
>>>> via creating 'Directory to DB Attr Map' Attribute Map

Property      Value(s)
-----
1) description -

?) help
f) finish - create the new Attribute Map
d) display the equivalent dsconfig arguments to create this object
b) back
q) quit

Enter choice [b]: f

```

To Configure a DN Map

Next, we set up a DN Map from DNs in the form of ***,ou=People,dc=example,dc=com** and map it to a column/row value of **"accountid={accountid}"** in the database using the **dsconfig** command.

1. On the Sync Class Property menu, enter 6 to set the **dn-map** property. On the DN Map Property menu, enter 2 to add one or more values. Since there are no existing maps, enter 1 to create a new DN Map. Enter a name for the DN Map. For this example, enter **ubid_to_oracle_accounts_dn_map**. Review the configuration on the DN Map Properties menu, and then enter **f** to save the configuration.

- Next, enter the name of the `from-dn-pattern` property on the source directory server. For example, enter `"*,ou=People,dc=example,dc=com."`
- Next, enter the name of the `to-dn-pattern` property to which it will be mapped to the destination database server. For example, enter `"accountid={accountid}."`
- On the DN Map Property menu, review the configuration, and then enter `f` to save and apply the changes.

```
>>>> Configure the properties of the DN Map
>>>> via creating 'ubid_to_oracle_accounts_dn_map' DN Map
>>>> via creating 'account_sync_class' Sync Class

Property      Value(s)
-----
1) description -
2) from-dn-pattern  "*,ou=People,dc=example,dc=com"
3) to-dn-pattern    accountid={accountid}

?) help
f) finish - create the new DN Map
d) display the equivalent dsconfig arguments to create this object
b) back
q) quit

Enter choice [b]:f
```

- On the DN Map Property menu, press `Enter` to use the value (`ubid_to_oracle_accounts_dn_map`) that you just entered.

To Configure the Ignore-Zero-Length-Values Property

- On the Sync Class Property menu, type `a` to show the advanced properties. Then, enter `15` to change the `ignore-zero-length-values` property. This property ignores attribute changes that result in an empty (zero-length) value. Set the value to `TRUE`.
- On the Sync Class Property menu, review the configuration, and type `f` to save and apply the changes. The advanced properties menu is displayed.

```
>>>> Configure the properties of the Sync Class
>>>> via creating 'account_sync_class' Sync Class

1) description "This Sync Class matches the site-user,
2) evaluation-order-index guest, and administrator objectclasses."
3) include-base-dn 5
The location of the entry in the Sync
Source is not taken into account when
determining whether an entry is
part of this Sync Class.
(accountID=*)
Directory to DB Attr Map
ubid_to_oracle_accounts_dn_map
-none-
No source attributes are excluded from
synchronization.
account_id
-
true
true
true
all-attributes
true
true
false
false

4) include-filter
5) attribute-map
6) dn-map
7) auto-mapped-source-attribute
8) excluded-auto-mapped-source-attributes

9) destination-correlation-attributes
10) destination-correlation-attributes-on-delete
11) synchronize-creates
12) synchronize-modifies
13) synchronize-deletes
14) attribute-synchronization-mode
15) ignore-zero-length-values
16) replace-all-attr-values
17) modifies-as-creates
18) creates-as-modifies

?) help
f) finish - create the new Sync Class
a) show advanced properties of the Sync Class
d) display the equivalent dsconfig arguments to create this object
```



```

b)    back
q)    quit

Enter choice [b]: f

```

You have successfully configured the `account_sync_class` Sync Class.

To Configure the the "Groups Sync Class" Sync Class

For this current example, we need to configure another Sync Class to handle the Groups objectclass. The procedures are similar to that of the configuration steps for the `account_sync_class` Sync Class.

1. On the Sync Class Management menu, enter 3 to view and edit an existing sync class, and then select `groups_sync_class`.

2. On the Sync Class Properties menu, configure the following properties:

- Set the `description` property to: "This Sync Class matches the Groups objectclass."
- Create and set the attribute map to: `Directory to DB Groups Map`
- Create and set the DN map to: `ubid_to_oracle_groups_dn_map`. The equivalent `dsconfig` command is as follows:

```
dsconfig create-dn-map --map-name ubid_to_oracle_groups_dn_map \
--set "from-dn-pattern:*" --set "to-dn-pattern:name={cn}"
```

- Set the `ignore-zero-length-values` property to: `true`

3. The specific property values for the Groups Sync Class can be seen below. When finished, review the configuration, and then enter `f` to save and apply the changes:

```

>>>> Configure the properties of the Sync Class
>>>> via creating 'Groups Sync Class' Sync Class

```

	Property	Value(s)
1)	<code>description</code>	This Sync Class matches the Groups objectclass.
2)	<code>evaluation-order-index</code>	10
3)	<code>include-base-dn</code>	The location of the entry in the Sync Source is not taken into account when determining whether an entry is part of this Sync Class.
4)	<code>include-filter</code>	(objectClass=groupOfUniqueNames)
5)	<code>attribute-map</code>	Directory to DB Groups Map
6)	<code>dn-map</code>	ubid_to_oracle_groups_dn_map
7)	<code>auto-mapped-source-attribute</code>	-none-
8)	<code>excluded-auto-mapped-source-attributes</code>	No source attributes are excluded from synchronization.
9)	<code>destination-correlation-attributes</code>	dn
10)	<code>destination-correlation-attributes-on-delete</code>	-
11)	<code>synchronize-creates</code>	true
12)	<code>synchronize-modifies</code>	true
13)	<code>synchronize-deletes</code>	true
14)	<code>ignore-zero-length-values</code>	true
15)	<code>replace-all-attr-values</code>	true
16)	<code>modifies-as-creates</code>	false
17)	<code>creates-as-modifies</code>	false
?)	<code>help</code>	
f)	<code>finish</code>	- create the new Sync Class
a)	<code>hide advanced properties of the Sync Class</code>	
d)	<code>display the equivalent dsconfig arguments to create this object</code>	
b)	<code>back</code>	
q)	<code>quit</code>	

Enter choice [b]: f

4. On the Sync Class Management menu, enter b to back out of this menu to return to the UnboundID Synchronization Server configuration console main menu.

Step 4. Configuring the Attribute Mappings

In a previous step, the attribute maps were configured and added to each Sync Class (see “To Configure an Attribute Map” on page 133). Attribute maps are containers for attribute mappings that map the source attributes to similar or other attributes in the destination server. Based on the example schema, we want to configure the following Accounts and Group Table attributes on the system as follows:

TABLE 5-1. Attribute Mappings to Synchronize the Accounts Table

from-attribute (DS)	to-attribute (DB)
accountID	account_id
address	address
email	email
firstName	first_name
lastName	last_name
lastLogin	last_login
middleName	middle_name
numLogins	num_logins
phone	phone

TABLE 5-2. Attribute Mappings to Synchronize the Group Table

from-attribute (DS)	to-attribute (DB)
cn	name
description	description
* uniqueMember	memberID

* = This is a DN attribute mapping

To Create the Attribute Mappings

1. On the UnboundID Synchronization Server configuration console main menu, enter the number corresponding to Attribute Mapping. On the Basic objects menu, enter 2 for Attribute Mapping.
2. On the Attribute Map Management menu, enter 2 to create a new attribute mapping.
3. Select the Attribute Map that will be the container for this attribute mapping. For this example, enter 1 for the **Directory to DB Attr Map**.

```
>>>> Select the Attribute Map from the following list:
```

- 1) Directory to DB Attr Map
- 2) Directory to DB Groups Map
- b) back
- q) quit

```
Enter choice [b]: 1
```

4. Next, select the type of attribute mapping that you want to create. In this example, enter 2 for Direct Attribute Mapping.
5. Next, enter the name of the "to-attribute" to which the entry's attribute will be mapped on the destination database server. For this example, enter "account_id".
6. Next, enter the name of the "from-attribute" from which it will be mapped to the "to-attribute" on the source directory server. For example, enter "accountID".
7. On the Direct Attribute Mapping Properties menu, review the configuration, and then type f to save the changes.

```
>>>> Configure the properties of the Direct Attribute Mapping
>>>> via creating 'account_id' Direct Attribute Mapping
```

	Property	Value(s)
1)	to-attribute	account_id
2)	description	-
3)	from-attribute	accountID
?)	help	
f)	finish - create the new Direct Attribute Mapping	
a)	show advanced properties of the Direct Attribute Mapping	
d)	display the equivalent dsconfig arguments to create this object	
b)	back	
q)	quit	

```
Enter choice [b]: f
```

8. Repeat steps 2–7 for the other attribute mappings.

Or, you can use the `dsconfig` batch file feature to configure the attribute mappings at one time. Quit the `dsconfig` interactive session, create a text file, copy-and-paste the following `dsconfig` commands in the file, save the file as "attr-mappings.txt." Run the `dsconfig` command using the `-F` (or `--batch-file`) option. You must also use the `--no-prompt` option with the command.

```
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name address --type direct --set from-attribute:address
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name email --type direct --set from-attribute:email
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name first_name --type direct --set from-attribute:firstName
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name last_name --type direct --set from-attribute:lastName
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name last_login --type direct --set from-attribute:lastLogin
... (continued on next page)
```

```
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name middle_name --type direct --set from-attribute:middleName
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name num_logins --type direct --set from-attribute:numLogins
dsconfig create-attribute-mapping --map-name "Directory to DB Attr Map" \
--mapping-name phone --type direct --set from-attribute:phone

# Create the Group Attribute Mappings and assign them to the "Directory to DB Groups
Map"
dsconfig create-attribute-mapping --map-name "Directory to DB Groups Map" \
--mapping-name description --type direct --set from-attribute:description

# Create the "Directory to Database Group Membership DN Map".
dsconfig create-dn-map --map-name "Directory to Database Group Membership DN Map" \
--set "from-dn-pattern:*,ou=people,dc=example,dc=com" --set "to-dn-pattern:{1}"

dsconfig create-attribute-mapping --map-name "Directory to DB Groups Map" \
--mapping-name memberID --type dn --set from-attribute:uniqueMember \
--set dn-map:"Directory to Database Group Membership DN Map"

dsconfig create-attribute-mapping --map-name "Directory to DB Groups Map" \
--mapping-name name --type direct --set from-attribute:cn

From the command line, run the following:

$ bin/dsconfig --port 7389 --bindPassword password --batch-file attr-mappings.txt \
--no-prompt
```

Step 5. Run the Resync Tool to Test the Configuration

The **resync** tool is used to test the sync configuration and connections. The tool has a **--dry-run** option that does not update the destination server but is convenient to test the configuration settings and report what is currently out of sync.

To Run Resync to Test the Configuration

- Run the **resync** command in "dry-run" mode to test the synchronization setup.

```
$ bin/resync --pipe-name ubid_to_oracle --dry-run
```

Step 6. Set the Startpoint in the Changelog

The **realtime-sync set-startpoint** command sets the starting point in the changelog to tell the Synchronization Server where to start when the Sync Pipe is started. This command provides a way to avoid syncing all of the changes that have happened in the past.

To Set the Starting Point

- Run the **realtime-sync set-startpoint** command to mark the point to start tracking changes in the change tracking mechanism.

```
$ bin/realtime-sync set-startpoint --end-of-changelog \
--pipe-name ubid_to_oracle
```

Step 7. Run the Resync Tool to Populate Data into One Endpoint

The `resync` tool is also used to populate a target server with data from the source.

To Run the Resync Tool to Populate Data onto a Target Server

- Run the `resync` command to populate data onto a newly configured target server. The Synchronization Server will make three passes to load data onto the server.

```
$ bin/resync --pipe-name ubid_to_oracle --numPasses 3
```

Step 8. Start the Sync Pipe

At this stage, we have configured everything necessary for the directory-to-database Sync Pipe. We only need to start it. Generally, it is preferable to use the `realtime-sync` tool to start and stop the Sync Pipes as well as start and stop the Sync configuration globally.

To Start the Sync Pipe

- Run the `realtime-sync` tool to start Sync Pipe.

```
$ bin/realtime-sync start --pipe-name ubid_to_oracle
```

Step 9. Debugging the Configuration

Typically, you will need to debug any problems after you run the `prepare-endpoint-server` and `resync` commands. There are a number of logging and tools options available when debugging the configuration as follows:

Check the Status

- Run the `status` tool to verify the synchronization. You should check if the servers are connected and that changes are being detected. You can enter your bindDN password and have the system use your bind DN and port as defaults.

```
$ status --bindPassword password
```

You can also restrict the status output to just list a single sync pipe using the `--pipe-name` option.

```
$ status --bindPassword password --pipe-name ubid_to_oracle
```

Check the Logs

- Increase the detail in the Sync log by changing the Sync Log Publisher handler's `logged-message-type` property to include: `change-applied-detailed`, `change-detected-detailed`, and `entry-mapping-details`.

```
$ dsconfig set-log-publisher-prop --publisher-name "File-Based Sync Logger" \
--set logged-message-type:change-applied-detailed \
--set logged-message-type:change-detected-detailed \
--set logged-message-type:change-failed-detailed \
--set logged-message-type:dropped-op-type-not-synchronized \
--set logged-message-type:dropped-out-of-scope \
--set logged-message-type:entry-mapping-details \
--set logged-message-type:no-change-needed
```

- Tail the **errors** log in the **logs** directory to locate any errors.
- Enable the debug logger (disabled by default), then rerun the **resync** command. You should disable the logger when no longer needed as it can impact performance.

```
# Enable the Debug Logger
dsconfig set-log-publisher-prop --publisher-name "File-Based Debug Logger" \
--set enabled:true

# Set the Debug Target and Verbosity Level
dsconfig create-debug-target --publisher-name "File-Based Debug Logger" \
--target-name com.unboundid.directory.sync.jdbc --set debug-level:verbose

# When finished with debugging, disable the logger
dsconfig set-log-publisher-prop --publisher-name "File-Based Debug Logger" \
--set enabled:false
```

- If your connections are working and the **resync** operation is working but you are seeing sync errors, tail the **sync-failed-ops** log. The problems could be in your attribute or DN maps.

Scripted Logging Methods

- The **ServerContext** class provides several logging methods which can be used to generate log messages and/or alerts from the scripted JDBC layer: **logMessage()**, **sendAlert()**, **debugCaught()**, **debugError()**, **debugInfo()**, **debugThrown()**, **debugVerbose()**, and **debugWarning()**. These are described in the *Server SDK API Javadocs*.

Testing One Entry at a Time

- Testing and debugging a configuration can be made more tractable if you test one entry at a time. When testing a directory-to-database sync configuration, the easiest way to do this is to use the **resync** tool's **"--sourceDNsFile"** option, which allows you to specify a list of one or more DNs to sync.

When to Restart the Sync Pipe

- Because **resync** is a separate process and independently loads the server configuration, it is not necessary to restart the sync pipe.
- Make sure to restart the Sync Pipes after modifying a script implementation.

```
$ bin/realtime-sync stop
$ bin/realtime-sync start
```

Note	Any Sync Server configuration changes automatically restart the Sync Pipe. Script implementation changes require a manual Sync Pipe restart.
------	--

Contact UnboundID

- If you require UnboundID Support assistance, the Support Staff usually requests that you run the `bin/collect-support-data` command so that they can locate the source of any problems. The command generates a zip file that you can send to UnboundID.

```
$ bin/collect-support-data --bindDN uid=admin,dc=example,dc=com \  
--bindPassword password
```

Configuring the Database-to-Directory Sync Pipe

The setup procedure for a Sync Pipe from a database to the directory server is similar to that of the directory-to-database sync configuration. However, there are slight differences in terms of enabling or setting properties for bidirectional synchronization.

To display the additional features of the `dsconfig` command, the following procedure uses `dsconfig` in non-interactive mode to set up the Database-to-Directory Sync Pipe. You can run each command from the command line, in scripts, or in a batch file when setting up multiple configurations.

The procedures assume that you have already set up the directory-to-database Sync Pipe and that it is fully operational and connected. Remember to include the connection parameters (hostname, port, bindDN, and bindPassword) with each `dsconfig` command.

General Tips When Syncing from a Database Source

When syncing from a database to a target endpoint server (directory server or RDBMS), remember to consider the following tips:

- It is advisable to identify the different database entry types that will be synced. There are two things that you need to do:
 - Set the `database-entry-type` property on the JDBC Sync Source (this is required), and make sure the entry types are what the triggers are inserting into the change tracking mechanism.
 - Create a Sync Class per entry type, so that you can easily distinguish between them and have different mappings and sync rules set up for each one.
- For each Sync Class, do the following:

- Make sure the `include-filter` property is set to match the entry type.
- Create a specific Attribute Mapping for every database column that you want to be synced to a LDAP attribute; add this to a single Attribute Map and set it on the Sync Class. This way, the script will not have to know about the schema on the directory side.
- Create a DN Map that recognizes the DNs generated by the script and map them to the correct location at the destination; set that on the Sync Class.
- Remove the default value of "-all-" from "auto-mapped-source-attrs" on the Sync Class, and replace it with the value "objectClass". The object class for the fetched source entry is determined by the scripted layer. You do not want any values from the database automatically mapped to an attribute with the same name, which is why we set up explicit Attribute Mappings. The exception to this rule is the `objectclass` attribute, which we want to directly map for CREATE operations. If this is not done, an error is generated due to the lack of structural object class in the entry.
- Change the "destination-correlation-attributes" to contain the attributes that uniquely represent the database entries on the directory server destination. This will likely be something other than the default, which is "dn".
- **Avoid Bidirectional Loopback.** Make sure to set the `ignore-changes-by-[user|dn]` property on both Sync Sources when configuring for bidirectional synchronization. This is important to make sure that changes are not looped back by the Sync Server.

To Create the Database-to-Directory Sync Pipe

1. Run the `create-sync-pipe-config` tool to configure the Database-to-Directory Sync Pipe. The steps are similar to those presented in the previous sections.
2. Run the `resync` tool to test the configuration. When testing a database-to-directory Sync Pipe, you must specify the `--entryType` of the database table that is synchronized.

```
$ bin/resync --pipe-name oracle_to_ubid --entryType account --dry-run
```

3. Run the `realtime-sync` tool with the `set-startpoint` subcommand to mark the point to start tracking changes in the change tracking mechanism.

```
$ bin/realtime-sync set-startpoint --end-of-changelog --pipe-name oracle_to_ubid
```

4. Run the `resync` tool to populate data onto a newly configured target server. The Synchronization Server will make three passes to load data onto the server.

```
$ bin/resync --pipe-name oracle_to_ubid --numPasses 3 --entryType account  
$ bin/resync --pipe-name oracle_to_ubid --numPasses 3 --entryType group
```

5. Run the `realtime-sync` tool to start Sync Pipe.

```
$ bin/realtime-sync start --pipe-name oracle_to_ubid
```


6. Troubleshoot the Sync Pipe as presented in “Step 9. Debugging the Configuration” on page 139.

You have successfully configured a bidirectional DBSync system.

6

Syncing Through Proxy Servers

Overview

The UnboundID Synchronization Server supports synchronization from one endpoint consisting of the UnboundID Directory Server (3.x), Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Servers (3.x), Sun™ Directory Server Enterprise Edition (DSEE 6.x, 7.x), Sun™ Directory Server (5.2 patch 3 or higher), Microsoft Active Directory, Oracle (10g, 11g), and Microsoft SQL Server (2005, 2008) systems. Because most data centers deploy their directory servers in a proxied environment, the UnboundID Synchronization Server can also synchronize data through a proxy server in both load-balanced and entry-balanced deployments. This feature is only available for deployments in combination with the UnboundID Directory Proxy Server (version 3.x) or the Alcatel-Lucent 8661 Directory Proxy Servers (3.x), and a backend set of standalone or replicated UnboundID Directory Servers (version 3.x) or Alcatel-Lucent 8661 Directory Servers (3.x).

This chapter presents the procedures to set up a Sync-through-Proxy deployment and provides some background information on how it works. Before setting up the Synchronization Server, review the section “Configuration Model” on page 13 to understand the important components of the Synchronization Server. Also, review the *UnboundID Directory Proxy Server Administration Guide* for background information on the proxy server.

This chapter presents the following topics:

- [Features](#)
- [How It Works](#)
- [About the Overall Sync-through-Proxy Configuration Process](#)
- [Configuring Sync Through Proxy](#)

Features

The UnboundID Synchronization Server (version 3.x) supports data synchronization through a proxy server from and to an endpoint consisting of the UnboundID Directory Proxy Server (version 3.0.0.0) with a backend set of UnboundID Directory Servers (version 3.0.0.0 or later). The three servers plus the Alcatel-Lucent 8661 Directory Server has been updated with additional components to provide seamless synchronization through the proxy.

The UnboundID Synchronization Server, UnboundID Directory Proxy Server, and UnboundID Directory Servers provide the following Sync-through-Proxy features:

- Synchronization is fully supported for load-balanced and entry-balanced proxy server deployments.
- The Directory Server and the Directory Proxy Server provide a common interface to detect and retrieve changes to be synchronized, including failover to an alternate source server.
- The Directory Proxy Server provides a built-in server affinity mechanism to ensure change-log searches are routed to the same directory server each time it is online. This allows for more efficient processing compared to load-balancing the searches across the backend directory servers.
- The UnboundID Synchronization Server uses the same configuration procedures as any other endpoint setup. The proxy server's operations are largely transparent to the Sync Server.
- Proxy transformations are not required as the Synchronization Server can handle them.

Note

If you are using the UnboundID Synchronization Server (version 3.x) with an earlier version of the UnboundID Directory Proxy Server and UnboundID Directory Server (versions 1.4.x, 2.2.x), you cannot run sync-through-proxy. The feature has not been backported to earlier versions.

How It Works


To handle data synchronization through a proxy server, the UnboundID Directory Server, UnboundID Directory Proxy Server, and the UnboundID Synchronization Server all have been updated with a new **cn=changelog** state management system that supports a token-based API and other components necessary for seamless data synchronization through the proxy. The tools have also been updated to handle these new components.

In a standard, non-proxied configuration, the Synchronization Server polls the source server for changes, determines if a change is necessary, and fetches the full entry from the source. Then, it finds the corresponding entry in the destination endpoint using flexible correlation rules and applies the minimal set of changes to bring any modified attributes into sync. The server fetches and compares the full entries to make sure it does not synchronize any stale data from the changelog.

In a proxied environment, the Synchronization Server essentially does the same thing but transparently to the user, it passes the request through a proxy server to the backend set of directory servers. The Synchronization Server uses the highest priority proxy server designated in its endpoint server configuration and can quickly use other proxy servers in the event of a failover. Figure 6-1 shows an example deployment with two endpoints consisting of a proxy server deployment in front of the backend set of directory servers. Remember that you can have one endpoint consisting of UnboundID Directory Proxy Server and UnboundID

Directory Servers while the other endpoint can be a directory server or RDBMS deployment (UnboundID Directory Server, Alcatel-Lucent 8661 Directory Server, Alcatel-Lucent 8661 Directory Proxy Servers, Sun DSEE 6.x, 7.x, Sun Directory Server 5.2 patch 3 or higher, Microsoft Active Directory, Oracle 10g, 11g, or Microsoft SQL Server 2005,2008).

FIGURE 6-1. Sync-Through-Proxy

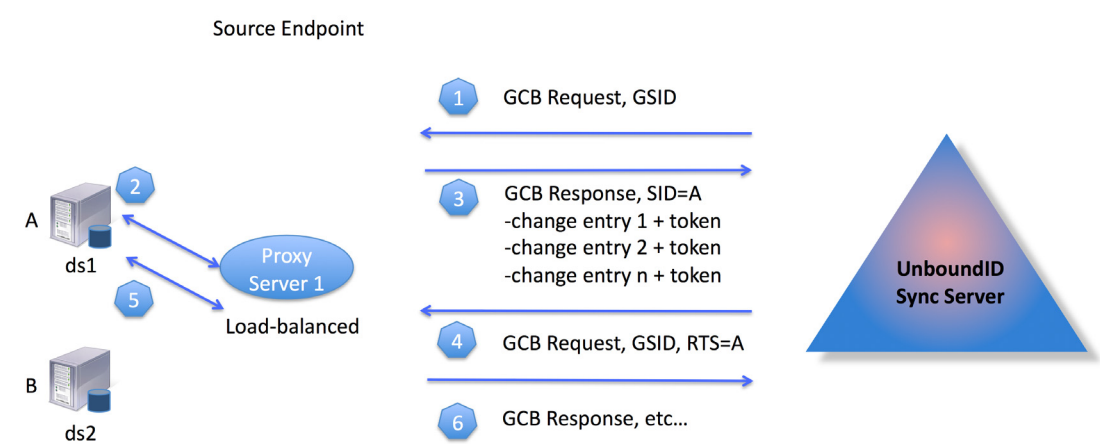


About the Get Changelog Batch and Get Server ID Controls

When the Synchronization Server runs a poll for any changes, it sends a Get Changelog Batch (GCB) Extended Request to the `cn=changelog` backend. The Get Changelog Batch looks for all changed entries in the changelog and asks for information on the server ID, change number, and replica state for each change. The Proxy Server routes the request to a directory server instance, which then returns a changed entry plus a token identifying the server ID, change number and replica state for each change. The proxy server then sends a Get Changelog Batch Extended Response back to the proxy server with this information. For entry-balanced deployments, the Directory Proxy Server must "re-package" the directory server tokens into its own proxy token to identify the specific data set. We will return to this a bit later.

To provide automatic server affinity in the proxied environment, the Synchronization Server uses the Get Server ID (GSID) Request Control together with the Get Changelog Batch (GCB) to identify the server ID of any fetched entry as illustrated in figure 6-2. The first time that the Synchronization Server issues a search request with the GCB, it also issues a GSID Request Control to identify the specific server ID that is processing the extended request. The Directory Proxy Server routes the request to the directory server instance, and then returns a server ID in the response. Upon the next GCB request, the Synchronization Server sends a Route to Server (RTS) Request Control specifying the server instance to access again (in this example, server A) in this batch session. It also issues a GSID Request Control to get an updated server in the event that the particular server (e.g., server A) is down. This method avoids round-robin server selection and provides more efficient overall change processing.

FIGURE 6-2. Get Changelog Batch Requests with Built-in Server Affinity



About the Directory Server and Directory Proxy Server Tokens

The Directory Server maintains a new changelog database index to determine at what point to resume sending changes (ADD, MODIFY, or DELETE operations) in its changelog. While a simple stand-alone directory server can track its resume point by the last change number sent, it is more difficult for replicated servers or servers deployed in entry-balanced environments. In replicated environments, each replica has a different change number ordering in its changelog as updates can come from a variety of sources: local write operations, changes from the other replication servers, or synchronized changes from other endpoints. Figure 6-3 illustrates a simple chart of two example changelogs in two replicated directory servers, server A and B. In the chart, *A* represents the replica identifier for a replicated subtree in Server A, and *B* represents the replica identifier for the same replicated subtree in server B. The replica identifiers with a hyphen ("-") mark any local, non-replicated but different changes. While the two replicas record all of the changes, you can see that the two changelogs have two different change number orderings as updates come in at different times.

FIGURE 6-3. Different Change Number Order in Two Replicated Changelogs

Server A			Server B		
ChangeNumber	ReplicaIdentifier	ReplicationCSN	ChangeNumber	ReplicaIdentifier	ReplicationCSN
1001	A _{ri}	10	2001	B _{ri}	11
1002	-	-	2002	A _{ri}	10
1003	A _{ri}	15	2003	-	-
1004	B _{ri}	11	2004	B _{ri}	12
1005	B _{ri}	12	2005	A _{ri}	15

To track this value, the directory server uses a changelog database index to identify the latest change number position corresponding to the replica with the highest replicationCSN number.

This information is encapsulated in a directory server token and returned in the Get Changelog Batch Response control to the Directory Proxy Server. The token has the following format:

Directory Server Token:

server ID, changeNumber, replicaState

For example, if the Proxy Server sends a request for any changed entries and the Directory Servers return the change number 1003 from server A and change number 2005 from server B, then each directory server token would return the following information:

Directory Server Token A:

serverID A, changeNumber 1003, replicaState {15 (A)}

Directory Server Token B:

serverID B, changeNumber 2005, replicaState {12 (B), 15 (A)}

Changelog Tracking in Entry-Balanced Deployments

Entry-balancing provides additional complexity in changelog tracking in that the shared areas of data can exist above the entry-balanced base DN in addition to each backend set having its own set of changes and tokens as mentioned previously. In Figure 6-4, the changelogs of two servers are shown with server A belonging to an entry-balanced set 1 and server C belonging to an entry-balanced set 2. Shared areas that exist above the entry-balanced base DN are assumed to be replicated to all servers. Thus, *SA* represents the replica identifier for that shared area on server A and *SC* represents the replica identifier for the same area on server C.

FIGURE 6-4. Changelogs in Entry-Balanced Sets

Set 1 - Server A			Set 2 - Server B		
ChangeNumber	ReplicaIdentifier	ReplicationCSN	ChangeNumber	ReplicaIdentifier	ReplicationCSN
1001	SA _{ri}	5	2001	SC _{ri}	10
1002	A _{ri}	10	2002	C _{ri}	20
1003	SC _{ri}	15	2003	SA _{ri}	5

The Directory Proxy Server cannot simply pass a directory server token from the client to the backend directory server backend and back again as each directory server has its own set of changes and its tokens. Thus, in an entry-balanced deployment, the Proxy Server must maintain its own token mechanism that associates a directory server token (changeNumber, replicaIdentifier, replicaState) to a particular backend set.

Proxy Token:

backendSetID 1: ds-token 1 (changeNumber, replicaIdentifier, replicaState)

backendSetID 2: ds-token 2 (changeNumber, replicaIdentifier, replicaState)

For example, if the Directory Proxy Server returned change 1002 from server A and change 2002 from server C, then the Proxy token would contain the following:

Proxy Token:

```
backendSetID 1: ds-token-1 {serverID A, changeNumber 1002, replicaState ( 5(SA), 15(A) }  
backendSetID 2: ds-token-2 {serverID C, changeNumber 2002, replicaState ( 10(SC),  
20(C) }
```

For each change entry returned by a backend, the Directory Proxy Server must also decide whether it is a duplicate of a change made to the backend set above the entry balancing base, since such changes appear in the changelog across all backend sets. If the change is a duplicate, then it is discarded. Otherwise, any new change is returned with a new value of the proxy token.

Finally, the Synchronization Server's `create-sync-pipe-config` tool has the logic to determine whether to use the token-based mechanism or the standard mechanism to process the changelog controls. It does so by accessing the server's root DSE entry, which advertises its token mechanism.

About the Overall Sync-through-Proxy Configuration Process

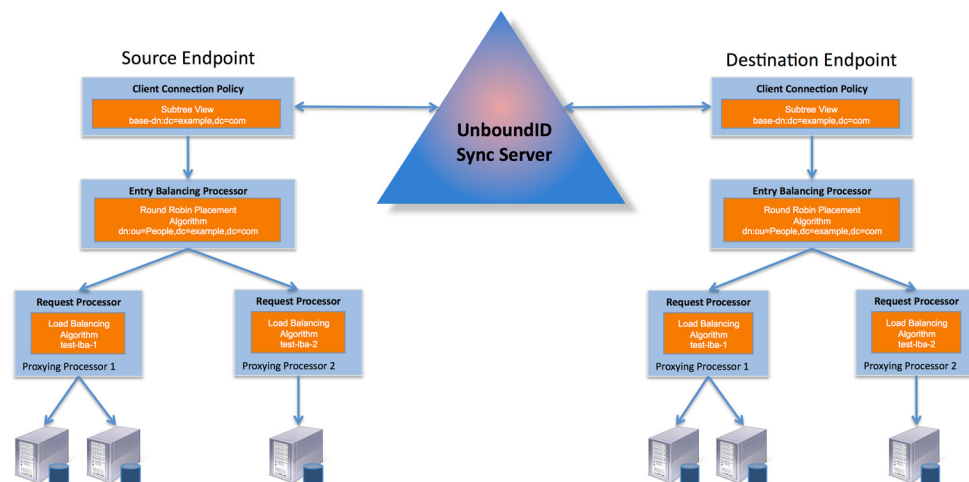
The procedure to configure a Sync-through-Proxy system follows the basic procedures for a standard Sync configuration. The overall configuration process is as follows:

1. Set up your proxy server with its backend set of directory servers at one endpoint or both endpoints.
2. Download the Synchronization Server zip build, and unpack it a directory of your choice.
3. From the server root directory of the Synchronization Server, run the `create-sync-pipe-config` command for your initial configuration. The command will interactively prompt you to input values necessary for your configuration.
4. Run the `prepare-external-server` command on the endpoint Directory Proxy Server instance and the backend set of directory servers. The Directory Proxy Server passes on a client request to the directory servers, which requires the `cn=Sync User` account be present on those servers for accessibility purposes.
5. Run the `resync --dry-run` command to test the configuration settings.
6. Run `realtime-sync set-startpoint` to initialize the starting point for synchronization. Note that you cannot use the `--change-number` option with a Sync-through-Proxy deployment but can use another option, such as, `--end-of-changelog` or `--change-sequence-number` options.
7. Run the `resync` command to populate data on a target endpoint.
8. Start the Sync Pipes using the `realtime-sync start` command.
9. Monitor the synchronization server using the status commands and logs.

Configuring Sync Through Proxy

This section presents the steps to configure a sync-through-proxy network. This section presents an example configuration that has its two endpoints consisting of an UnboundID Directory Proxy Server with a backend set of UnboundID Directory Servers: one set is replicated, the other is a standalone server. The Directory Proxy Server uses an entry-balanced environment for the **DN: ou=People, dc=example, dc=com** and provides a subtree view for **dc=example, dc=com** in its client connection policy. For this example, we assume that communication will be over standard LDAP and that failover servers are not installed or designated in the Synchronization Server.

FIGURE 6-5. Example Sync-through-Proxy Configuration



Configure the Directory Servers

The following procedures present the basic **dsconfig** command-line instructions in non-interactive mode to set up this example's backend set of directory servers. The specific setup procedures may differ based on your particular environment. For more detailed background information, please review the *UnboundID Directory Server Administration Guide*.

To Configure the Directory Servers

1. To begin installing and configuring the directory servers, unzip the directory server file in a location of your choice.

```
$ unzip UnboundID-DS-<version>.zip
```

2. If you plan to use SSL or StartTLS for communication, copy any keystore and truststore files to the **<server-root>/config** directory. For this example, we do not use SSL or StartTLS. All communication will be over standard LDAP.
3. If you have an existing schema file, copy the file to the **<server-root>/config/schema** directory.

4. Run the setup command from the root installation directory. Select your memory size options for your machine. For this example, create the base entry for the first directory server instance at host name `ldap-west-01.example.com`. Set the maximum JVM heap size to 4 GB.

```
$ ./setup --cli --no-prompt --listenAddress ldap-west-01.example.com \  
--ldapPort 389 --rootUserPassword password --baseDN dc=example,dc=com \  
--aggressiveJVMtuning --maxHeapSize 4g --acceptLicense
```

5. Configure the directory server. Here you can configure your local DB indexes, virtual attributes, log files, password policies, SASL mechanisms and global configuration properties. Minimally, you must enable the changelog database backend on your server instance, either from the command line or using a `dsconfig` batch file.

```
$ dsconfig set-backend-prop --backend-name changelog --set enabled:true
```

Note	Note that if you do not plan to have the specific directory server instance participate in synchronization, you do not need to enable its changelog.
------	--

6. Import your dataset into the first server instance. You must stop the server if it is running prior to the import.

```
$ bin/stop-ds  
$ bin/import-ldif --backendID userRoot --ldifFile ../backup.ldif  
$ bin/start-ds
```

7. Repeat steps 1–5 for the other two instances. Make sure to specify the hostname and port for each server instance.

8. On the first server instance, configure replication between server 1 and 2.

```
$ bin/dsreplication enable ---host1 ldap-west-01.example.com \  
--port1 389 --bindDN1 "cn=Directory Manager" --bindPassword1 password \  
--replicationPort1 8989 --host2 ldap-west-01.example.com --port2 389 \  
--bindDN2 "cn=Directory Manager" --bindPassword2 password \  
--replicationPort2 9989 --adminUID admin --adminPassword admin \  
--baseDN dc=example,dc=com --no-prompt
```

9. Change to the second server, and then initialize the second server with data from the first instance.

```
$ bin/dsreplication initialize --hostSource ldap-west-01.example.com \  
--portSource 389 --hostDestination ldap-west-02.example.com \  
--portDestination 389 --baseDN "dc=example,dc=com" --adminUID admin \  
--adminPassword admin --no-prompt
```

10. Run `dsreplication status` to check your replicas.

```
$ bin/dsreplication status --hostname ldap-west-01.example.com \  
--port 389 --adminPassword admin --no-prompt
```

Configuring the Proxy Servers

The following procedures present the basic **dsconfig** command-line instructions in non-interactive mode to set up your proxy servers. The procedures configure the proxy servers from a bottom-up perspective: from defining the external servers to configuring the client-connection policy. If you are configuring the proxy servers for the first time, we recommend using the **create-initial-proxy-config** tool. The tool provides a command-line wizard presenting the interactive steps to configure your proxy server. For additional changes, you can use the **dsconfig** tool to fine-tune your proxy server. For more detailed background information, please review the *UnboundID Directory Proxy Server Administration Guide*.

1. To begin installing and configuring the directory servers, unzip the UnboundID Directory Proxy Server file in a location of your choice.

```
$ unzip UnboundID-Proxy-<version>.zip
```

2. Run the **setup** command from the proxy server root installation directory. For this example, the default bind DN will be "cn=Directory Manager" and bind DN (or root user) password is set to "pxy-pwd." You can also use the **--aggressiveJVMtuning** with the **--maxHeapSize** options to set the amount of JVM memory for this application.

```
$ setup --cli --no-prompt --ldapPort 389 --rootUserPassword pxy-pwd \
--acceptLicense
```

3. From the Directory Proxy Server root directory, run the **prepare-external-server** command to set up the **cn=Proxy User** account and its privileges to give the proxy server access to the backend directory servers. After you press Enter, the command tests the connection to the server, creates the "cn=Proxy User" account, tests the connection to the account again, and checks the backend.

```
$ bin/prepare-external-server --no-prompt \
--hostname ldap-west-01.example.com \
--port 389 --bindDN "cn=Directory Manager" --bindPassword pxy-pwd \
--proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" \
--proxyBindPassword pass --baseDN "dc=example,dc=com"
```

4. Repeat step 3 for the other two directory server instances in this example. Make sure to specify the specific hostname and port.
5. Next, run the **dsconfig** command to define the external servers and their types. The Directory Proxy Server communicates with these external servers through the **cn=Proxy User** account. Normally, you may want to set up any health checks and designate your server locations using this command. However, for this example, we use round-robin load-balancing algorithms, which do not require any health checks or locations to be specified.

```
$ bin/dsconfig create-external-server --server-name ldap-west-01 \
--type "unboundid-ds" --set "server-host-name:ldap-west-01.example.com" \
--set "server-port:389" --set "bind-dn:cn=Proxy User" \
--set "password:password" --bindDN "cn=Directory Manager" \
--bindPassword pxy-pwd --no-prompt
```

```
$ bin/dsconfig create-external-server --server-name ldap-west-02 \
--type "unboundid-ds" --set "server-host-name:ldap-west-02.example.com" \
--set "server-port:389" --set "bind-dn:cn=Proxy User" \
--set "password:password" --bindDN "cn=Directory Manager" \
--bindPassword pxy-pwd --no-prompt
```

```
$ bin/dsconfig create-external-server --server-name ldap-west-03 \
--type "unboundid-ds" --set "server-host-name:ldap-west-03.example.com" \
--set "server-port:389" --set "bind-dn:cn=Proxy User" \
--set "password:password" --bindDN "cn=Directory Manager" \
--bindPassword pxy-pwd --no-prompt
```

6. Next, create the load-balancing algorithm for the external servers. In this example, create one algorithm for the two replicated servers, and another for the third server.

```
$ bin/dsconfig create-load-balancing-algorithm \
--algorithm-name "test-lba-1" \
--type "round-robin" --set "enabled:true" \
--set "backend-server:ldap-west-01" \
--set "backend-server:ldap-west-02" --set "use-location:false" \
--set "operation-to-retry:add" --bindDN "cn=Directory Manager" \
--bindPassword pxy-pwd --no-prompt
```

```
$ bin/dsconfig create-load-balancing-algorithm \
--algorithm-name "test-lba-2" \
--type "round-robin" --set "enabled:true" \
--set "backend-server:ldap-west-03" --set "use-location:false" \
--set "operation-to-retry:add" --bindDN "cn=Directory Manager" \
--bindPassword pxy-pwd --no-prompt
```

7. Next, configure the proxying request processor. The request processor provides the logic to either process the operation directly, forward the request to another server, or hand off the request to another request processor. You will define two request processors, one for each load-balanced directory server set.

```
$ bin/dsconfig create-request-processor \
--processor-name "proxying-processor-1" --type "proxying" \
--set "load-balancing-algorithm:test-lba-1" \
--bindDN "cn=Directory Manager" --bindPassword pxy-pwd --no-prompt
```

```
$ bin/dsconfig create-request-processor \
--processor-name "proxying-processor-2" --type "proxying" \
--set "load-balancing-algorithm:test-lba-2" \
--bindDN "cn=Directory Manager" --bindPassword pxy-pwd --no-prompt
```

8. At this stage, we define an entry-balancing-algorithm request processor. This request processor is used to distribute entries under a common parent entry among multiple backend sets. A backend set is a collection of replicated directory servers that contain identical portions of the data. This request processor uses multiple proxying request processors to process operations for the various backend LDAP servers.

```
$ bin/dsconfig create-request-processor \
--processor-name "entry-balancing-processor" \
--type "entry-balancing" \
--set "entry-balancing-base-dn:ou=People,dc=example,dc=com" \
--set "subordinate-request-processor:proxying-processor-1" \
--set "subordinate-request-processor:proxying-processor-2" \
--bindDN "cn=Directory Manager" --bindPassword pxy-pwd --no-prompt
```

9. Next, define the placement algorithm, which selects the server set to use for new add operations to create new entries. In this example, we define a placement algorithm with a round-robin algorithm that forwards LDAP add requests to backends sets in a round-robin manner.

```
$ bin/dsconfig create-placement-algorithm \  
--processor-name "entry-balancing-processor" \  
--algorithm-name "round-robin-placement" \  
--set "enabled:true" --type "round-robin" \  
--bindDN "cn=Directory Manager" --bindPassword pxy-pwd --no-prompt
```

10. Define the subtree view that specifies the base DN for the entire deployment.

```
$ bin/dsconfig create-subtree-view --view-name "test-view" \  
--set "base-dn:dc=example,dc=com" \  
--set "request-processor: entry-balancing-processor" \  
--bindDN "cn=Directory Manager" --bindPassword pxy-pwd --no-prompt
```

11. Finally, define a client connection policy that specifies how the client connects to the proxy server.

```
$ bin/dsconfig set-client-connection-policy-prop --policy-name "default" \  
--add "subtree-view:test-view" --bindDN "cn=Directory Manager" \  
--bindPassword pxy-pwd --no-prompt
```

You have successfully configured the first endpoint topology.

Configuring the Destination Proxy Deployment

To configure the destination proxy deployment, follow the procedures in the previous two sections. From the server root directory of the second proxy server, make sure to specify the hostnames and ports for the external servers. A summary of the example configuration commands are listed in Table 6-1.

TABLE 6-1. Summary of Proxy Configuration Commands for the Source and Destination Deployments

Component	Source Proxy Topology	Destination Proxy Topology
Prepare External Servers	<pre>prepare-external-server --no-prompt --hostname "ldap-west-01.example.com" --port "389" --bindDN "cn=Directory Manager" --bindPassword "password" --proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" --base DN" dc=example,dc=com" prepare-external-server --no-prompt --hostname "ldap-west-02.example.com" --port "389" --bindDN "cn=Directory Manager" --bindPassword "password" --proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" --base DN" dc=example,dc=com" prepare-external-server --no-prompt --hostname "ldap-west-03.example.com" --port "389" --bindDN "cn=Directory Manager" --bindPassword "password" --proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" --base DN" dc=example,dc=com"</pre>	<pre>prepare-external-server --no-prompt --hostname "ldap-east-01.example.com" --port "389" --bindDN "cn=Directory Manager" --bindPassword "password" --proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" --base DN" dc=example,dc=com" prepare-external-server --no-prompt --hostname "ldap-east-02.example.com" --port "389" --bindDN "cn=Directory Manager" --bindPassword "password" --proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" --base DN" dc=example,dc=com" prepare-external-server --no-prompt --hostname "ldap-east-03.example.com" --port "389" --bindDN "cn=Directory Manager" --bindPassword "password" --proxyBindDN "cn=Proxy User,cn=Root DNs,cn=config" --base DN" dc=example,dc=com"</pre>
External Servers	<pre>dsconfig create-external-server --server-name: "ldap-west-01" --type "unboundid-ds" --set "server-host-name:ldap-west-01.example.com" --set "server-port:389" --set "bind-dn:cn=Proxy User" --set "password:password" dsconfig create-external-server --server-name: "ldap-west-02" --type "unboundid-ds" --set "server-host-name:ldap-west-02.example.com" --set "server-port:389" --set "bind-dn:cn=Proxy User" --set "password:password" dsconfig create-external-server --server-name: "ldap-west-03" --type "unboundid-ds" --set "server-host-name:ldap-west-03.example.com" --set "server-port:389" --set "bind-dn:cn=Proxy User" --set "password:password"</pre>	<pre>dsconfig create-external-server --server-name: "ldap-east-01" --type "unboundid-ds" --set "server-host-name:ldap-east-01.example.com" --set "server-port:389" --set "bind-dn:cn=Proxy User" --set "password:password" dsconfig create-external-server --server-name: "ldap-east-02" --type "unboundid-ds" --set "server-host-name:ldap-east-02.example.com" --set "server-port:389" --set "bind-dn:cn=Proxy User" --set "password:password" dsconfig create-external-server --server-name: "ldap-east-03" --type "unboundid-ds" --set "server-host-name:ldap-east-03.example.com" --set "server-port:389" --set "bind-dn:cn=Proxy User" --set "password:password"</pre>

TABLE 6-1. Summary of Proxy Configuration Commands for the Source and Destination Deployments

Component	Source Proxy Topology	Destination Proxy Topology
Load-Balancing Algorithm	<pre>dsconfig create-load-balancing-algorithm --algorithm-name "test-lba-1" --type "round-robin" --set "enabled:true" --set "backend-server: ldap-west-01" --set "backend-server: ldap-west-02" --set "use-location:false" --set "operation-to-retry: add" dsconfig create-load-balancing-algorithm --algorithm-name "test-lba-2" --type "round-robin" --set "enabled:true" --set "backend-server: ldap-west-03" --set "use-location:false" --set "operation-to-retry: add"</pre>	<pre>dsconfig create-load-balancing-algorithm --algorithm-name "test-lba-1" --type "round-robin" --set "enabled:true" --set "backend-server: ldap-east-01" --set "backend-server: ldap-east-02" --set "use-location:false" --set "operation-to-retry: add" dsconfig create-load-balancing-algorithm --algorithm-name "test-lba-2" --type "round-robin" --set "enabled:true" --set "backend-server: ldap-east-03" --set "use-location:false" --set "operation-to-retry: add"</pre>
Request Processors (load-balancing)	<pre>dsconfig dsconfig create-request-processor --processor-name "proxying-processor-1" --type "proxying" --set "load-balancing-algorithm:test-lba-1" dsconfig create-request-processor --processor-name "proxying-processor-2" --type "proxying" --set "load-balancing-algorithm:test-lba-2"</pre>	Same as source
Request Processors (entry-balancing)	<pre>dsconfig create-request-processor --processor-name "entry-balancing-processor" --type "entry-balancing" --set "entry-balancing-base-dn:ou=People,dc=example,dc=com" --set "subordinate-request-processor:proxying-processor-1" --set "subordinate-request-processor:proxying-processor-2"</pre>	Same as source
Placement Algorithm	<pre>dsconfig create-placement-algorithm --processor-name "entry-balancing-processor" --algorithm-name "round-robin-placement" --set "enabled:true" --type "round-robin"</pre>	Same as source
Subtree View	<pre>dsconfig create-subtree-view --view-name: "test-view" --set "base-dn:dc=example,dc=com" --set "request-processor:entry-balancing-processor"</pre>	Same as source
Client Connection Policy	<pre>dsconfig set-client-connection-policy-prop --policy-name: "default" --add "subtree-view:test-view"</pre>	Same as source

Configuring the Synchronization Server

At this stage, the UnboundID Directory Proxy Server and its backend set of UnboundID Directory Server instances should be configured and fully functional for each endpoint, which is labelled as ldap-west and ldap-east in this example.

1. Download the UnboundID Synchronization build zip. Unzip the code in a directory of your choice.

```
$ unzip UnboundID-Sync-<version>.zip
```

2. If you are configuring the Synchronization Server from scratch, you must ensure that you have JDK1.6 update 20. Set the **JAVA_HOME** environment variable and your **PATH** or **CLASSPATH** variables accordingly.
3. From the Synchronization Server root directory, run the **setup** tool. For this example, the default bind DN will be "**cn=Directory Manager**" and bind DN (or root user) password is set to "password". You can also use the **--aggressiveJVMtuning** with the **--maxHeapSize** options to set the amount of JVM memory for this application.

```
$ setup --no-prompt --ldapPort 389 --rootUserPassword password --acceptLicense
```

4. From the Synchronization Server root directory, run the **create-sync-pipe-config** tool, and then, press Enter to continue.

```
$ bin/create-sync-pipe-config
```

5. On the Synchronization Directory menu, enter the number associated with the type of synchronization you want to configure: 1 for One-Way, 2 for bidirectional. For this example, type 2 for bidirectional, which will require that you configure two Sync Pipes: one for each direction (e.g., "proxy 1 to proxy 2" and "proxy 2 to proxy 1").
6. Next, you will be prompted to configure the first endpoint server, which will be the first Directory Proxy Server topology. On the First Endpoint Type menu, enter the number for the type of backend datastore for the first endpoint. In this example, type 2 for the UnboundID Proxy Server.

```
>>>> First Endpoint Type
```

```
Enter the type of data store for the source endpoint:
```

- 1) UnboundID Directory Server
- 2) UnboundID Proxy Server
- 3) Alcatel-Lucent Directory Server
- 4) Alcatel-Lucent Proxy Server
- 5) Sun Directory Server
- 6) Microsoft Active Directory
- 7) Microsoft SQL Server
- 8) Oracle Database
- 9) Generic JDBC

- b) back
- q) quit

```
Enter choice [1]: 2
```

7. Next, enter a descriptive name for the first endpoint. For this example, use "UnboundID Proxy 1".
8. Next, enter the base DN where the Sync Server can search for the entries on the first endpoint server. For this example, press Enter to accept the default, **dc=example,dc=com**.

9. Specify the type of security when communicating with the endpoint server. For this example, select **None**.
10. Enter the hostname and port of the endpoint server. The Sync Server will automatically test the connection to the endpoint server. Repeat the step if you are configuring another server for failover.
11. Next, enter the Sync User account that will be used to access the endpoint server (i.e., proxy server 1). Enter **cn=Sync User,cn=Root DNs,cn=config**, then, enter a password for the account.
12. At this point, you have defined the first endpoint deployment using the Proxy Server (e.g., ldap-west). Repeat steps 3–8 to define the second proxy deployment (e.g., ldap-east) on the Synchronization Server.
13. At this point, you will be prompted to "prepare" the endpoint servers in the topology. The endpoint servers here refer to the proxy servers in this example. This step ensures that the Sync User account is present on each server and that it has the proper privileges to allow communication between the Sync Server and the proxy servers.

In addition to preparing the proxy server, the Synchronization Server must also prepare the backend set of directory servers as the proxy server passes through the authorization to access these servers. If they have not been prepared, you will see the following messages to invoke the commands prior to starting synchronization. Also note that each endpoint is a source and a destination in a bidirectional sync network; therefore, you must use **--isSource** and **--isDestination** options. If you are configuring a one-way Sync Pipe, you must specify **--isSource** for the first endpoint.

Discovering additional servers that require preparation

Server ldap-west-01.example.com:389 requires preparation. Before starting synchronization you must invoke the following command, substituting the correct password for [password]:

```
prepare-endpoint-server --hostname ldap-west-01.example.com --port 389 \  
--baseDN dc=example,dc=com --isSource --isDestination \  
--syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \  
--syncServerBindPassword "[password]"
```

Server ldap-west-02.example.com:389 requires preparation. Before starting synchronization you must invoke the following command, substituting the correct password for [password]:

```
prepare-endpoint-server --hostname ldap-west-02.example.com --port 389 \  
--baseDN dc=example,dc=com --isSource --isDestination \  
--syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \  
--syncServerBindPassword "[password]"
```

Server ldap-west-03.example.com:389 requires preparation. Before starting synchronization you must invoke the following command, substituting the correct password for [password]:

```
prepare-endpoint-server --hostname ldap-west-03.example.com --port 389 \  
--baseDN dc=example,dc=com --isSource --isDestination \  
--syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \  
--syncServerBindPassword "[password]"
```

Discovering additional servers that require preparation Done

14. Next, repeat step 13 to prepare the second endpoint server (i.e., in this example, the second proxy server). Again, if you have not prepared the underlying directory servers (e.g., ldap-east-01, ldap-east-02, ldap-east-03), you will need to run the commands prior to starting synchronization.
15. Define the Sync Pipe from proxy 1 to proxy 2. First, enter a descriptive name for the Sync Pipe. In this example, accept the default "UnboundID Proxy 1 to UnboundID Proxy 2."
16. Next, if you want to customize on a per-entry basis how attributes get synchronized, you must define one or more sync classes. Type **yes** if you have specific attribute or DN mappings, create a sync class for the special cases, and use *default* sync class for all other mappings. For this example, press Enter to accept the default (no).
17. For the default Sync Class Operations, specify the operations that will be synchronized for the default sync class. For this example, accept the default ([1,2,3]) for Creates, Deletes, and Modifies.
18. Repeat steps 12–14 to define the second sync pipe, "UnboundID Proxy 2 to UnboundID Proxy 1." For this example, accept the default prompts.
19. Finally, review the configuration settings, and then accept the default (write configuration) to the Synchronization Server. The Synchronization Server writes your configuration settings to a file, **sync-pipe-cfg.txt**, so that you can apply these configurations to other failover Sync Servers if necessary. Connect to the Synchronization Server so that the server will be updated with your settings.

Set the Proxy-Server and Use-Changelog-Batch-Request Properties

1. Depending on how you configured your Sync configuration, there are two properties that you need to verify on each endpoint: **proxy-server** and **use-changelog-batch-request**. The **proxy-server** property should specify the name of the proxy server, while the **use-changelog-batch-request** should be set to **true**. Remember to add the connection parameters to your Synchronization Server (hostname, port, bind DN, and bind password). The following commands check the properties on the Sync Source. (You can run a similar command on the Sync Destination.)

```
$ bin/dsconfig get-sync-source-prop --source-name "Proxy 1" \
--property "proxy-server" --property "use-changelog-batch-request" --no-prompt
```

2. From the server root directory, run the **dsconfig** command to set a flag indicating that the endpoints are proxy servers. Remember to add the connection parameters for the Synchronization Server (hostname, port, bind DN, and bind password) with the following commands:

```
$ bin/dsconfig set-sync-source-prop --source-name "Proxy 1" \
--set proxy-server:ldap-west-01 --set use-changelog-batch-request:true
```

Run prepare-external-server on the Backend Set of Directory Servers

1. From the server root directory, run the **prepare-external-server** command on each of directory server instances in the first endpoint topology that you want to have participate in synchronization.

```
$ prepare-endpoint-server --hostname ldap-west-01.example.com --port 389 \
--baseDN dc=example,dc=com --isSource --isDestination \
--syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \
--syncServerBindPassword "password"

$ prepare-endpoint-server --hostname ldap-west-02.example.com --port 389 \
--baseDN dc=example,dc=com --isSource --isDestination \
--syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \
--syncServerBindPassword "password"

$ prepare-endpoint-server --hostname ldap-west-03.example.com --port 389 \
--baseDN dc=example,dc=com --isSource --isDestination \
--syncServerBindDN "cn=Sync User,cn=Root DNs,cn=config" \
--syncServerBindPassword "password"
```

2. Repeat the previous step on the other endpoint topology (e.g., ldap-east).

Testing and Starting the Configuration

1. Run the **resync --dry-run** command to test the configuration settings. We recommend running it for each sync pipe, debug any issues, then run the command again for the other sync pipe.

```
$ bin/resync --pipe-name "UnboundID Proxy 1 to UnboundID Proxy 2" --dry-run
$ bin/resync --pipe-name "UnboundID Proxy 2 to UnboundID Proxy 1" --dry-run
```

2. Run **realtime-sync set-startpoint** to initialize the starting point for synchronization.

```
$ realtime-sync set-startpoint --end-of-changelog --pipe-name "UnboundID Proxy 1 to
UnboundID Proxy 2"
$ realtime-sync set-startpoint --end-of-changelog --pipe-name "UnboundID Proxy 2 to
UnboundID Proxy 1"
```

Note	For Sync-through-Proxy deployments, you cannot use the --change-number option with the realtime-sync set-startpoint command as the Sync Server cannot retrieve specific change numbers from the backend set of directory servers. You can use the --change-sequence-number , --end-of-changelog or the other options available for the tool.
------	--

3. Run the **resync** command to populate data on the endpoint destination server if necessary.

```
$ bin/resync --pipe-name "UnboundID Proxy 1 to UnboundID Proxy 2" --numPasses 3
```

4. Start the Sync Pipes using the **realtime-sync start** command.

```
$ bin/realtime-sync start --pipe-name "UnboundID Proxy 1 to UnboundID Proxy 2"
$ bin/realtime-sync start --pipe-name "UnboundID Proxy 2 to UnboundID Proxy 1"
```

5. Monitor the synchronization server using the **status** commands and logs.

You have successfully configured a Sync-through-Proxy deployment.

7

Managing Logging and Monitoring

Overview

The Synchronization Server supports extensive logging features to track any aspect of your Synchronization topology. You can also set up administrative alert handlers to notify of any specific events.

This chapter presents the following information:

- [Working with Logs](#)
- [Default Synchronization Server Logs](#)
- [Sync Log Messaging Types](#)
- [Creating New Log Publishers](#)
- [Configuring Log Rotation](#)
- [Configuring Log Retention](#)
- [Working with Administrative Alert Handlers](#)
- [Running the Status Tool](#)
- [Monitoring the Synchronization Server](#)
- [Profiling Server Performance Using the Periodic Stats Logger](#)

Working with Logs

The UnboundID Synchronization Server supports different classes of log publishers that can be used to provide the monitoring information for sync, access, debug, and error messages that occur during normal server processing. The Synchronization Server provides a standard set of default log files as well as mechanisms to configure custom log publishers with its own log rotation and retention policies.

Types of Log Publishers

The UnboundID Synchronization Server provides a number of different types of loggers that can be used to get processing information about the server. There are several primary types of loggers:

- **Sync loggers** provide information about synchronization actions that occur within the server. Specifically, the Sync Log records all changes applied, detected or failed; dropped operations that were not synchronized; changes dropped due to being out of scope, or no changes needed for synchronization. The log also shows the entries that were involved in the synchronization process.
- **Resync loggers** provide summaries or details of synchronized entries and any missing entries in the Sync Destination.
- **Error loggers** provide information about warnings, errors, or significant events that occur within the server.
- **Debug loggers** can provide detailed information about processing performed by the server, including any exceptions caught during processing, detailed information about data read from or written to clients, and accesses to the underlying database.
- **Access loggers** provide information about LDAP operations processed within the server. This log only applies to operations performed in the server. This includes configuration changes, searches of monitor data, and bind operations for authenticating administrators using the command-line tools and the UnboundID Sync Management console.

By default, the following log publishers are enabled on the system:

- File-based sync logger
- File-based access logger
- File-based error logger

The UnboundID Synchronization Server also provides a File-based Audit Logger, which is a special type of access logger that can provide detailed information about changes processed within the server, and a File-based Debug Logger. Both are disabled by default.

Default Synchronization Server Logs

The Synchronization Server provides a standard set of default log files to monitor the server activity. You can view this set of logs in the `UnboundID-Sync/logs` directory. The following default log files are available:

TABLE 7-1. Synchronization Server Logs

Log File	Description
access	File-based Access Log that records LDAP operations processed by the Synchronization Server. Access log records can be used to provide information about problems during operation processing (for example, unindexed searches, failed requests, etc.), and provide information about the time required to process each operation.
config-audit.log	Records information about changes made to the Synchronization server configuration in a format that can be replayed using the <code>dsconfig</code> tool

TABLE 7-1. Synchronization Server Logs

Log File (Continued)	Description
errors	File-based Error Log. Provides information about warnings, errors, and significant events that occur during server processing.
failed-ops	Lists all failed sync operations.
server.out	Records anything written to standard output or standard error, which includes startup messages. If garbage collection debugging is enabled, then the information will be written to server.out.
server.pid	Stores the server's process ID.
server.status	Stores the timestamp, a status code, and an optional message providing additional information on the server status.
setup.log	Records messages that occur during the initial configuration of a Synchronization server with the setup command.
sync	File-based Sync Log that records synchronization operations processed by the server. Specifically, the log records all changes applied, detected or failed; dropped operations that were not synchronized; changes dropped due to being out of scope, or no changes needed for synchronization.
sync-pipe-cfg.txt	Records the configuration changes used with the <code>bin/create-sync-pipe-config</code> tool. The file is placed wherever the tool is run. Typically, this is in server-root or in the bin directory.
tools	Holds logs for long running utilities. Current and previous copies of the log are present in the directory.
update.log	Records messages that occur during a Synchronization Server upgrade.

Viewing the List of Log Publishers

You can quickly view the list of log publishers on the Synchronization Server using the `dsconfig` tool.

To View the List of Log Publishers

Use `dsconfig` to view the log publishers.

```
$ bin/dsconfig list-log-publishers \
  --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret \
  --port 1389 --no-prompt
```

Log Publisher	: Type	: enabled
Expensive Operations Access Logger	: file-based-access	: false
Failed Operations Access Logger	: file-based-access	: true
File-Based Access Logger	: file-based-access	: true
File-Based Audit Logger	: file-based-audit	: false
File-Based Debug Logger	: file-based-debug	: false
File-Based Error Logger	: file-based-error	: true
File-Based Sync Logger	: file-based-sync	: true
Successful Searches with No Entries Returned	: file-based-access	: false

Sync Log Messaging Types

The Synchronization Server logs certain types of log messages with the sync log. You can control which message types can be included or excluded from the logger in a custom log publisher.

TABLE 7-2. Sync Log Messaging Types

Messaging Type	Description
change-applied	Default summary message. Logged each time a change is applied successfully.
change-detected	Default summary message. Logged each time a change is detected.
change-failed-detailed	Default detail message. Logged when a change cannot be applied. It includes the reason for the failure and details about the change that can be used to manually repair the failure.
dropped-op-type-not-synchronized	Default summary message. Logged when a change is dropped because the operation type (for example, ADD) is not synchronized for the matching Sync Class.
dropped-out-of-scope	Default summary message. Logged when a change is dropped because it does not match any Sync Class.
no-change-needed	Default summary message. Logged each time a change is dropped because the modified source entry is already in-sync with the destination entry.
change-detected-detailed	Optional detail message. Logged each time a change is detected. It includes attribute values for added and modified entries. This level of information is often useful for diagnosing problems, but it causes log files to grow faster, which impacts performance.
entry-mapping-details	Optional detail message. Logged each time a source entry (attributes and DN) are mapped to a destination entry. This level of information is often useful for diagnosing problems, but it causes log files to grow faster, which impacts performance.
change-applied-detailed	Optional detail message. Logged each time a change is applied. It includes attribute values for added and modified entries. This level of information is often useful for diagnosing problems, but it causes log files to grow faster, which impacts performance.
change-failed	Optional summary message. Logged when a change cannot be applied. It includes the reason for the failure but not enough information to manually repair the failure.
intermediate-failure	Optional summary message. Logged each time an attempt to apply a change fails. Note that a subsequent retry of applying the change might succeed.

Creating New Log Publishers

The UnboundID Synchronization Server provides options to help you create your own custom log publishers with the `dsconfig` command.

When you create a new log publisher, you must also configure the log retention and rotation policies for each new publisher. For more information, see “Configuring Log Rotation” on page 168 and “Configuring Log Retention” on page 168.

To Create a New Log Publisher

1. Use the `dsconfig` command in non-interactive mode to create and configure the new log publisher. This example shows how to create a logger that only logs entry mapping details.

```
$ bin/dsconfig create-log-publisher \  
--publisher-name "Entry Mapper" \  
--type file-based-sync \  
--set enabled:true \  
--set logged-message-type:entry-mapping-details \  
--set log-file:logs/sync \  
--set "rotation-policy:24 Hours Time Limit Rotation Policy" \  
--set "rotation-policy:Size Limit Rotation Policy" \  
--set "retention-policy:File Count Retention Policy" \  
--set log-file:logs/entry-mapper.log \  
--bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret \  
--port 1389 --no-prompt
```

2. View the Log Publishers.

```
$ bin/dsconfig list-log-publishers \  
--bindDN "uid=admin,dc=example,dc=com" \  
--bindPassword secret --port 1389 --no-prompt
```

Log Publisher	: Type	: enabled
Entry Mapper	: file-based-sync	: true
Expensive Operations Access Logger	: file-based-access	: false
Failed Operations Access Logger	: file-based-access	: true
File-Based Access Logger	: file-based-access	: true
File-Based Audit Logger	: file-based-audit	: false
File-Based Debug Logger	: file-based-debug	: false
File-Based Error Logger	: file-based-error	: true
File-Based Sync Logger	: file-based-sync	: true
Successful Searches with No Entries Returned	: file-based-access	: false

To Command a Log Publisher Using dsconfig Interactive Command-Line Mode

1. On the command line, type `bin/dsconfig`.
2. Authenticate to the server by following the prompts.
3. On the UnboundID Synchronization Server Configuration console main menu, select the option to configure the log publisher.
4. On the Log Publisher Management menu, select the option to create a new log publisher.
5. Select the Log Publisher type. In this case, select File-Based Sync Logger.
6. Type a name for the log publisher.
7. Enable it.
8. Type the path to the log file, relative to the Synchronization Server root. For example, `logs/entry-mapper.log`.

9. On the File Based Sync Log Publisher Properties menu, select the option for `logged-message-type`. Remove or add properties.
10. On the Log Publisher Properties menu, change the rotation policy and/or the retention policy if necessary.
11. Type `Y` to save and apply the changes.

Configuring Log Rotation

The Synchronization Server allows you to configure the log rotation policy for the server. When the rotation limits are reached, the Synchronization Server rotates the current log and starts a new log. If you create a new log publisher, you must configure at least one log rotation policy.

You can select the following properties:

- **Time Limit Rotation Policy.** Rotates the log based on the length of time since the last rotation. Default implementations are provided for rotation every 24 hours and every 7 days.
- **Fixed Time Rotation Policy.** Rotates the logs every day at a specified time (based on 24-hour time). The default time is 2359.
- **Size Limit Rotation Policy.** Rotates the logs when the file reaches the maximum size for each log. The default size limit is 100 MBytes.
- **Never Rotate Policy.** Used in a rare event that does not require log rotation.

To Configure the Log Rotation Policy

Use `dsconfig` to modify the log rotation policy for the Sync Logger.

```
$ bin/dsconfig set-log-publisher-prop \
  --publisher-name "File-Based Sync Logger" \
  --remove "rotation-policy:24 Hours Time Limit Rotation Policy" \
  --add "rotation-policy:7 Days Time Limit Rotation Policy" \
  --hostname host1 --port 1389 --bindDN "uid=admin,dc=example,dc=com" \
  --bindPassword secret --no-prompt
```

Configuring Log Retention

The Synchronization Server allows you to configure a log retention policy for each log on the server. When the retention limits are reached, the Synchronization Server removes the oldest archived log prior to creating a new log. Log retention is only effective if you have a log rota-

tion policy in place. If you create a new log publisher, you must configure at least one log retention policy.

- **File Count Retention Policy.** Sets the number of log files you want the server to retain. The default file count is 10 logs. If the file count is set to 1, then the log will continue to grow indefinitely without being rotated.
- **Free Disk Space Retention Policy.** Sets the minimum amount of free disk space. The default free disk space is 500 MBytes.
- **Size Limit Retention Policy.** Sets the maximum size of the combined archived logs. The default size limit is 500 MBytes.
- **Never Delete Retention Policy.** Used in a rare event that does not require log deletion.

To Configure the Log Retention Policy

Use `dsconfig` to modify the log retention policy for the Sync Logger.

```
$ bin/dsconfig set-log-publisher-prop \  
  --publisher-name "File-Based Sync Logger" \  
  --set "retention-policy:Free Disk Space Retention Policy" \  
  --hostname host1 --port 1389 --bindDN "uid=admin,dc=example,dc=com" \  
  --bindPassword secret --no-prompt
```

Working with Administrative Alert Handlers

The UnboundID Synchronization Server provides mechanisms to send alert notifications to administrators when significant problems or events occur during processing, such as problems during server startup or shutdown. The Synchronization Server provides four types of alert handlers:

- **Admin Alert Log Publisher.** Sends administrative alerts to a Log Publisher. Specifically, the Admin Alert Access Log Publisher generates administrative alerts when certain criteria appear in the File-Based Access Log.
- **JMX™ Alert Handler.** Sends administrative alerts to clients using the JMX (Java Management Extensions) protocol. The Synchronization Server uses JMX for monitoring entries and requires that the JMX connection handler be enabled.
- **SMTP Alert Handler.** Sends administrative alerts to clients via email using SMTP (Simple Mail Transfer Protocol). The server requires that one or more SMTP servers be defined in the global configuration.

Configuring the JMX Connection Handler and Alert Handler

You can configure the JMX connection handler and alert handler using the `dsconfig` tool. Any user allowed to receive JMX notifications must have the `jmx-read`, `jmx-write`, and `jmx-notify` privileges. By default, these privileges are not granted to any users (including root users or global administrators). For security reasons, we recommend that you create a separate user account that does not have any other privileges but those listed. Although not shown in this section, you can configure the JMX connection handler and alert handler using `dsconfig` in interactive command-line mode, which is visible on the "Standard" object complexity menu.

To Configure the JMX Connection Handler

1. Use `dsconfig` to enable the JMX Connection Handler.

```
$ bin/dsconfig set-connection-handler-prop \  
  --handler-name "JMX Connection Handler" \  
  --set enabled:true --set listen-port:1689 \  
  --hostname host1 --port 1389 \  
  --bindDN "uid=admin,dc=example,dc=com" \  
  --bindPassword secret \  
  --no-prompt
```

2. Add a new non-root user account with the `jmx-read`, `jmx-write`, and `jmx-notify` privileges. This account can be added using the `ldapmodify` tool using an LDIF representation like:

```
dn: uid=jmx-user,dc=example,dc=com  
changetype: add  
objectClass: top  
objectClass: person  
objectClass: organizationalPerson  
objectClass: inetOrgPerson  
uid: jmx-user  
givenName: JMX  
sn: User  
cn: JMX User  
userPassword: password  
ds-privilege-name: jmx-read  
ds-privilege-name: jmx-write  
ds-privilege-name: jmx-notify
```

To Configure the JMX Alert Handler

Use `dsconfig` to configure the JMX Alert Handler.

```
$ bin/dsconfig --no-prompt set-alert-handler-prop \  
  --hostname host1 --port 1389 \  
  --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret \  
  --handler-name "JMX Alert Handler" --set enabled:true
```

Configuring the SMTP Alert Handler

By default, there is no configuration entry for an SMTP alert handler. To create a new instance of an SMTP Alert Handler, use the `dsconfig` tool.

To Configure the SMTP Alert Handler

By default, there is no configuration entry for an SMTP alert handler. To create a new instance of an SMTP Alert Handler, use the `dsconfig` tool.

```
$ bin/dsconfig create-alert-handler \
  --hostname host1 \
  --port 1389 \
  --bindDN "uid=admin,dc=example,dc=com" \
  --bindPassword secret \
  --no-prompt \
  --handler-name "SMTP Alert Handler" \
  --type smtp \
  --set enabled:true \
  --set "sender-address:alerts@example.com" \
  --set "recipient-address:administrators@example.com" \
  --set "message-subject:Sync Admin Alert %%alert-type%%" \
  --set "message-body:Administrative alert:\\n\\n%%alert-message%%"
```

Running the Status Tool

The Synchronization Server provides a command-line tool, `status`, that outputs the health of the Synchronization server. The `status` tool is a command-line utility that polls the current health of the server and displays summary information about the number of operations processed in the network. The tool provides different component categories as shown below.

TABLE 7-3. Status Information

Staus Component	Description
Server Status	Displays the server start time, operation status, number of connections (open, max, and total). --- Server Status --- Server Run Status: Started 17/May/2010:15:26:47.000 -0500 Operational Status: Available Open Connections: 6 Max Connections: 8 Total Connections: 24
Server Details	Displays the server details including host name, administrative users, install path, Sync Server version, and Java version. --- Server Details --- Host Name: sync1.example.com Administrative Users: cn=ADSync User Administrative Users: cn=Directory Manager Administrative Users: cn=IntraSync User Installation Path: /UnboundID-Sync Server Version: UnboundID Synchronization Server 3.0.0 Java Version: 1.6.0_20

TABLE 7-3. Status Information

Staus Component	Description
Connection Han- dlers	<p>Displays the state of the connection handlers including address, port, protocol and current state.</p> <pre> --- Connection Handlers --- Address:Port : Protocol : State -----:-----:----- 0.0.0.0:1689 : JMX : Disabled 0.0.0.0:636 : LDAPS : Disabled 0.0.0.0:7389 : LDAP : Enabled </pre>
Sync Topology	<p>Displays information about the connected Sync topology and any standby sync server instances.</p> <pre> --- Sync Topology --- Host:Port : Status : Priority Index -----:-----:----- sync1.example.com:7389 (this server) : Active : 1 sync2.example.com:8389 : Standby : 2 </pre>
Summary for sync pipe	<p>Displays the health status of each sync pipe configured on the topology, including current status, percent busy, changes detected, operations completed, number of operations processing, number of operations waiting, source unretrieved changes, failed operation attempts, source changes count. The important stats to view is the Source Unretrieved Changes and the Failed Op Attempts.</p> <pre> --- Summary for 'UBID1 to UBID2' Sync Pipe --- Summary Stat : Count -----:----- Started : true Percent Busy : 0 Changes Detected : 27299 Ops Completed Total : 26746 Num Ops In Flight : 0 Num Ops In Queue : 0 Source Unretrieved Changes : 10218 Failed Op Attempts : 5 Poll For Source Changes Count : 480 </pre>
Operations Com- pleted for the Sync Pipe	<p>Displays the completed operation statistics for the sync pipe, including the number of successful operations, out of scope, op type not synced, no change needed, entry already exists, no match found, multiple matches found, failed during mapping, failed at resource, unexpected exception, total operations.</p> <pre> --- Ops Completed for 'UBID1 to UBID2' Sync Pipe --- Op Result : Count -----:----- Success : 4559 Out Of Scope : 0 Op Type Not Synced : 0 No Change Needed : 22181 Entry Already Exists : 0 No Match Found : 0 Multiple Matches Found : 0 Failed During Mapping : 0 Failed At Resource : 0 Unexpected Exception : 0 Total : 26746 </pre>

TABLE 7-3. Status Information

Staus Component	Description
Sync Pipe Source Stats	<p>Displays the source statistics for the external server, including the current connection status, successful connect attempts, failed connect attempts, forced disconnects, unretrieved changed, failed to decode changelog entry.</p> <pre> --- Source Stats for 'UBID1 to UBID2' Sync Pipe --- Source Stat : Value ----- Is Connected : true Connected Server : ldap://sync1.example.com:1389 Successful Connect Attempts : 1 Failed Connect Attempts : 8 Forced Disconnects : 1 Root DSE Polls : 366 Unretrieved Changes : 10218 Entries Fetched : 26740 Failed To Decode Changelog Entry : 0 Ops Excluded By Modifiers Name : 0 Num Backtrack Batches Retrieved : 0 </pre>
Sync Pipe Source Stats	<p>Displays the destination statistics for the external server, including the current connection status, successful connect attempts, failed connect attempts, forced disconnects, unretrieved changed, failed to decode changelog entry.</p> <pre> --- Destination Stats for 'UBID1 to UBID2' Sync Pipe --- Destination Stat : Value ----- Is Connected : true Connected Server : ldap://sync1.example.com:3389 Successful Connect Attempts : 1 Failed Connect Attempts : 0 Forced Disconnects : 0 Entries Fetched : 26740 Entries Created : 0 Entries Modified : 4559 Entries Deleted : 0 </pre>
Admin Alerts	<p>Displays the 15 administrative alerts that were generated over the last 48 hour period. You can limit the number of displayed alerts using the <code>--maxAlerts</code> option. For example, <code>status --maxAlerts 0</code> suppresses any displayed alerts.</p> <pre> --- Administrative Alerts --- Severity : Time : Message ----- Informational : 17/May/2010 15:28:07 -0500 : A configuration change has been made in the Synchronization Server: [17/May/2010:15:28:07: -0500] conn=1 op=7 dn='cn=Directory Manager,cn=Root DNs,cn=config' authtype=[Simple]from=10.2.1.232 to=101.6.1.232 command='dsconfig createexternal-server -- server-name sync1.example.com:1389 --type unboundid --set server-host- name:sync1.example.com --set server-port:1389 --set 'bind-dn:cn=Sync User,cn=Root DNs,cn=config' --set 'password:AAB8sw4PRiOfLXNI4cu+5Saa'' Informational : 17/May/2010 15:26:47 -0500 : The Synchronization Server has started successfully Shown are info messages, warnings and errors from the past 48 hours </pre>

To Run the Status Tool

Go to the installation directory. Run the `status` command on the command line.

```
$ bin/status --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret
```

```

--- Server Status ---
Server Run Status:   Started 13/May/2010:17:51:33.000 -0500
Operational Status:  Available
Open Connections:    2
Max Connections:     2
Total Connections:   4
--- Server Details ---
Host Name:           server1.example.com

```

Administrative Users: uid=admin,dc=example,dc=com
Installation Path: UnboundID-Sync
Server Version: UnboundID Synchronization Server 3.0.0
Java Version: 1.6.0_20

```
--- Summary for 'sun-to-unboundid-sync-pipe' Sync Pipe ---
Summary Stat          : Count
-----:-----
Started               : true
Percent Busy          : 0
Changes Detected       : 2
Ops Completed Total    : 2
Num Ops In Flight      : 0
Num Ops Pending Queue Size : 0
Source Unretrieved Changes : 0
Failed Op Attempts     : 10
Poll For Source Changes Count : 1138
--- Ops Completed for 'sun-to-unboundid-sync-pipe' Sync Pipe ---
Op Result              : Count
-----:-----
Success                : 0
Out Of Scope           : 0
Op Type Not Synced     : 0
No Change Needed       : 0
Entry Already Exists   : 0
No Match Found         : 0
Multiple Matches Found : 0
Failed During Mapping   : 0
Failed At Resource     : 2
Unexpected Exception    : 0
Total                  : 2

--- Source Stats for 'sun-to-unboundid-sync-pipe' Sync Pipe ---
Source Stat            : Value
-----:-----
Is Connected           : true
Connected Server       : ldap://ds1.server.com:1389
Successful Connect Attempts : 1
Failed Connect Attempts : 14
Forced Disconnects     : 0
Root DSE Polls         : 1125
Unretrieved Changes    : 0
Entries Fetched        : 10
Failed To Decode Changelog Entry : 0
Ops Excluded By Modifiers Name : 0
Num Backtrack Batches Retrieved : -1

--- Destination Stats for 'sun-to-unboundid-sync-pipe' Sync Pipe ---
Destination Stat       : Value
-----:-----
Is Connected           : true
Connected Server       : ldap://ds3.server.com:3389
Successful Connect Attempts : 1
Failed Connect Attempts : 0
Forced Disconnects     : 0
Entries Fetched        : 10
Entries Created        : 0
Entries Modified       : 0
Entries Deleted        : 0

--- Connection Handlers ---
Address:Port : Protocol : State
-----:-----
0.0.0.0:1389 : LDAP      : Enabled
0.0.0.0:1689 : JMX       : Disabled
```



```
0.0.0.0:636 : LDAPS : Disabled
```

```
--- Administrative Alerts ---  
-No Errors or Warnings Less Than 48 Hours Old-
```

Monitoring the Synchronization Server

The UnboundID Synchronization Server exposes its monitoring information under the `cn=monitor` entry for easy access to its information. Administrators can use various means to monitor the server's information including the Synchronization Management Console, JConsole, LDAP command-line tools.

The following monitoring components are accessible:

TABLE 7-4. Synchronization Server Monitoring Components

Component	Description
Active Operations	Provides information about the operations currently being processed by the Synchronization Server. Shows the number of operations, information on each operation, and the number of active persistent searches.
Backend	Provides general information about the state of a Synchronization Server backend, including the backend ID, base DN(s), entry counts, entry count for the <code>cn=admin</code> data, writability mode, and whether it is a private backend. The following backend monitors are provided: <ul style="list-style-type: none">• adminRoot• ads-truststore• alerts• backup• config• monitor• schema• tasks• userRoot
Berkeley DB JE Environment	Provides information about the state of the Oracle Berkeley DB Java Edition database used by the Synchronization Server backend. Most of the statistics are obtained from the Berkeley DB JE and are not under the control of the Synchronization Server.
Client Connections	Provides information about all client connections to the Synchronization Server. The client connection information contains a name followed by an equal sign and a quoted value (e.g., <code>connID="15"</code> , <code>connect-Time="20100308223038Z"</code> , etc.)
Disk Space Usage	Provides information about the disk space available to various components of the Synchronization Server.
Connection Handler	Provides information about the available connection handlers on the Synchronization Server, which includes the LDAP and LDIF connection handlers. These handlers are used to accept client connections and to read requests and send responses to those clients.
General	Provides general information about the state of the Synchronization Server, including product name, vendor name, server version, etc.
JVM Stack Trace	Provides a stack trace of all threads processing within the JVM..

TABLE 7-4. Synchronization Server Monitoring Components

Component	Description
LDAP Connection Handler Statistics	Provides statistics about the interaction that the associated LDAP connection handler has had with its clients, including the number of connections established and closed, bytes read and written, LDAP messages read and written, operations initiated, completed, and abandoned, etc.
Processing Time Histogram	Categorizes operation processing times into a number of user-defined buckets of information, including the total number of operations processed, overall average response time (ms), number of processing times between 0ms and 1ms, etc.
System Information	Provides general information about the system and the JVM on which the Synchronization Server is running, including system host name, operation system, JVM architecture, Java home, Java version, etc.
Version	Provides information about the Synchronization Server version, including build ID, version, revision number, etc.
Work Queue	Provides information about the state of the Synchronization Server work queue, which holds requests until they can be processed by a worker thread, including the requests rejected, current work queue size, number of worker threads, number of busy worker threads, etc.

Profiling Server Performance Using the Periodic Stats Logger

The Synchronization Server ships with a built-in Periodic Stats Logger that is useful for profiling server performance for a given configuration. At specified intervals, the Periodic Stats logger writes server statistics to a log file in a comma-separated format (.csv), which can be ready by most spreadsheet applications. The logger has a negligible impact on server performance unless the `log-interval` property is set to a very small value.

You can control what statistics are logged and their verbosity. We recommend that you see what stats are available by experimenting with configurations or values. For example, you can make a configuration change in a test environment and see what stats are logged in the file.

To Enable the Periodic Stats Logger

By default, the Synchronization Server ships with the built-in 'Periodic Stats Logger' disabled. To enable it using the `dsconfig` tool or the web console, go to Plugins menu (available on the Advanced object menu), and then, select "Stats Logger."

1. Run `dsconfig` in interactive mode. Enter the LDAP or LDAPS connection parameters when prompted.

```
$ bin/dsconfig
```

2. Enter "o" to change to the Advanced Objects menu.

3. On the Synchronization Server configuration console menu, enter the number for Plugins. In this example, enter 33.
4. On the Plugin management menu, enter 3 to view and edit an existing plugin.
5. On the Plugin selection list, enter 9 for Stats Logger.
6. On the Periodic Stats Logger Plugin menu, enter 2 to set the **enabled** property to TRUE. When done, enter **f** to save and apply the configuration. The default logger will log information about the server every second to `<server-root>/logs/dsstats.csv`. If the server is idle, nothing will be logged, but this can be changed by setting the **suppress-if-idle** property to FALSE (**suppress-if-idle=false**).

```
>>>> Configure the properties of the Periodic Stats Logger Plugin
```

	Property	Value(s)
1)	description	Logs performance stats to a log file periodically.
2)	enabled	true
3)	log-interval	1 s
4)	collection-interval	200 ms
5)	suppress-if-idle	true
6)	header-prefix-per-column	false
7)	empty-instead-of-zero	true
8)	lines-between-header	50
9)	included-ldap-stat	active-operations, num-connections, op-count-and-latency, work-queue
10)	included-resource-stat	memory-utilization
11)	histogram-format	count
12)	histogram-op-type	all
13)	local-db-backend-info	basic
14)	replication-info	basic
15)	log-file	logs/dsstats.csv
16)	log-file-permissions	640
17)	append	true
18)	rotation-policy	Fixed Time Rotation Policy, Size Limit Rotation Policy
19)	retention-policy	File Count Retention Policy
?)	help	
f)	finish - apply any changes to the Periodic Stats Logger Plugin	
a)	hide advanced properties of the Periodic Stats Logger Plugin	
d)	display the equivalent dsconfig command lines to either re-create this object or only to apply pending changes	
b)	back	
q)	quit	

```
Enter choice [b]: f
```

7. Run the Synchronization Server. For example, if you are running in a test environment, you can run the **search-and-mod-rate** tool to apply some searches and modifications to the server. You can run **search-and-mod-rate --help** to see an example command.

- View the Stats log output at `<server-root>/logs/dsstats.csv`. You can open the file in a spreadsheet. The following image displays a portion of the file's output. On the actual file, you will need to scroll right for more stats.

Date	Time	All Ops	Avg Time (ms)	Search Ops	Avg Time (ms)	Add Ops	Avg Time (ms)	Delete Ops	Avg Time (ms)	Modify Ops	Avg Time (ms)	Bind Ops	Avg Time (ms)	Compare Ops	Avg Time (ms)	Num Current	Num New	Num Deleted	Active	Cleaner
2011/01/12	10:54:16.000	5300	0.715	2051	0.609	2640	0.821	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:17.000	5700	0.465	2053	0.276	2643	0.821	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:18.000	6190	0.446	3097	0.208	3099	0.683	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:19.000	4022	0.301	3132	0.192	3132	0.653	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:20.001	6207	0.432	3103	0.221	3104	0.653	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:21.000	6243	0.433	3121	0.221	3122	0.653	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:22.000	5737	0.652	2889	0.289	2889	1.015	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:23.000	6374	0.393	3187	0.199	3187	0.586	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:24.000	6330	0.437	3185	0.221	3185	0.654	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:25.000	6212	0.325	3197	0.138	3197	0.471	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:26.000	6352	0.333	3149	0.172	3153	0.493	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:27.001	3519	0.664	1161	0.153	1190	0.609	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:28.018	6309	0.403	3154	0.231	3155	0.576	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:29.000	6057	0.364	3038	0.244	3039	0.484	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:30.009	6210	0.380	3109	0.221	3110	0.551	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:31.001	6255	0.371	3130	0.199	3130	0.543	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:32.000	6410	0.44	3208	0.247	3210	0.634	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:33.001	5068	0.631	2535	0.413	2535	1.25	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:34.000	6191	0.410	3094	0.231	3097	0.607	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:35.001	6279	0.380	3140	0.232	3139	0.524	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:36.000	6413	0.467	3206	0.223	3207	0.711	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:37.000	5872	0.357	2937	0.184	2937	0.531	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:38.000	6357	0.396	3178	0.252	3179	0.54	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:39.000	5845	0.535	2923	0.287	2922	0.783	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:40.000	6204	0.38	3102	0.222	3102	0.538	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:41.000	6354	0.392	3179	0.219	3179	0.564	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:42.000	5270	0.756	2633	0.24	2637	1.272	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:43.000	6386	0.45	3178	0.234	3178	0.666	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:44.000	6332	0.387	3177	0.185	3178	0.589	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:45.000	6272	0.413	3136	0.207	3136	0.52	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:46.003	6399	0.46	3201	0.286	3190	0.635	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:47.001	5168	0.688	2582	0.388	2586	0.987	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:48.001	6395	0.368	3196	0.163	3199	0.569	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:49.000	6205	0.351	3134	0.18	3131	0.522	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:50.000	6294	0.454	3149	0.238	3145	0.67	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:51.000	5762	0.618	2677	0.294	2680	0.94	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:52.000	6294	0.412	3148	0.185	3146	0.639	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:53.001	6241	0.476	3112	0.254	3112	0.702	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:54.000	5959	0.404	2996	0.285	2996	0.722	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:55.000	6072	0.467	3034	0.266	3037	0.667	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:56.000	6037	0.36	3020	0.203	3021	0.578	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:57.001	4896	0.615	2447	0.389	2440	0.941	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:58.000	6012	0.367	3007	0.171	3005	0.565	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:54:59.000	4827	0.454	2413	0.184	2415	0.723	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:55:00.000	5140	0.428	2571	0.201	2568	0.657	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:55:01.000	5250	0.476	2625	0.175	2625	0.777	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:55:02.000	5240	0.464	2631	0.191	2631	0.657	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:55:03.000	4900	0.628	2450	0.427	2450	0.83	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:55:04.000	6300	0.272	3152	0.148	3148	0.398	0	0	0	0	0	0	0	0	10	0	0	0	0
2011/01/12	10:55:05.000	6260	0.374	3127	0.226	3133	0.512	0	0	0	0	0	0	0	0	10	0	0	0	0

To Configure Multiple Periodic Stats Loggers

Multiple Periodic Stats Loggers can be created to log different stats or to create a log at different intervals (such as logging cumulative operations statistics every hour). To create a new log, use the existing Stats Logger as a template to get reasonable settings, including rotation and retention policy.

- Run `dsconfig` by repeating steps 1–3 in “To Enable the Periodic Stats Logger” on page 176 (the previous procedure).
- From the Plugin management menu, enter 2 to create a new plugin.
- From the How to Create a New Plugin menu, enter `t` to use an existing plugin as a template.
- Enter 9 to select the existing stats logger as a template.
- Next, enter a descriptive name for the new stats logger. For this example, type `Stats Logger-10s`.
- Enter the log file path to the file. For this example, type `logs/dsstats2.csv`.
- On the Periodic Stats Logger Plugin menu, make any other change to your logger. For this example, change the `log-interval` to 10s, and the `suppress-if-idle` to `false`. When finished, enter `x` to save and apply the configuration.
- You should now see two loggers `dsstats.csv` and `dsstats2.csv` in the logs directory.

Configuring a Custom Periodic Stats Logger

You can add custom statistics based on any attribute in an entry under **cn=monitor** using the Custom Logged Stats object. This configuration object provides powerful controls for how monitor attributes are written to the log. For example, you can extract a value from a monitor attribute using a regular expression.

Besides allowing a straight passthrough of the values using the 'raw' statistic-type, you can configure attributes to be treated as a *counter* (where the interval includes the difference in the value since the last interval), an *average*, a *minimum*, or a *maximum* value held by the attribute during the specified interval. The value of an attribute can also be scaled by a fixed value or by the value of another monitor attribute.

Note

Custom third-party server extensions that were written using the UnboundID Server SDK can also expose interval statistics using the Periodic Stats Logger. The extension must first implement the SDK's MonitorProvider interface and register with the server. The monitor attributes produced by this custom MonitorProvider are then available to be referenced by a Custom Logged Stats object.

To illustrate how to configure a Custom Period Stats Logger, the following procedure reproduces the built-in "Consumer Total GB" column that shows up in the output when the **included-resource-stat** property is set to **memory-utilization** on the Periodic Stats Logger. The column is derived from the **total-bytes-used-by-memory-consumers** attribute of the **cn=JVM Memory Usage,cn=monitor** entry as follows:

```
dn: cn=JVM Memory Usage,cn=monitor
objectClass: top
objectClass: ds-monitor-entry
objectClass: ds-memory-usage-monitor-entry
objectClass: extensibleObject
cn: JVM Memory Usage
...
total-bytes-used-by-memory-consumers: 3250017037
```

To Configure a Custom Periodic Stats Logger Using dsconfig Interactive

1. Run **dsconfig** and enter the LDAP/LDAPS connection parameters when prompted.

```
$ bin/dsconfig
```

2. On the Synchronization Server configuration main menu (Advanced Objects menu), enter 8 for **Custom Logged Stats**.

```
>>>> UnboundID Synchronization Server configuration console main menu
```

```
What do you want to configure?
```

- | | |
|---|-----------------------------|
| 1) Access Control Handler | 18) Log Rotation Policy |
| 2) Alert Handler | 19) Monitor Provider |
| 3) Attribute Map | 20) Password Generator |
| 4) Attribute Mapping | 21) Password Policy |
| 5) Backend | 22) Password Validator |
| 6) Connection Criteria | 23) Plugin |
| 7) Connection Handler | 24) Plugin Root |
| 8) Custom Logged Stats | 25) Result Code Map |
| 9) Debug Target | 26) Root DN |
| 10) DN Map | 27) Root DSE Backend |
| 11) External Server | 28) Sync Class |
| 12) Global Configuration | 29) Sync Destination |
| 13) Global Sync Configuration | 30) Sync Destination Plugin |
| 14) LDAP SDK Debug Logger | 31) Sync Pipe |
| 15) Log Field Mapping | 32) Sync Pipe Plugin |
| 16) Log Publisher | 33) Sync Source |
| 17) Log Retention Policy | 34) Sync Source Plugin |
| o) 'Advanced' objects are shown - change this | |
| q) quit | |

```
Enter choice: 12
```

3. On the Custom Logged Stats menu, enter 2 to create a new custom logged stats.

```
>>>> Custom Logged Stats management menu
```

```
What would you like to do?
```

- | |
|--|
| 1) List existing Custom Logged Stats |
| 2) Create a new Custom Logged Stats |
| 3) View and edit an existing Custom Logged Stats |
| 4) Delete an existing Custom Logged Stats |
| b) back |
| q) quit |

```
Enter choice [b]: 2
```

4. Select the Periodic Stats Logger Plugin from the list if more than one is present on the system. If you only have one stats logger, press Enter to confirm that you want to use the existing plugin.

```
>>>> There is only one Periodic Stats Logger Plugin: 'Stats Logger'. Are you sure
that this is the correct one? (yes / no) [yes]:
```

5. Enter a descriptive name for the Custom Logged Stats. For this example, enter **Memory Usage**.

```
>>>> Enter a name for the Custom Logged Stats that you want to create: Memory Usage
```

6. From the monitor-objectclass property menu, enter the objectclass attribute to monitor. For this example, enter **ds-memory-usage-monitor-entry**. You can run **ldapsearch** using the base DN "**cn=JVM Memory Usage,cn=monitor**" entry to view the entry.

```
>>>> Configuring the 'monitor-objectclass' property
>>>> via creating 'Memory Usage' Custom Logged Stats

The objectclass name of the monitor entries to examine for generating these
statistics.

Syntax: monitor-entry-objectclass - The objectclass name of the monitor entries
to examine for generating these statistics.

Enter a value for the 'monitor-objectclass' property: ds-memory-usage-monitor-entry
```

7. Next, specify the attributes of the monitor entry that you want to log in the stats logger. In this example, enter **total-bytes-used-by-memory-consumers**, and then press Enter again to continue.

```
>>>> Configuring the 'attribute-to-log' property
>>>> via creating 'Memory Usage' Custom Logged Stats

Specifies the attributes on the monitor entries that should be included in the
output.

The full column name, which appears in the header, is built from the
header-prefix, the header-prefix-attribute, and the column name, which is
either determined automatically from the monitor entry attribute name or is
specified explicitly in the column-name property. Both the header-prefix, and
the header-prefix-attribute column names are optional.

Syntax: monitor-entry-objectclass - The name of the attribute to log.

Enter a value for the 'attribute-to-log' property: total-bytes-used-by-memory-consumers

Enter another value for the 'attribute-to-log' property [continue]:
```

8. Next, specify the type of statistics for the monitored attribute that will appear in the log file. In this example, enter 5 for **raw** statistics as recorded by the logger.

```
>>>> Configuring the 'statistic-type' property
>>>> via creating 'Memory Usage' Custom Logged Stats

Specifies the type of statistic to include in the output for each monitored
attribute.

The statistics listed here will be applied to each attribute listed in the
attribute-to-log list. The statistic is computed after all other optional
processing such as regular expression manipulation and value scaling.

If a single monitor entry includes attributes of different statistic-types,
then multiple Custom Logged Stats objects must be specified.

Select one or more values for the 'statistic-type' property:

1) average
2) counter
3) maximum
4) minimum
5) raw

?) help
c) cancel
q) quit

Enter one or more choices separated by commas [c]: 5
```

9. In the Custom Logged Stats menu, review the configuration. At this point, we want to set up a column name that lists the Memory Usage. Enter 6 to change the `column-name` property.

```
>>>> Configure the properties of the Custom Logged Stats
>>>> via creating 'Memory Usage' Custom Logged Stats
```

	Property	Value(s)
1)	description	-
2)	enabled	true
3)	monitor-objectclass	ds-memory-usage-monitor-entry
4)	include-filter	-
5)	attribute-to-log	total-bytes-used-by-memory-consumers
6)	column-name	-
7)	statistic-type	raw
8)	header-prefix	-
9)	header-prefix-attribute	-
10)	regex-pattern	-
11)	regex-replacement	-
12)	divide-value-by	-
13)	divide-value-by-attribute	-
14)	decimal-format	###
15)	non-zero-implies-not-idle	false
?)	help	
f)	finish - create the new Custom Logged Stats	
a)	hide advanced properties of the Custom Logged Stats	
d)	display the equivalent dsconfig arguments to create this object	
b)	back	
q)	quit	

```
Enter choice [b]: 6
```

10. Next, we want to add a specific label for the column name. Enter 2 to add a value, and then enter **Memory Consumer Total (GB)**, and press Enter again to continue.

```
>>>> Configuring the 'column-name' property
>>>> via creating 'Memory Usage' Custom Logged Stats
```

Optionally, specifies an explicit name for each column header instead of having these names automatically generated from the monitored attribute name.

The full column name, which appears in the header, is built from the header-prefix, the header-prefix-attribute, and the column name, which is either determined automatically from the monitor entry attribute name or is specified explicitly in the column-name property. The column-name, header-prefix, and header-prefix-attribute properties are optional.

The number of items in this property must line up with the attribute-to-log property. That is, the first value in this property will be used as the column name for the first value in the attribute-to-log property.

Syntax: STRING

```
Do you want to modify the 'column-name' property?
```

```
1) Leave undefined
2) Add one or more values

?) help
q) quit
```

```
Enter choice [1]: 2
```

```
Enter a value for the 'column-name' property [continue]: Memory Consumer Total (GB)
```

```
Enter another value for the 'column-name' property [continue]:
```


11. Confirm that you want to use the column-name value that you entered in the previous step, and then press Enter to use the value.

```
>>>> Configuring the 'column-name' property (Continued)
```

```
Do you want to modify the 'column-name' property?
```

- 1) Use the value: Memory Consumer Total (GB)
- 2) Add one or more values
- 3) Remove one or more values
- 4) Leave undefined
- 5) Revert changes
- ?) help
- q) quit

```
Enter choice [1]:
```

12. Next, we want to scale the Memory Consumer Totals by one gigabyte. On the Custom Logged Stats menu, enter 12 to change the `divide-value-by` property.

```
>>>> Configure the properties of the Custom Logged Stats
```

```
>>>> via creating 'Memory Usage' Custom Logged Stats
```

	Property	Value(s)
1)	description	-
2)	enabled	true
3)	monitor-objectclass	ds-memory-usage-monitor-entry
4)	include-filter	-
5)	attribute-to-log	total-bytes-used-by-memory-consumers
6)	column-name	Memory Consumer Total (GB)
7)	statistic-type	raw
8)	header-prefix	-
9)	header-prefix-attribute	-
10)	regex-pattern	-
11)	regex-replacement	-
12)	divide-value-by	-
13)	divide-value-by-attribute	-
14)	decimal-format	###
15)	non-zero-implies-not-idle	false
?)	help	
f)	finish - create the new Custom Logged Stats	
a)	hide advanced properties of the Custom Logged Stats	
d)	display the equivalent dsconfig arguments to create this object	
b)	back	
q)	quit	

```
Enter choice [b]: 12
```

13. On the "divide-value-by" property menu, enter 2 to change the value, and then enter 1073741824 (i.e., 1073741824 bytes = 1 gigabytes).

```
>>>> Configuring the 'divide-value-by' property
```

```
>>>> via creating 'Memory Usage' Custom Logged Stats
```

```
An optional floating point value that can be used to scale the resulting value.
```

```
This is used to scale the resulting value by a fixed amount. This is typically used to reduce large values to a more manageable length. For instance, free disk space could be represented in GB instead of just bytes. If a value for this property is specified, then it will be applied to the monitored attribute after any regular expression replacement is done and before it is scaled by the divide-value-by-attribute.
```

```
Syntax: divide-by-amount - A floating point value used to scale the result.
```

```
Do you want to modify the 'divide-value-by' property?
```

- 1) Leave undefined
- 2) Change the value

```
? ) help
q ) quit

Enter choice [1]: 2

Enter a value for the 'divide-value-by' property [continue]: 1073741824
```

14. On the Custom Logged Stats menu, review your configuration. When finished, enter **f** to save and apply the settings.

```
>>>> Configure the properties of the Custom Logged Stats
>>>> via creating 'Memory Usage' Custom Logged Stats

Property                                     Value(s)
-----
1) description                             -
2) enabled                                 true
3) monitor-objectclass                     ds-memory-usage-monitor-entry
4) include-filter                          -
5) attribute-to-log                       total-bytes-used-by-memory-consumers
6) column-name                             Memory Consumer Total (GB)
7) statistic-type                          raw
8) header-prefix                           -
9) header-prefix-attribute                 -
10) regex-pattern                          -
11) regex-replacement                      -
12) divide-value-by                        1073741824
13) divide-value-by-attribute              -
14) decimal-format                         #.##
15) non-zero-implies-not-idle              false

? ) help
f ) finish - create the new Custom Logged Stats
a ) hide advanced properties of the Custom Logged Stats
d ) display the equivalent dsconfig arguments to create this object
b ) back
q ) quit

Enter choice [b]: f

The Custom Logged Stats was created successfully
```

To Configure a Custom Periodic Stats Logger Using dsconfig Interactive

Use the **dsconfig** non-interactive command-line equivalent to create your custom stats logger. The following one-line command replicates the procedure in the previous section.

This command produces a column named "Memory Consumer Total (GB)" that contains the value of the **total-bytes-used-by-memory-consumers** attribute pulled from the entry with the **ds-memory-usage-monitor-entry** objectclass. This value is scaled by 1073741824 to get to a value represented in GBs.

```
$ bin/dsconfig create-custom-logged-stats --plugin-name "Stats Logger" \
--stats-name "Memory Usage" --type custom \
--set monitor-objectclass:ds-memory-usage-monitor-entry \
--set attribute-to-log:total-bytes-used-by-memory-consumers \
--set "column-name:Memory Consumer Total (GB)" --set statistic-type:raw \
--set divide-value-by:1073741824
```

8

Troubleshooting the Synchronization Server

Overview

The UnboundID Synchronization Server provides a highly-available background synchronization solution for all types of network configurations. However, problems can arise from issues in the Synchronization Server itself or from a supporting component, like the JVM, operating system, or hardware. The Synchronization Server provides tools to diagnose any problems quickly to determine the underlying cause and the best course of action to take towards a resolution.

This chapter provides information on how to perform this analysis to help ensure that the problem is resolved as quickly as possible. It targets cases in which the Synchronization Server is running on Solaris or Linux systems, but much of the information can be useful on other platforms.

This chapter presents the following information:

- [About Synchronization Troubleshooting](#)
- [Working with the Troubleshooting Tools](#)
- [Troubleshooting Process Flow](#)
- [Using the Sync Log](#)
- [Troubleshooting Sync Failures](#)
- [Working with the Collect Support Data Tool](#)

About Synchronization Troubleshooting

The majority of synchronization problems involve issues around the connection state of the external servers and the synchronization of the data between the two endpoints. Administrators should check if the Synchronization Server properly failed over to another endpoint instance if the connection was down on the highest priority external server. Further, if the main Synchronization Server is down for any reason, administrators should check if the Synchronization Server properly failed over to another Sync Server instance.

When troubleshooting synchronization information, administrators must determine if the DN and attribute mappings were properly configured and if the information is properly being syn-

chronized across the network. Typical scenarios involve checking for any entry sync failures and mapping issues.

Working with the Troubleshooting Tools

The Synchronization Server provides utilities to troubleshoot the synchronization state of your server and to locate the causes of any problems that have occurred. The following tools are available for diagnosing any problems and are located in the `bin` directory on UNIX or Linux systems, or the `bat` directory on Windows systems:

TABLE 8-1. Troubleshooting Tools

Tool	Description
status	The <code>status</code> tool provides a high-level view of the current operational state of the Synchronization Server and displays any recent alerts that have occurred in past 24 hours.
ldap-diff	The <code>ldap-diff</code> tool can be used to compare one or more entries across two server end-points to determine any data sync issues.
ldapsearch	The <code>ldapsearch</code> tool is used to get the full entries from two different servers if you want to review the exact content of an entry from each server.
logs	<p>The <code>logs</code> directory provides important logs that should be used to troubleshoot or monitor any issue with the Synchronization Server:</p> <ul style="list-style-type: none"> • Sync log provides information about the synchronization operations that occur within the server. Specifically, the Sync Log records all changes applied, detected or failed; dropped operations that were not synchronized; changes dropped due to being out of scope, or no changes needed for synchronization. The log also shows the entries that were involved in the synchronization process. • Failed Operations Log provides information about all operations that failed including replication problems from other servers. • Resync log provides summaries or details of synchronized entries and any missing entries in the Sync Destination. • Error log provides information about warnings, errors, or significant events that occur within the server. • Debug log can provide detailed information, if enabled, about processing performed by the server, including any exceptions caught during processing, detailed information about data read from or written to clients, and accesses to the underlying database. • Access loggers provide information about LDAP operations processed within the server. This log only applies to operations performed in the server. This includes configuration changes, searches of monitor data, and bind operations for authenticating administrators using the command-line tools and the UnboundID Sync Management console. <p>For more information, see “Managing Logging and Monitoring” on page 163.</p>
resync	The <code>resync</code> tool can be used to validate your sync classes and your data mappings from one endpoint to another (DN or attribute maps). The tool provides a dry-run mode that sees what could happen to data using an operation without actually affecting the data.
collect-support-data	The <code>collect-support-data</code> tool is used to aggregate the results of various support tools for the UnboundID Support team to diagnose. For more information, see “Working with the Collect Support Data Tool” on page 194.

Troubleshooting Process Flow

The general troubleshooting flow involves checking the status of the Synchronization Server, and then looking at the log files for information. The general flow is as follows:

1. **Run Status.** Run the `status` command to get the synchronization state information for your synchronization network.
2. **Check the Sync Log.** Depending on the nature of the problem, check the `sync log` file to diagnose any potential problems.
3. **Check the Failed Operations Log.** If you believe that the issue is data synchronization-related, then check the `failed-ops.log` to look at the cause of an issue.
4. **Check Sync Server Error Logs.** If the issue is a connectivity problem related to the source or destination servers, check the Synchronization Server error logs and the external server error logs.
5. **Check Endpoint Server Logs.** Look at the `access` and `error` logs on the source and destination servers.
6. **Run Collect-Support-Data.** If the Synchronization Server is experiencing issues that require assistance from UnboundID Support, then run the `collect-support-data` tool right away while the server is up and running to gather as much information as possible.

Using the Sync Log

The Sync log, located in the logs directory (`<server-root>/logs/sync`), provides useful troubleshooting information on the type of operation that was processed or completed. Most log entries provide the following common elements in their messages:

TABLE 8-2. Sync Logs Elements

Sync Log Element	Description
category	Indicates the type of operation, which will always be SYNC.
severity	Indicates the severity type of the message: INFORMATION, MILD_WARNING, SEVERE_WARNING, MILD_ERROR, SEVERE_ERROR, FATAL_ERROR, DEBUG, or NOTICE.
msgID	Specifies the unique ID number assigned to the message.
op	Specifies the operation number specific to sync.
changeNumber	Specifies the change number from the source server assigned to the modification.
replicationCSN	Specifies the replication change sequence number from the source server.
replicaID	Specifies the replica ID from the source server if there are multiple backend databases.

TABLE 8-2. Sync Logs Elements

Sync Log Element	Description
pipe	Specifies the sync pipe that was used to sync this operation.
msg	Displays the result of the sync operation.

Sync Log Example 1

The following example displays an informational message that a modification to an entry was detected on the source server.

```
$ tail -f logs/sync
```

```
[17/May/2010:15:46:19 -0500] category=SYNC severity=INFORMATION msgID=1893728293 op=14
changeNumber=15 replicationCSN=00000128A7E3C7D31E96000000F replicaID=7830 pipe="DS1
to DS2" msg="Detected MODIFY of uid=user.993,ou=People,dc=example,dc=com at ldap://
server1.example.com:1389"
```

Sync Log Example 2

The next example shows a successful synchronization operation that resulted from a MODIFY operation on the source server and synchronized to the destination server.

```
[18/May/2010:13:54:04 -0500] category=SYNC severity=INFORMATION msgID=1893728306
op=701 changeNumber=514663 replicationCSN=00000128ACC249A31E960007DA67 replicaID=7830
pipe="DS1 to DS2" class="DEFAULT" msg="Synchronized MODIFY of uid=user.698,ou=Peo-
ple,dc=example,dc=com at ldap://server1.example.com:1389 by modifying entry
uid=user.698,ou=People,dc=example,dc=com at ldap://server3.example.com:3389"
```

Sync Log Example 3

The next example shows a failed synchronization operation on a MODIFY operation from the source server that could not be synchronized on the destination server. The log displays the LDIF-formatted modification that failed, which came from a schema violation that resulted from an incorrect attribute mapping (telephoneNumber -> telephone) from the source to destination server.

```
[18/May/2010:11:29:49 -0500] category=SYNC severity=SEVERE_WARNING msgID=1893859389
op=71831 changeNumber=485590 replicationCSN=00000128AC3DE8D51E96000768D6 replic-
aID=7830 pipe="DS1 to DS2" class="DEFAULT" msg="Detected MODIFY of uid=user.941,ou=Peo-
ple,dc=example,dc=com at ldap://server1.example.com:1389, but failed to apply this
change because: Failed to modify entry uid=user.941,ou=People,dc=example,dc=com on des-
tination 'server3.example.com:3389'. Cause: LDAPException(resultCode=65(object class
violation), errorMessage='Entry uid=user.941,ou=People,dc=example,dc=com cannot be
modified because the resulting entry would have violated the server schema: Entry
uid=user.941,ou=People,dc=example,dc=com violates the Directory Server schema configu-
ration because it includes attribute telephone which is not allowed by any of the
objectclasses defined in that entry') (id=1893859386 ResourceOperationFailedExcep-
tion.java:125 Build revision=6226). Details: Source change detail:
```

```
dn: uid=user.941,ou=People,dc=example,dc=com
changetype: modify
replace: telephoneNumber
telephoneNumber: 027167170433915
-
```

```

replace: modifiersName
modifiersName: cn=Directory Manager,cn=Root DNs,cn=config
-
replace: ds-update-time
ds-update-time:: AAABKKw96NU=
  Equivalent destination changes:
dn: uid=user.941,ou=People,dc=example,dc=com
changetype: modify
replace: telephone
telephone: 818002279103216
  Full source entry:
dn: uid=user.941,ou=People,dc=example,dc=com
objectClass: person
... (more output)
Mapped destination entry:
dn: uid=user.941,ou=People,dc=example,dc=com
telephone: 818002279103216
objectClass: person
objectClass: inetOrgPerson
... (more output) ...

```

Troubleshooting Sync Failures

While many Synchronization Server issues are deployment-related and are directly affected by the hardware, software, and network structure used in the synchronization topology, most sync failures usually fall into one of three categories:

- **Entry Already Exists.** Indicates that when an add operation was attempted on the destination server, an entry with the same DN already exists.
- **No Match Found.** Indicates that a match was not found at the destination based on the current sync classes and correlation rules (i.e., DN and attribute mapping). When this value has a high count, it is likely that there were correlation rule problems. For example, use `bin/status` and look for "No Match Found".
- **Failure at Resource.** Indicates that some other error happened during the sync process that does not fall into the above categories. Typically, these errors are communication problems with a source or destination server.

Troubleshooting "Entry Already Exists" Failures

The `status` utility provides a comprehensive view of your synchronization network and displays the operation statistics to diagnose any potential problems with the Synchronization Server or the external servers. If you see that there is a count for the Entry Already Exists statistic using the `status` tool, then verify the problem in the sync log. For example, the `status` tool displays the following information:

```

      --- Ops Completed for 'DS1 to DS2' Sync Pipe ---
Op Result      : Count
-----
Success        : 0
Out Of Scope   : 0

```

```
Op Type Not Synced      : 0
No Change Needed       : 0
Entry Already Exists    : 1
No Match Found         : 1
Multiple Matches Found  : 0
Failed During Mapping   : 0
Failed At Resource      : 0
Unexpected Exception    : 0
Total                  : 2
```

Then verify the change by viewing the `<server-root>/logs/sync` file to see the specific operation, which could be due to someone manually adding the entry on the target server:

```
[18/May/2010:15:14:30 -0500] category=SYNC severity=SEVERE_WARNING msgID=1893859372
op=2 changeNumber=529277 replicationCSN=00000128AD0D9BA01E960008137D replicaID=7830
pipe="DS1 to DS2" class="DEFAULT" msg="Detected ADD of uid=user.1001,ou=Peo-
ple,dc=example,dc=com at ldap://server1.example.com:1389, but cannot create this entry
at the destination because an equivalent entry already exists at ldap://server3.exam-
ple.com:3389. Details: Search using [search-criteria dn: uid=user.1001,ou=Peo-
ple,dc=example,dc=com attrsToGet: [*, dn]] returned results;
[uid=user.1001,ou=People,dc=example,dc=com]. "
```

However, in the following example, a client attempted a MODIFY operation on an entry (uid=1234) on the source server, but the Sync Server could not find the entry on the destination server when it ran an initial search. The Sync Server then changed the MODIFY request to an ADD operation request to add the entry to the destination server. The ADD operation subsequently failed because an entry with the same DN already existed on the target server. In a case like this, the main problem could be due to an incorrectly-formed correlation rule (DN mapping) defined in the Sync Class used in the Sync Pipe.

```
[12/May/2010:00:00:53 -0500] category=SYNC severity=SEVERE_WARNING msgID=1893859389
op=2827888 changeNumber=5317162 replicationCSN=4bea4af3000b21140000 replic-
aID=8468,dc=example,dc=com pipe="DS1 to DS2" class="FullSync" msg="Detected MODIFY of
uid=1234,ou=People,dc=example,dc=com at ldap://server1.example.com:389, but failed to
apply this change because: Failed to create entry uid=1234,ou=People,dc=example,dc=com
on destination 'server1:389'. Cause: LDAPException(resultCode=entry already exists,
errorMessage='The entry uid=1234,ou=People,dc=dest,dc=com cannot be added because an
entry with that name already exists') (id=1893859385)"
```

To Diagnose an Entry Already Exists Problem

1. Assuming that a possible DN mapping is ill-formed, you should first run the `ldap-diff` utility to compare the entries on the source and destination servers. Then look at the `ldap-diff` results with your mapping rules to see why the original search did not find a match.

```
$ bin/ldap-diff --outputLDIF config-difference.ldif --baseDN "dc=example,dc=com" \
--sourceHost server1.example.com --targetHost server2.example.com \
--sourcePort 1389 --targetPort 3389 --sourceBindDN "cn=Directory Manager" \
--sourceBindPassword password --searchFilter "(uid=1234)"
```

2. Next, look at the destination server access logs to verify the search and filters it used to find the entry. Typically, you will find that your key correlation attributes are out-of-sync, which is why the search failed.
3. If the mapping rule attributes are out-of-sync, then you need to determine why that happened. Review your sync classes and mapping rules, and use the information from the

`ldap-diff` results to determine why a specific attribute may not be getting updated. Some questions to answer are as follows:

- Do you have more than one sync class that the operation could be matched with?
 - If you use an "include-base-dn" or "include-filter" in your mapping rules, does this exclude this operation by mistake?
 - If you use an attribute map, are your mappings correct? Usually, the cause in these type of messages are the destination mapping attribute settings. For example, if you define a set of correlation attributes as follows: `dn, mobile, accountNumber`. And the `accountNumber` changes for some reason, this will cause future operations on this entry to fail. To resolve this, you would either remove `accountNumber` from the rule, or add a second rule as follows: `dn, mobile`. The second rule will only be used if the search using the first set of attributes fails. In this case, the entry will be found and the `accountNumber` information will also be updated.
4. If you have deletes being synced, check to see if there was a previous delete of this entry that did not sync properly. In some cases, you will have to use simpler logic for deletes than other operations due to the available attributes in the changelogs. This scenario could cause an entry to not be deleted for some reason, which would cause an issue when a new entry with the same DN is added later. You can then use this information with your mapping rules to see why the original search did not find a match.
 5. Look at the destination directory server access logs to verify the search and filters it used to find the entry. Typically, you will find that your key attribute mappings are out-of-sync.

Troubleshooting "No Match Found" Failures

If you see that there is a count for the No Match Found statistic using the `status` tool, then verify the problem in the sync log. For example, the `status` tool displays the following information:

```

--- Ops Completed for 'DS1 to DS2' Sync Pipe ---
Op Result          : Count
-----
Success            : 0
Out Of Scope       : 0
Op Type Not Synced : 0
No Change Needed   : 0
Entry Already Exists : 1
No Match Found     : 1
Multiple Matches Found : 0
Failed During Mapping : 0
Failed At Resource  : 0
Unexpected Exception : 0
Total              : 2

```

Then verify the change by viewing the `<server-root>/logs/sync` file to see the specific operation:

```

[12/May/2010:10:30:45 -0500] category=SYNC severity=MILD_WARNING msgID=1893793952
op=4159648 changeNumber=6648922 replicationCSN=4beadaf4002f21150000 replicaID=8469-
ou=test,dc=example,dc=com pipe="DS1 to DS2" class="Others" msg="Detected DELETE of

```

```
'uid=1234,ou=test,dc=example,dc=com' at ldap://server1.example.com:389, but cannot
DELETE this entry at the destination because no matching entries were found at ldap://
server2.example.com:389. Details: Search using [search-criteria dn:
uid=1234,ou=test,dc=alu,dc=com filter: (nsUniqueId=3a324c60-5ddb11df-80ffe681-
717b93af) attrsToGet: [* , accountNumber, dn, entryuuid, mobile, nsUniqueId, object-
Class]] returned no results."
```

To Troubleshoot "No Match at Destination" Failures

1. First, test the search using the filter in the error message if displayed. For example, the sync log specifies "filter: (nsUniqueId=3a324c60-5ddb11df-80ffe681-717b93af) ". Use the **ldapsearch** tool to test the filter. Did this succeed? If yes, can you see anything in the your attribute mappings that would exclude this from working properly?
2. Next, test the search using the full DN as the base. For example, use **ldapsearch** with the full DN (**uid=1234,ou=People,dc=example,dc=com**). Did this succeed? If yes, then does the entry contain the attribute used in the mapping rule?
3. If the attribute is not in the entry, then determine if there is a reason why this attribute value was not synced in the first place. Look at the attribute mappings and the filters used in the sync classes.

Troubleshooting "Failure at Resource" Failures

If you see that there is a count for the "Failed at Resource" statistic using the **status** tool, then verify the problem in the sync log. For example, the **status** tool displays the following information:

```
--- Ops Completed for 'DS1 to DS2' Sync Pipe ---
Op Result          : Count
-----:-----
Success            : 0
Out Of Scope       : 0
Op Type Not Synced : 0
No Change Needed   : 0
Entry Already Exists : 0
No Match Found     : 0
Multiple Matches Found : 0
Failed During Mapping : 0
Failed At Resource  : 1
Unexpected Exception : 0
Total              : 1
```

You will see this stat after a change has been detected at the source in any of the following cases:

- If the fetch of the full source entry fails. In this case, the entry exists but there is a connection failure, server down, timeout, etc.
- If the fetch of the destination entry fails or if the modification to the destination fails for an exceptional reason (but not for cases "Entry Already Exists," "Multiple Matches Found," "No Match Found").

Verify the change by viewing the `<server-root>/logs/sync` file to see the specific operation. If you see any of the following resultCodes, then your server is seeing timeout errors:

- `resultCode=timeout: errorMessage=A client-side timeout was encountered while waiting 60000ms for a search response from server server1.example.com:1389`
- `resultCode=timeout: errorMessage=An I/O error occurred while trying to read the response from the server`
- `resultCode=server down: errorMessage=An I/O error occurred while trying to read the response from the server`
- `resultCode=server down: errorMessage=The connection to server server1.example.com:1389 was closed while waiting for a response to search request SearchRequest ...`
- `resultCode=object class violation: errorMessage='Entry device=1234,dc=example,dc=com violates the Directory Server schema configuration because it contains undefined object class`

To Troubleshoot "Failure at Destination" Failures

With the Failure at Destination timeout errors, you can look at the following settings in the Sync Server to see if they need adjustments:

1. **For External Server Properties.** Check the `connect-timeout` property. This property specifies the maximum length of time to wait for a connection to be established before giving up and considering the server unavailable.
2. **For the Sync Destination/Sync Source Properties.** Check the `response-timeout` property. This property specifies the maximum length of time that an operation should be allowed to be blocked while waiting for a response from the server. A value of zero indicates that there should be no client-side timeout. In this case, the server's default will be used.

```
$ bin/dsconfig --port 7389 --bindDN "cn=Directory Manager" --bindPassword password
list-external-servers --property connect-timeout --no-prompt
```

External Server	Type	connect-timeout	response-timeout
server1.example.com:389	sundsee-ds	10 s	-
server2.example.com:389	sundsee-ds	10 s	-
server3.example.com:389	unboundid-ds	10 s	-
server4.example.com:389	unboundid-ds	10 s	-

3. **For Sync Pipe Properties.** Check the `max-operation-attempts`, `retry-backoff-initial-wait`, `retry-backoff-max-wait`, `retry-backoff-increase-by`, `retry-backoff-percentage-increase`. These Sync Pipe Properties provide tuning parameters that are used in conjunction with the timeout settings. When a sync pipe experiences an error, then it will use these settings to determine how often and quickly it will retry the operation.

```
$ bin/dsconfig --port 7389 --bindDN "cn=Directory Manager" --bindPassword password
list-sync-pipes --property max-operation-attempts --property retry-backoff-initial-wait
--property retry-backoff-max-wait --property retry-backoff-increase-by --property
retry-backoff-percentage-increase --no-prompt
```

Working with the Collect Support Data Tool

The Synchronization Server itself provides a significant amount of information about its current state including any problems that it has encountered during processing. If a problem occurs, the final step is to run the `collect-support-data` tool in the `bin` directory. The tool aggregates all relevant support files into a zip file that administrators can send to UnboundID for analysis. The tool also runs data collector utilities, such as `jps`, `jstack`, and `jstat` plus other diagnostic tools for Solaris and Linux machines, and bundles the results in the zip file.

Table 7-1 presents a summary of the data collectors that the `collect-support-data` tool archives in zip format. If an error occurs during processing, you can rerun the specific data collector command and send the results to UnboundID.

TABLE 8-3. Data Collectors for the collect-support-data Tool

Data Collector	Description
config Directory	Archives the config directory, which holds the servers configuration settings.
logs Directory	Archives the logs directory.
db Directory	Archives the database directory.
lib Directory	Archives the lib directory, which holds the various libraries used by the directory server. The lib directory holds the core server code plus the jar file as well as Oracle Berkeley DB Java Edition (JE) libraries.
monitor.ldif	Archives the monitor entry.
rootdse.ldif	Archives the rootDSE entry.
Directory Listing	Archives the server root and tmp directories.
jps	Java Virtual Machine Process status tool. Reports information on the JVM.
jstack	Java Virtual Machine Stack Trace. Prints the stack traces of threads for the Java process.
jstat	Java Virtual Machine Statistics Monitoring Tool. Displays performance statistics for the JVM.
jinfo	Displays the Java configuration information for the Java process.
system-info	Prints system, machine, and operating system information. Solaris, Linux: <code>uname -a</code>
ps	Prints a snapshot of the current active processes. Solaris, Linux: <code>ps -elyZf</code>
zonename	Prints the name of the current zone. Solaris: <code>zonename</code> .
zoneadm	Prints the name of the current configured in verbose mode. Solaris: <code>zoneadm list -ev</code>
df	Prints the amount of available disk space for filesystems in 1024-byte units. Solaris: Linux. <code>df -k</code>
zpool	Print a zpool's status. Solaris: <code>zpool status</code> , <code>zpool iostat -v [interval][count]</code>
fmdump	Prints the log files managed by the Solaris Fault Manager. Solaris: <code>fmdump -eV</code>

TABLE 8-3. Data Collectors for the collect-support-data Tool

Data Collector	Description
prtconf	Prints the system configuration information. Solaris: prtconf
prtdiag	Prints the system diagnostic information. Solaris: prtdiag
etc-system	Prints the kernel parameters. Solaris: cat /etc/system
messages	Archives the messages generated by the directory server. Solaris: tail -100000c /var/adm/messages Linux: tail -100000c /var/log/messages The Messages data collector requires root or admin group privileges on Linux machines. UnboundID recommends that you run the tool with these privileges. If administrators running the tool do not have the privileges to run a particular data collector, they will get an error in the <code>collector.out</code> log or zip file.
netstat-int	Prints the state of all network interfaces. Solaris, Linux: netstat -i
netstat-sum	Prints the statistics for each protocol. Solaris, Linux: netstat -s
netstat-rn	Prints the kernel routing table. Solaris, Linux: netstat -rn
ifconfig	Prints information on all interfaces. Solaris: ifconfig -au Linux: ifconfig -a
uptime	Prints the time the server has been up and active. Solaris, Linux: uptime
dmesg	Prints the message buffer of the kernel. Solaris, Linux: dmesg
patchadd-p	Prints the patches added to the system. Solaris: patchadd -p
cpuinfo	Prints processor information. Linux: cat /proc/cpuinfo
meminfo	Prints the memory usage information. Linux: cat /proc/meminfo
vmstat	Prints information about virtual memory statistics. Solaris: vmstat
iostat	Prints disk I/O and CPU utilization information. Solaris: iostat -xn Linux (through sysstat RPM on Red Hat): iostat -x {reportIntervalSeconds} {reportCount}
mpstat	Prints performance statistics for all logical processors. Solaris: mpstat Linux (through sysstat RPM on Red Hat): mpstat -P ALL {reportIntervalSeconds} {reportCount}

TABLE 8-3. Data Collectors for the collect-support-data Tool

Data Collector	Description
pstack	Prints an execution stack trace on an active process specified by the pid. Solaris: pstack {pid} Linux: pstack {pid}
prstat	Prints out information on all active processes on the system. Solaris: prstat -L -n 1000 -p {pid} {reportIntervalSeconds} {reportCount}
top	Prints a list of active processes and how much CPU and memory each process is using. Linux: top -b -H -d {reportIntervalSeconds} -n {reportCount} -p {pid}

Available Tool Options

The `collect-support-data` tool has some important options that you should be aware of:

- **-R, --serverRoot {serverRootDir}**. Specifies the server root of the server from which to gather information. By default, this is the localhost server.
- **--noLdap**. Specifies that no effort should be made to collect any information over LDAP. This option should only be used if the server is completely unresponsive or will not start and only as a last resort.
- **--pid {pid}**. Specifies the ID of an additional process from which information is to be collected. This option is useful for troubleshooting external server tools and can be specified multiple times for each external server, respectively.
- **--sequential**. Use this option to diagnose “Out of Memory” errors. The tool collects data in parallel to minimize the collection time necessary for some analysis utilities. This option specifies that data collection should be run sequentially as opposed to in parallel. This action has the effect of reducing the initial memory footprint of this tool at a cost of taking longer to complete.
- **--reportCount {count}**. Specifies the number of reports generated for commands that support sampling (for example, vmstat, iostat, or mpstat). A value of 0 (zero) indicates that no reports will be generated for these commands. If this option is not specified, it defaults to 10.
- **--reportInterval {interval}**. Specifies the number of seconds between reports for commands that support sampling (for example, mpstat). This option must have a value greater than 0 (zero). If this option is not specified, it default to 1.

To Run the Collect Support Data Tool

1. Go to the installation directory.

```
$ cd UnboundID-Sync
```

2. Use the `collect-support-data` tool. Make sure to include the host, port number, bind DN, and bind password.

```
$ bin/collect-support-data --hostname 127.0.0.1 --port 389 \  
  --bindDN "uid=admin,dc=example,dc=com" --bindPassword secret \  
  --serverRoot /opt/UnboundID-Sync --pid 1234
```

3. Email the zip file to UnboundID Support.

Note	You might need to change the .zip file extension to something like .zap to get through your mail servers.
------	---

Index

A

- access log
 - described 164
- Active Directory
 - configuring password encryption 107
 - configuring sync 102–107
 - installing the Password Sync Agent 110
 - using SSL 103
 - using the Password Sync Agent 102
- Administration Guide
 - overview 1
- alert handlers
 - JMX
 - configuring 170
 - described 169
 - SMTP
 - configuring 171
 - described 169
- attribute mappings
 - configuring using dsconfig interactive 90–91
- attribute maps
 - configuring using dsconfig non-interactive 89
 - defined 88
 - described 15

C

- change flow 16
- changelog-deleted-entry-include-attribute 91
- config-audit
 - described 164
- Configuration Model
 - described 13
- connection handlers
 - JMX
 - configuring 170
- create-sync-pipe-config
 - logs 107
- create-sync-pipe-config utility 102
 - configuring 42–50
 - described 44

D

- DBSync

- attribute-synchronization-mode
 - property
 - described 126
- configure changelog 117
- configuring 113–143
- configuring attribute mappings 136
- configuring directory-to-database sync
 - pipe 125
- configuring sync classes 129
- configuring the database 122
- creating dn maps 132
- database-to-directory
 - general tips 141
- debugging the configuration 139
- directory-to-database
 - creating sync destination 127
 - creating sync source 127
 - general tips 125
- process overview 113
- running resync 138
- running status 140
- starting the sync pipe 139
- using dsconfig non-interactive 141

- deployments
 - advanced replication 4
 - migrations 4
- DN maps
 - configuring using dsconfig interactive 92
 - configuring using dsconfig non-interactive 93
 - described 15, 91
 - using regular expressions 92
 - using wildcards 92
- dsconfig utility
 - configuring 68

E

- error codes
 - ADMIN_LIMIT_EXCEEDED 97
 - ALIAS_DEREFERENCING_PROBLEM 97
 - BUSY 96
 - CANCELED 97
 - CLIENT_SIDE_CLIENT_LOOP 97
 - CLIENT_SIDE_DECODING_ERROR 97
 - CLIENT_SIDE_ENCODING_ERROR 97

- CLIENT_SIDE_LOCAL_ERROR 97
- CLIENT_SIDE_NO_MEMORY 97
- CLIENT_SIDE_REFERRAL_LIMIT_EXCEEDED 97
- CLIENT_SIDE_TIMEOUT 97
- CLIENT_SIDE_USER_CANCELLED 97
- CONNECT_ERROR 96
- DECODING_ERROR 97
- ENCODING_ERROR 97
- INTERACTIVE_TRANSACTION_ABORTED 97
- LOCAL_ERROR 97
- LOOP_DETECT 97
- max-backtrack-replication-latency 100
- max-failover-error-code-frequency 99
- max-operation-attempts 98
- NO_MEMORY 97
- OPERATIONS_ERROR 97
- OTHER 97
- PROTOCOL_ERROR 97
- response-timeout 98
- SERVER_CONNECTION_CLOSED 96
- TIME_LIMIT_EXCEEDED 98
- TIMEOUT 98
- UNAVAILABLE 96
- UNWILLING_TO_PERFORM 98

- Error log

- described 165
- External Server
 - described 14

F

- Failed Ops log
 - described 165
- fractional replication
 - configuring 93

G

- Groovy 1.7 119

J

- Java 6
 - recommended version 22
 - using with JAVA_HOME 24
 - using --XX-UseSSE42Intrinsics 37
- JDBC 3.0 or higher driver 117

- downloading 118
- JDBC Adapter Script
 - creating 118

L

- log publishers
 - creating 167
 - using interactive mode 167–168
- logs
 - access
 - described 164
 - config-audit
 - described 164
 - error
 - described 165
 - failed-ops
 - described 165
 - periodic stats 176
 - retention
 - configuring 169
 - Custom Retention Policy 169
 - described 168
 - File Count Retention Policy 169
 - Free Disk Space Retention Policy 169
 - Never Delete Retention Policy 169
 - Size Limit Retention Policy 169
 - rotation
 - configuring 168
 - Custom Rotation Policy 168
 - described 168
 - Fixed Time Rotation Policy 168
 - Never Rotate Policy 168
 - Size Limit Rotation Policy 168
 - Time Limit Rotation Policy 168
 - server.out
 - described 165
 - server.pid
 - described 165
 - server.status
 - described 165
 - setup.log
 - described 165
 - sync
 - described 165
 - sync-pipe-cfg.txt
 - described 165
 - tools
 - described 165
 - update.log
 - described 165

M

- meta-directories 2
 - limitations
 - cost 3
 - functionality 3
 - scalability 3

P

- Password Sync Agent. See UnboundID Password Sync Agent.

- Periodic Stats Logger
 - configuring custom 179
 - configuring multiple 178
 - description 176
 - enabling 176
- pre-deployment checklist 39
 - external servers 40
 - sync pipes 40
- prepare-endpoint-server utility
 - described 75

R

- realtime mode 12
 - scheduling a sync task 88
 - setting startpoints 86
 - starting globally 85
 - starting individual pipes 85
 - stopping globally 86
 - stopping individual pipes 86
- realtime-sync utility
 - options
 - scheduling a task 88
 - set-startpoint 85
 - rewinding 87
 - setting a changelog number 87
 - setting globally 86
 - start 84
 - stop 85
- replication
 - advanced 4, 8
 - filtered 6
 - fraction
 - configuring 93
 - fractional 5
 - local data 6
 - subtree 6
- resync mode
 - described 12
- resync utility
 - described 79
 - options
 - dry-run 79
 - logFilePath 79
 - logLevel 79
 - numPasses 79
 - secondsBetweenPass 80
 - sourceDNsFile 80
- populating an empty sync topology 81–82
- synchronizing a list of DN's 83
- testing attribute maps 80
- testing DN maps 80
- verifying the configuration 80

- revert-update utility
 - reverting an update 34

S

- Server SDK 114
- server.out log
 - described 165
- server.pid log
 - described 165
- server.status
 - described 165
- described 165
- setup.log
 - described 165
- Sync Class
 - described 14
- Sync Destination
 - described 14
- sync log
 - described 165
- Sync Management Console
 - configuring 50–67
 - installing 30–32
 - uninstalling 33
 - upgrading 35
- Sync Pipe
 - described 14
- Sync Source
 - described 14
- Synchronization Server
 - advanced replication 8
 - bulk resync mode 12
 - change flow 16
 - Configuration Model 13
 - configuring using create-sync-pipe-config 42–50
 - configuring using dsconfig 68
 - configuring using Sync Console 50–67
 - control flow scenarios 17
 - example of 19
 - filtered replication 6
 - subtree 6
 - fractional replication 5
 - local data 6
 - installing
 - redundant instance 35
 - using interactive command-line mode 23
 - using non-interactive command-line mode 23
 - overview 1
 - point-to-point bidirectional 9
 - pre-deployment checklist 39
 - external servers 40
 - sync pipes 40
 - preparing external servers 75
 - realistic test environments 9
 - realtime mode 12
 - synchronizing with Active Directory 8
 - using SSL/StartTLS
 - PKCS#11 token 24
 - PKCS#12 key store 24
- sync-pipe-cfg.txt
 - described 165
- Sync-through-Proxy
 - configuring 151–161
 - entry-balancing 149
 - features 145
 - Get Changelog Batch Request/Response Control 147
 - Get Server ID Request Control 147
 - how it works 146
 - proxy-server 160
 - tokens 148
 - use-changelog-batch-request 160

T

- tools log
 - described 165
- translate-ldif utility
 - described 82
 - populating an empty sync topology 82

U

- UnboundID LDAP SDK for Java
 - 2.0.0 119
- UnboundID Password Sync Agent
 - installation prerequisites 109
 - installing 108–110
 - manually configuring the registry 111
 - supported platforms 109
 - uninstalling 111
 - upgrading 110
- update utility
 - updating 34
- update.log
 - described 165

V

- virtual directories 3

