SymmetricDS User Guide

v3.3

Table of Contents

Preface

SymmetricDS is an open-source, web-enabled, database independent, data synchronization software application. It uses web and database technologies to replicate tables between relational databases in near real time. The software was designed to scale for a large number of databases, work across low-bandwidth connections, and withstand periods of network outages.

This User Guide introduces SymmetricDS and its uses for data synchronization. It is intended for users who want to be quickly familiarized with the software, configure it, and use its many features. This version of the guide was generated on 2013-02-08 at 14:42:17.

Chapter 1. Introduction

This User Guide will introduce both basic and advanced concepts in the configuration of SymmetricDS. By the end of this chapter, you will have a better understanding of SymmetricDS' capabilities, and many of its basic concepts.

1.1. System Requirements

SymmetricDS is written in Java 5 and requires a Java SE Runtime Environment (JRE) or Java SE Development Kit (JDK) version 5.0 or above.

Any database with trigger technology and a JDBC driver has the potential to run SymmetricDS. The database is abstracted through a *Database Dialect* in order to support specific features of each database. The following Database Dialects have been included with this release:

- MySQL version 5.0.2 and above

- Oracle version 10g and above

- PostgreSQL version 8.2.5 and above

- Sql Server 2005 and above

- Sql Server Azure

- HSQLDB 2.x

- H2 1.x

- Apache Derby 10.3.2.1 and above

- IBM DB2 9.5

- Firebird 2.0 and above

- Interbase 2009 and above

- Greenplum 8.2.15 and above

- SQLite 3 and above

See Appendix C, *Database Notes* (p. 133), for compatibility notes and other details for your specific database.

1.2. Overview

The following is an overview of how SymmetricDS works.

1.2.1. A Node is Born

SymmetricDS is a Java-based application that hosts a synchronization engine which acts as an agent for data synchronization between a single database instance and other synchronization engines in a network.

The SymmetricDS engine is also referred to as a *node* . SymmetricDS is designed to be able to scale out to many thousands of nodes. The database connection is configured by providing a database connection string, database user, and database password in a properties file. SymmetricDS can synchronize any table that is accessible by the database connection, given that the database user has been assigned the appropriate database permissions.
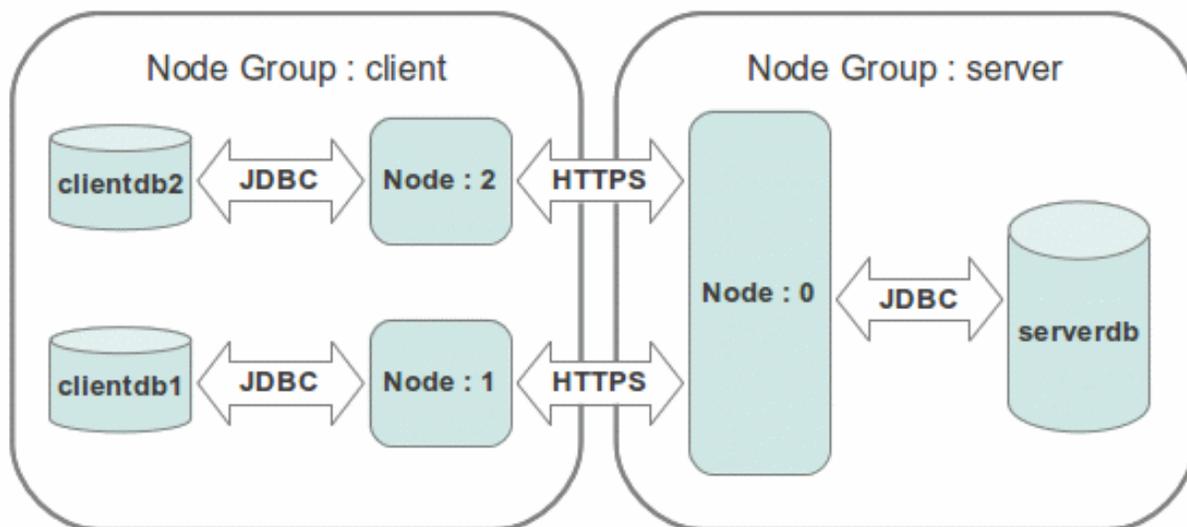
Figure 1.1. Simple Configuration

A SymmetricDS node is assigned an external id and a node group id. The external id is a meaningful, user-assigned identifier that is used by SymmetricDS to understand which data is destined for a given node. The node group id is used to identify groupings or tiers of nodes. It defines where the node fits into the overall node network. For example, one node group might be named "corporate" and represent an enterprise or corporate database. Another node group might be named "local_office" and represent databases located in different offices across a country. The external id for a "local_office" could be an office number or some other identifying alphanumeric string. A node is uniquely identified in a network by a node id that is automatically generated from the external id. If local office number 1 had two office databases and two SymmetricDS nodes, they would probably have an external id of "1" and node ids of "1-1" and "1-2."

SymmetricDS can be deployed in a number of ways. The most common option is to deploy it as a standalone process running as a service on your chosen server platform. When deployed in this manner

SymmetricDS can act as either a client, a multi-tenant server or both depending on where the SymmetricDS database fits into the overall network of databases. Although it can run on the same server as its database, it is not required to do so. SymmetricDS can also be deployed as a web application in an application server such as Apache Tomcat, JBoss Application Server, IBM WebSphere, or others.

SymmetricDS was designed to be a simple, approachable, non-threatening tool for technology personnel. It can be thought of and dealt with as a web application, only instead of a browser as the client, other SymmetricDS engines are the clients. It has all the characteristics of a web application and can be tuned using the same principles that would be used to tune user facing web applications.

1.2.2. Capturing Changes

Changes are captured at a SymmetricDS enabled database by database triggers that are installed automatically by SymmetricDS based on configuration settings that you specify. The database triggers record data changes in the DATA table. The database triggers are designed to be as noninvasive and as lightweight as possible. After SymmetricDS triggers are installed, changes are captured for any Data Manipulation Language (DML) statements performed by external applications. Note that no additional libraries or changes are needed by the applications that use the database and SymmetricDS does not have to be online for data to be captured.

Database tables that need to be replicated are configured in a series of SymmetricDS configuration tables. The configuration for the entire network of nodes is typically managed at a central node in the network, known as the registration server node. The registration server node is almost always the same node as the root node in a tree topology. When configuring "leaf" nodes, one of the start-up parameters is the URL of the registration server node. If the "leaf" node has not yet registered, it contacts the registration server and requests to join the network. Upon acceptance, the node downloads its configuration. After a node is registered, SymmetricDS can also provide an initial load of data before synchronization starts.

SymmetricDS will install or update its database triggers at start-up time and on a regular basis when a scheduled "sync triggers" job runs (by default, each night at midnight). The "sync triggers" job detects changes to your database structure or trigger configuration when deciding whether a trigger needs to be rebuilt. Optionally, the "sync triggers" job can be turned off and the database triggers DDL script can be generated and run by a DBA.

After changed data is inserted by the database trigger into the DATA table, it is batched and assigned to a node by the "router" job. Routing data refers to choosing the nodes in the SymmetricDS network to which the data should be sent. By default, data is routed to other nodes based on the node group. Optionally, characteristics of the data or of the target nodes can also be used for routing. A batch of data is a group of data changes that are transported and loaded together at the target node in a single database transaction. Batches are recorded in the OUTGOING_BATCH . Batches are node specific. DATA and OUTGOING_BATCH are linked by DATA_EVENT . The delivery status of a batch is maintained in OUTGOING_BATCH . After the data has been delivered to a remote node the batch status is changed to 'OK.'

1.2.3. Change Delivery

Data is delivered to remote nodes over HTTP or HTTPS. It can be delivered in one of two ways

depending on the type of transport link that is configured between node groups. A node group can be configured to push changes to other nodes in a group or pull changes from other nodes in a group. Pushing is initiated from the "push" job at the source node. If there are batches that are waiting to be transported, the pushing node will reserve a connection to each target node using an HTTP HEAD request. If the reservation request is accepted, then the source node will fully extract the data for the batch. Data is extracted to a memory buffer in CSV format until a configurable threshold is reached. If the threshold is reached, the data is flushed to a file and the extraction of data continues to that file. After the batch has been extracted, it is transported using an HTTP PUT to the target node. The next batch is then extracted and sent. This is repeated until the maximum number of batches have been sent for each channel or there are no more batches available to send. After all the batches have been sent for one push, the target returns a list of the batch statuses.

Pull requests are initiated by the "pull" job from at the target node. A pull request uses an HTTP GET. The same extraction process that happens for a "push" also happens during a "pull."

After data has been extracted and transported, the data is loaded at the target node. Similar to the extract process, while data is being received the data loader will cache the CSV in a memory buffer until a threshold is reached. If the threshold is reached the data is flushed to a file and the receiving of data continues. After all of the data in a batch is available locally, a database connection is retrieved from the connection pool and the events that had occurred at the source database are played back against the target database.

1.2.4. Channeling Data

Data is always delivered to a remote node in the order it was recorded for a specific channel. A channel is a user defined grouping of tables that are dependent on each other. Data that is captured for tables belonging to a channel is always synchronized together. Each trigger must be assigned a channel id as part of the trigger definition process. The channel id is recorded on SYM_DATA and SYM_OUTGOING_BATCH. If a batch fails to load, then no more data is sent for that channel until the failure has been addressed. Data on other channels will continue to be synchronized, however.

If a remote node is offline, the data remains recorded at the source database until the node comes back online. Optionally, a timeout can be set where a node is removed from the network. Change data is purged from the data capture tables by SymmetricDS after it has been sent and a configurable purge retention period has been reached. Unsent change data for a disabled node is also purged.

The default behavior of SymmetricDS in the case of data integrity errors is to attempt to repair the data. If an insert statement is run and there is already a row that exists, SymmetricDS will fall back and try to update the existing row. Likewise, if an update that was successful on a source node is run and no rows are found to update on the destination, then SymmetricDS will fall back to an insert on the destination. If a delete is run and no rows were deleted, the condition is simply logged. This behavior can be modified by tweaking the settings for conflict detection and resolution.

SymmetricDS was designed to use standard web technologies so it can be scaled to many clients across different types of databases. It can synchronize data to and from as many client nodes as the deployed database and web infrastructure will support. When a two-tier database and web infrastructure is maxed out, a SymmetricDS network can be designed to use N-tiers to allow for even greater scalability. At this point we have covered what SymmetricDS is and how it does its job of replicating data to many databases

using standard, well understood technologies.

1.3. Features

At a high level, SymmetricDS comes with a number of features that you are likely to need or want when doing data synchronization. A majority of these features were created as a direct result of real-world use of SymmetricDS in production settings.

1.3.1. Two-Way Table Synchronization

In practice, much of the data in a typical synchronization requires synchronization in just one direction. For example, a retail store sends its sales transactions to a central office, and the central office sends its stock items and pricing to the store. Other data may synchronize in both directions. For example, the retail store sends the central office an inventory document, and the central office updates the document status, which is then sent back to the store. SymmetricDS supports bi-directional or two-way table synchronization and avoids getting into update loops by only recording data changes outside of synchronization.

1.3.2. Data Channels

SymmetricDS supports the concept of *channels* of data. Data synchronization is defined at the table (or table subset) level, and each managed table can be assigned to a *channel* that helps control the flow of data. A channel is a category of data that can be enabled, prioritized and synchronized independently of other channels. For example, in a retail environment, users may be waiting for inventory documents to update while a promotional sale event updates a large number of items. If processed in order, the item updates would delay the inventory updates even though the data is unrelated. By assigning changes to the item tables to an *item* channel and inventory tables' changes to an *inventory* channel, the changes are processed independently so inventory can get through despite the large amount of item data. Channels are discussed in more detail in Section 3.5, Choosing Data Channels (p. 27) .

1.3.3. Change Notification

After a change to the database is recorded, the SymmetricDS nodes interested in the change are notified. Change notification is configured to perform either a *push* (trickle-back) or a *pull* (trickle-poll) of data. When several nodes target their changes to a central node, it is efficient to push the changes instead of waiting for the central node to pull from each source node. If the network configuration protects a node with a firewall, a pull configuration could allow the node to receive data changes that might otherwise be blocked using push. The frequency of the change notification is configurable and defaults to once per minute.

1.3.4. HTTP(S) Transport

By default, SymmetricDS uses web-based HTTP or HTTPS in a style called Representation State Transfer (REST). It is lightweight and easy to manage. A series of filters are also provided to enforce

authentication and to restrict the number of simultaneous synchronization streams. The `ITransportManager` interface allows other transports to be implemented.

1.3.5. Data Filtering and Rerouting

Using SymmetricDS, data can be filtered as it is recorded, extracted, and loaded.

- Data routing is accomplished by assigning a router type to a ROUTER configuration. Routers are responsible for identifying what target nodes captured changes should be delivered to. Custom routers are possible by providing a class implementing `IDataRouter` .

- In addition to synchronization, SymmetricDS is also capable of performing fairly complex transformations (see Section 4.8) of data as the synchronization data is loaded into a target database. The transformations can be used to merge source data, make multiple copies of source data across multiple target tables, set defaults in the target tables, etc. The types of transformation can also be extended to create even more custom transformations.

- As data changes are loaded in the target database, data can be filtered, either by a simple bean shell load filter (see Section 4.9 data-load-filter) or by a class implementing IDatabaseWriterFilter. You can change the data in a column, route it somewhere else, trigger initial loads, or many other possibilities. One possible use might be to route credit card data to a secure database and blank it out as it loads into a centralized sales database. The filter can also prevent data from reaching the database altogether, effectively replacing the default data loading process.

1.3.6. Transaction Awareness

Many databases provide a unique transaction identifier associated with the rows that are committed together as a transaction. SymmetricDS stores the transaction identifier, along with the data that changed, so it can play back the transaction exactly as it occurred originally. This means the target database maintains the same transactional integrity as its source. Support for transaction identification for supported databases is documented in the appendix of this guide.

1.3.7. Remote Management

Administration functions are exposed through Java Management Extensions (JMX) and can be accessed from the Java JConsole or through an application server. Functions include opening registration, reloading data, purging old data, and viewing batches. A number of configuration and runtime properties are available to be viewed as well.

SymmetricDS also provides functionality to send SQL events through the same synchronization mechanism that is used to send data. The data payload can be any SQL statement. The event is processed and acknowledged just like any other event type.

1.4. Origins

The idea of SymmetricDS was born from a real-world need. Several of the original developers were, several years ago, implementing a commercial Point of Sale (POS) system for a large retailer. The development team came to the conclusion that the software available for trickling back transactions to corporate headquarters (frequently known as the 'central office' or 'general office') did not meet the project needs. The list of project requirements made finding the ideal solution difficult:

- Sending and receiving data with up to 2000 stores during peak holiday loads.

- Supporting one database platform at the store and a different one at the central office.

- Synchronizing some data in one direction, and other data in both directions.

- Filtering out sensitive data and re-routing it to a protected database.

- Preparing the store database with an initial load of data from the central office.

The team ultimately created a custom solution that met the requirements and led to a successful project. From this work came the knowledge and experience that SymmetricDS benefits from today.

1.5. Why Database Triggers?

There are several industry recognized techniques to capture changing data for replication, synchronization and integration in a relational database.

- *Lazy data capture* queries changed data from a source system using some SQL condition (like a time stamp column).

- *Trigger-based data capture* installs database triggers to capture changes.

- *Log-based data capture* reads data changes from proprietary database recovery logs.

All three of these techniques have advantages and disadvantages, and all three are on the road map for SymmetricDS. At present time, SymmetricDS supports trigger-based data capture and basic lazy data capture. These two techniques were implemented first for a variety of reasons, not the least of which is that the majority of use cases that SymmetricDS targets can be solved using trigger-based and conditional replication in a way that allows for more database platforms to be supported using industry standard technologies. This fact allowed our developers' valuable time and energy to be invested in designing a product that is easy to install, configure and manage versus spending time reverse engineering proprietary and not well documented database log files.

Trigger-based data capture does introduce a measurable amount of overhead on database operations. The amount of overhead can vary greatly depending on the processing power and configuration of the database platform, and the usage of the database by applications. With nonstop advances in hardware and database technology, trigger-based data capture has become feasible for use cases that involve high data throughput or require scaling out.

Trigger-based data capture is easier to implement and support than log-based solutions. It uses well known database concepts and is very accessible to software and database developers and database

administrators. It can usually be installed, configured, and managed by application development teams or database administrators and does not require deployment on the database server itself.

1.6. Support

SymmetricDS is backed by JumpMind, Inc.

SymmetricDS is, and always will be, open source, which means free community support is available online, through the forums and the issue tracker. In a production environment, we have found that clients demand fast, more experienced help from the original architects and engineers — people who have the knowledge and experience to design, tune, troubleshoot, and shape future versions of the product.

To meet this demand, JumpMind provides Support Subscriptions designed to provide your organization with expert, dependable support from development to mission critical production support.

1.7. Whats New in SymmetricDS 3

SymmetricDS 3 builds upon the existing SymmetricDS 2.x software base and incorporates a number of architectural changes and performance improvements. If you are brand new to SymmetricDS, you can safely skip this section. If you have used SymmetricDS 2.x in the past, this section summarizes the key differences you will encounter when moving to SymmetricDS 3.

One optimization that effects both routing and data extraction is a change to the routing process to reuse batches across nodes if all of the data in the batches is going to be the same. SymmetricDS will automatically reuse batches if the default router is being used and there are NO inbound routers that have sync_on_incoming_batch turned on. If the same data is being sent to all nodes then a great deal of processing, during both routing and extraction, can be avoided. This is especially useful when data is being delivered to thousands of nodes. As a result of this change, the primary key of OUTGOING_BATCH has changed. This means that during an upgrade the table will be rebuilt.

Another optimization that effects data transport is the change to load batches as soon as they have been delivered to a target node. In 2.x all batches for a synchronization run were delivered, and then data was loaded. When errors occurred early on and there were several big batches or hundreds of batches to deliver, this was inefficient because all the batches were transported before the loading started.

Yet another optimization allows SymmetricDS to scale better when it is initiating communication with nodes. The pulling and pushing of data now happens from a configurable, but fixed size thread pool so that multiple nodes can be pulled and pushed to concurrently. This means that now, a centralized node can reach out to many child nodes in an efficient manner where in the past, the child nodes were relied upon to initiate communication.

The 2.x series allowed multiple nodes to be hosted in one standalone SymmetricDS instance. This feature (called multiServerMode) was off by default. In SymmetricDS 3 this feature is now the preferred mode of operation. It formalizes where properties file are configured and allows multiple nodes to be hosted on one JVM which saves on system resources.

SymmetricDS 3 introduces a long anticipated feature: Conflict Detection and Resolution. Please see Section 3.8, Planning Conflict Detection and Resolution (p. 31) for more information.

Transformations are now friendlier. They allow columns to be implied. This means that when configuring transformations, not all of the columns have to be specified which makes transformations much more maintainable.

An architectural change to the data loader subsystem allows the data loader to now be pluggable by channel. This will allow more efficient data loaders to be built if necessary. It will also make it straight forward to load data into non-relational data stores.

Several properties and extension points have been deprecated or renamed. Please see Appendix E, *Upgrading from 2.x* (p. 143) for a list of deprecated features.

Chapter 2. Quick Start Tutorial

Now that an overview of SymmetricDS has been presented, a quick working example of SymmetricDS is in order. This section contains a hands-on tutorial that demonstrates how to synchronize two databases with a similar schema between two nodes of SymmetricDS. This example models a retail business that has a central office database (which we'll call the "root" or "corp" node) and multiple retail store databases (which we'll call the "client" or "store" nodes). For the tutorial, we will have only one "client" or store node, as shown in Figure 2.1, although by the end of the tutorial you could extend the example and configure a second store, if desired.
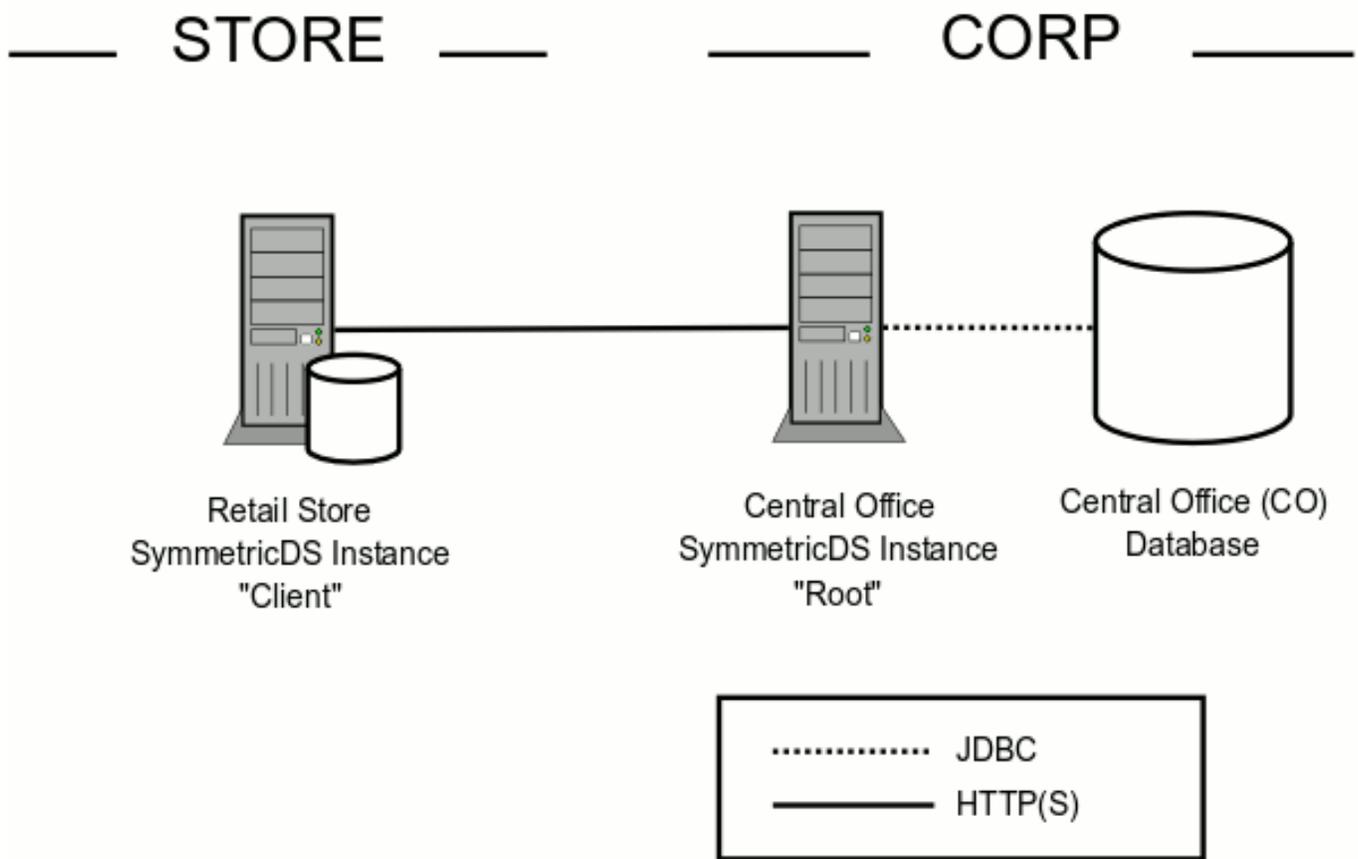
Figure 2.1. Simplified Two-Tier Retail Store Tutorial Example

For this tutorial, we will install two separate copies of SymmetricDS to represent the two different servers. One will represent the store server and one will represent the corp server. Each installed copy of SymmetricDS will be responsible for one database, and thus each copy acts as a single "node" in SymmetricDS terminology. This is the most common configuration of SymmetricDS - one installed copy of the software is responsible for one single database and represents one node. (Incidentally, there is also an option to configure a single installed copy of SymmetricDS to be responsible for both nodes. This is called "multi-homing" and will be discussed at the very end of the tutorial.) Since you are most likely going to run both SymmetricDS copies on a single machine, we will run the the two copies of SymmetricDS on two separate ports. We will use port 8080 for the corp server and 9090 for the store server, as shown in Figure 2.2.
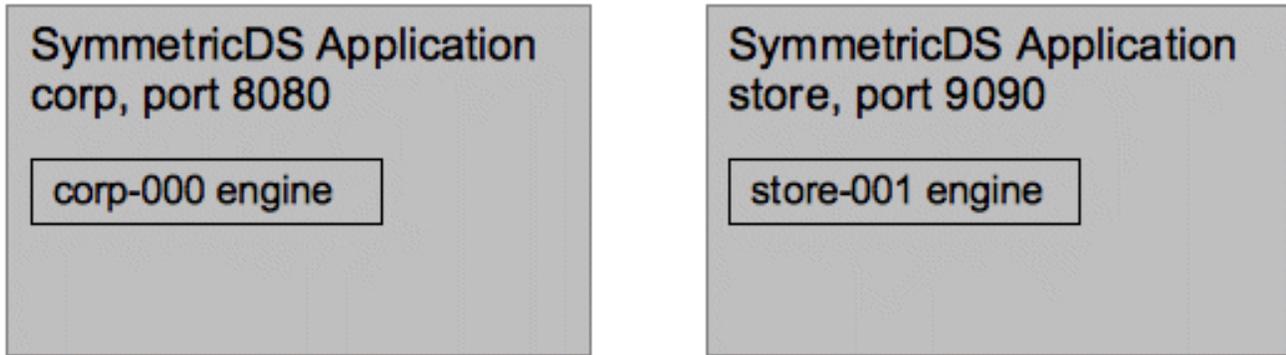
Figure 2.2. Two SymmetricDS applications - one for corp, one for store

Functionally, the corp SymmetricDS application will be responsible for capturing item data changes for the client, such as item number, description, and prices by store. The client SymmetricDS application (our store, specifically our first store, store # 001) captures sale transaction data changes for the root, such as time of sale and items sold. The pricing information is sent only to the specific store for which the price is relevant, thereby minimizing the amount of pricing data sent to each store. In other words, item pricing specific to store 001 will only be sent to the database for store 001 and not to store 002's database, for example.

The sample configuration has the client always initiating communication with the root node, which is a fairly common configuration. In this configuration, the client will attach to the root on a periodic basis to pull data from the server, and the client will also push captured changes to the root when changes are available.

Enough overview. Let's get started. We will next walk through:

1. Installing and configuring the two SymmetricDS applications,

2. Creating SymmetricDS configuration and sample tables as needed for the root and client, used to hold corp data and store data, respectively,

3. Creating sample retail data in the corp database,

4. Starting the SymmetricDS servers and registering the store with the corp node,

5. Sending an initial load of data to the store node,

6. Causing a data push and data pull operation, and

7. Verifying information about the batches that were sent and received.

2.1. Installing SymmetricDS

First, we will install two copies of the SymmetricDS software and configure it with your database connection information:

1. Download the symmetric-ds-3.x.x-server.zip file from http://www.symmetricds.org/

2. Create two directories to represent your two "machines". One will hold the corp installation of SymmetricDS and one to hold the store installation. For example, you could name the directories `sym-corp` and `sym-store001`, and we'll assume you used these names below (but feel free to update the steps below with your directory names as needed). Unzip the above zip file into both directories. This will create a `symmetric-ds-3.x.x` directory, which corresponds to the version you downloaded.

3. Properties files are use to store the minimal configuration information needed to start SymmetricDS. Copy the corp sample properties file to the corp engines directory, and the store one to the store engines directory. If you used the suggested directory names above, you would do the following copies:

    `samples/corp-000.properties` to `sym-corp/symmetric-ds-3.x.x/engines/`

    and

    `samples/store-001.properties` to `sym-store001/symmetric-ds-3.x.x/engines/`

4. Browse both properties files and explore the various settings. For exampl, notice that the root node is given a group id of corp, and that the store node is given a group id of store. Notice also that the root node is given an external id of 000, and the store node is given an external id of 001.

    Set the following properties in *both* properties files now present in the engines directories to specify how to connect to your particular database (the values below are just examples):

```
# The class name for the JDBC Driver
db.driver=com.mysql.jdbc.Driver

# The JDBC URL used to connect to the database
db.url=jdbc:mysql://localhost/sample

# The user to login as who can create and update tables
db.user=symmetric

# The password for the user to login as
db.password=secret
```

5. Next, set the following property in the `store-001.properties` file to specify where the root node can be contacted:

```
# The HTTP URL of the root node to contact for registration
registration.url=http://localhost:8080/sync/corp-000
```

Tip

Note that the URL for an engine is in the following general format:

```
http://{hostname}:{port}/sync/{engine.name}
```

where the engine.name portion of the URL comes from a node's properties file.

2.2. Creating and Populating Your Databases

Important

You must first create the databases for your root and client nodes using the administration tools provided by your database vendor. Make sure the name of the databases you create match the settings in the properties files you modified in the previous step.

See Appendix C, *Database Notes* (p. 133) for compatibility with your specific database.

First, create the sample tables in the *root* node database, load the sample data, and load the sample configuration, by doing the following:

1. Open a command prompt and navigate to the `samples` subdirectory of your *corp* SymmetricDS installation (for example, navigate to `sym-corp/symmetric-ds-3.x.x/samples`)

2. Create the sample tables for items, prices, and sales, in the root database by executing the following command:

    ../bin/dbimport --engine corp-000 --format XML create_sample.xml

    Note that the warning messages from the command are safe to ignore.

    Another quick comment about properties files. At startup, SymmetricDS looks for one or more properties files in the `engines` directory. Since we have specified a `--engine` parameter on the command line, it will look only for the specific file listed, namely `corp-000.properties`. Technically, the `--engine corp-000` part is optional in our particular tutorial example. Since there's only one properties file in the engines directory, SymmetricDS would just default to using that one file, after all. By including it, though, it will reduce errors while running the tutorial, because if you run the command from the wrong SymmetricDS installation, SymmetricDS will complain about the missing engines property file you specified.

3. Next, create the SymmetricDS-specific tables in the corp node database. These tables will contain the configuration for synchronization. The following command uses the auto-creation feature to create all the necessary SymmetricDS system tables.

    ../bin/symadmin --engine corp-000 create-sym-tables

4. Finally, load the sample item and transaction data and SymmetricDS configuration into the root node database by executing:

../bin/dbimport --engine corp-000 insert_sample.sql

> ⚠️ **Important**
> Please note that for MySQL, you will need to use the file `insert_sample_mysql.sql` in the above command. MySql uses back ticks (i.e., `) instead of double quotes (") for case-sensitive table and column names. The MySQL version of the file has the necessary change.

We have now created the corp database tables and populated them with our SymmetricDS configuration and sample data. Next, we will create the sample tables in the *store* node database to prepare it for receiving data.

1. Open a command prompt and navigate to the `samples` subdirectory of your *store #001* SymmetricDS installation (for example, navigate to `sym-store001/symmetric-ds-3.x.x/samples`)

2. Create the empty, sample tables in the client database by executing:

../bin/dbimport --engine store-001 --format XML create_sample.xml

Note that the warning messages from the command are safe to ignore. Also, feel free to review the `create_sample.xml` file to see what it contains.

Please verify *both* databases by logging in and listing the tables.

1. Find the item tables that sync from root to client (that is, from corp to store): `item` and `item_selling_price`.

2. Find the sales tables that sync from store to corp: `sale_transaction` and `sale_return_line_item`.

3. Find the SymmetricDS system tables, which have a prefix of "sym_", such as `sym_channel`, `sym_trigger`, `sym_router`, and `sym_trigger_router`.

4. Validate the corp item tables have sample data.

2.3. Starting SymmetricDS

Database setup and configuration for the tutorial is now complete. Time to put SymmetricDS into action. We will now start both SymmetricDS nodes and observe the logging output.

1. If they are not already open, open two command prompts and navigate to the samples directory of each installed SymmetricDS application (for example, navigate to `sym-corp/symmetric-ds-3.x.x/samples` and `sym-store001/symmetric-ds-3.x.x/samples`).

2.   From the corp samples directory, start the corp SymmetricDS by executing:

../bin/sym --engine corp-000 --port 8080

Upon startup for the first time, the corp node creates all the triggers that were configured by the sample configuration. It listens on port 8080 for synchronization and registration requests for the corp-000 engine.

3.   From the store001 samples directory, start the store SymmetricDS by executing:

../bin/sym --engine store-001 --port 9090

This command starts the store node server for the first time and uses the auto-creation feature to create the SymmetricDS system tables. It begins polling the corp node to try to register (it knows where to contact the corp node via the registration URL you configured in the previous steps). Since registration is not yet open, the store node receives an authorization failure (HTTP response of 403). We discuss registration next.

2.4. Registering a Node

When an unregistered node starts up, it will attempt to register with the node specified by the registration URL (which is our root node, in almost every case). The registration node centrally controls nodes on the network by allowing registration and returning configuration to a node once it has registered. In this tutorial, the registration node is the root node or 'corp' node, and it also participates in synchronization with other nodes.

So, we next need to open registration for the store node so that it may receive its initial load of data and so that it may receive and send data from and to the corp node. There are several ways to do this. We will use an administration feature available in SymmetricDS and issue a command on the corp node (since it is the node responsible for registration).

1.   Leave the corp and store SymmetricDS applications that you started in the previous step running, and open a command prompt and navigate to corp's samples subdirectory of your corp SymmetricDS installation.

Open registration for the store node server by executing:

../bin/symadmin --engine corp-000 open-registration store 001

The registration is now opened for a node group called "store" with an external identifier of "001". This information matches the settings in store-001.properties for the store node. In SymmetricDS, each node is assigned to a node group and is given an external ID that makes sense for the application. In this tutorial, we have retail stores that run SymmetricDS, so we named our node group representing stores as "store" and we used numeric identifiers for external ids starting with "001" ("000" is used to represent the corp node). More information about node groups will be covered in the next chapter.

2.   Watch the logging output of the store node to see it successfully register with the corp node. The

store is configured to attempt registration at a random time interval up to every minute. Once registered, the corp and store nodes are enabled for synchronization!

2.5. Sending an Initial Load

Next, we will send an initial load of data to our store, again using a node administration feature run on the corp node.

1. Open a command prompt and navigate to the corp `samples` subdirectory of the corp SymmetricDS installation. (Note that, in general, most system commands are issued using the corp server directly. All configuration, for example, is entered at the corp and synchronized to any clients.)

2. Send an initial load of data to the store node server by executing:

    ../bin/symadmin --engine corp-000 reload-node 001

    With this command, the server node queues up an initial load for the store node that will be sent the next time the store performs its pull. The initial load includes data for each table that is configured for synchronization (assuming its initial load order is a non-negative number, as discussed in later chapters).

3. Watch the logging output of both nodes to see the data transfer. The store is configured to pull data from the corp node every minute.

2.6. Pulling Data

Next, we will make a change to the item data in the central office corp node database (we'll add a new item), and observe the data being pulled down to the store.

1. Open an interactive SQL session with the *corp* database.

2. Add a new item for sale, with different prices at store 001 and store 002:

    insert into "item" ("item_id", "name") values (110000055, 'Soft Drink');

    insert into "item_selling_price" ("item_id", "store_id", "price") values (110000055, '001', 0.65); insert into "item_selling_price" ("item_id", "store_id", "price") values (110000055, '002', 1.00);

    > ⚠️ **Important**
    > Please note that for MySQL, you'll need to change the double quotes (") in the above commands to back ticks (i.e., `) since MySQL uses back ticks instead of double quotes for case-sensitive table and column names.

Once the statements are committed, the data change is captured by SymmetricDS and queued for the store node to pull.

3. Watch the logging output of both nodes to see the data transfer. The store is configured to pull data from the corp every minute.

4. Since `item_selling_price` is configured with a column match router in this tutorial, specific pricing data changes will be sent (or "routed", in SymmetricDS terms) only to nodes whose `store_id` matches the node's external ID (see Section 4.6.2, Router (p. 39) for details of the various routing options available). Verify that the new data arrives in the store database using another interactive SQL session. In this case, the first pricing row will be routed to store 001 only, and the second row would be routed to store 002 (which doesn't exist currently, so in this case the data change is recorded but routed nowhere and therefore discarded.)

2.7. Pushing Data

We will now simulate a sale at the store and observe how SymmetricDS pushes the sale transaction to the central office.

1. Open an interactive SQL session with the *store* node database.

2. Add a new sale to the store node database:

   insert into "sale_transaction" ("tran_id", "store_id", "workstation", "day", "seq") values (1000, '001', '3', '2007-11-01', 100);

   insert into "sale_return_line_item" ("tran_id", "item_id", "price", "quantity") values (1000, 110000055, 0.65, 1);

   Once the statements are committed, the data change is captured and queued for the store node to push.

3. Watch the logging output of both nodes to see the data transfer. The store is configured to push data to the corp node every minute.

2.8. Verifying Outgoing Batches

Now that we have pushed and pulled data, we will demonstrate how you can obtain information about what data has been batched and sent. A batch is used for tracking and sending one or more data changes to a given node. The sending node creates a batch and the receiving node receives and then acknowledges it.

In addition, in SymmetricDS tables are grouped into data "Channels" for, among many reasons, the purpose of allowing different types of data to synchronize even when other types of data might be in

error. For example, if a batch for a given channel is in error, that batch will be retried with each synchronization for that channel until the batch is no longer in error. Only after the batch is no longer in error will additional batches for that channel be sent. In this way, the order of the data changes that have occurred for a given channel are guaranteed to be sent to the destination in the same order they occurred on the source. Batches on a channel without batch errors, however, will not be blocked by the existence of a batch in error on a different channel. In this way, data changes for one channel are not blocked by errors present in another channel.

Explore the outgoing batches by doing the following:

1.   Open an interactive SQL session with either the corp or store database.

2.   Verify that the data change you made was captured:

   select * from sym_data order by data_id desc;

   Each row represents a row of data that was changed. Data Ids are sequentially increasing, so one of the most recent (highest) data ids should be related to your data insert SQLs. The `event_type` is "I" for insert, "U" for update", or "D" for delete. For insert and update, the captured data values are listed in `row_data`. For update and delete, the primary key values are listed in `pk_data`.

3.   Verify that the data change was included in a batch, using the data_id from the previous step:

   select * from sym_data_event where data_id = ?;

   Batches are created based on the needed routing to nodes as part of a background job, called the Route Job. As part of the Route Job, the data change is assigned to a batch using a `batch_id` which is used to track and synchronize the data. The links between batches and data are managed by this `sym_data_event` table.

4.   Verify that the data change was batched, sent to the destination, and acknowledged, using the `batch_id` from the previous step:

   select * from sym_outgoing_batch where batch_id = ?;

   Batches initially have a status of "NE" when they are new and not yet sent to a node. Once a receiving node acknowledges the batch, the batch status is changed to a status of "OK" for success or "ER" for error (failure). If the batch failed, the `error_flag` on the batch is also sent to 1, since the status of a batch that failed can change as it's being retried.

Understanding these three tables, along with a fourth table discussed in the next section, is key to diagnosing any synchronization issues you might encounter. As you work with SymmetricDS, either when experimenting or starting to use SymmetricDS on your own data, spend time monitoring these tables to better understand how SymmetricDS works. Exploring and solving any synchronization issues is discussed later in far greater depth in Section 6.1, Solving Synchronization Issues (p. 81).

2.9. Verifying Incoming Batches

The receiving node keeps track of the batches it acknowledges and records statistics about loading the data. Duplicate batches are skipped by default, but this behavior can be changed with the `incoming.batches.skip.duplicates` runtime property.

Explore incoming batches by doing the following:

1. Open an interactive SQL session with either the corp or store database.

2. Verify that the batch was received and acknowledged, using a batch_id from the previous section:

    select * from sym_incoming_batch where batch_id = ?;

    A batch represents a collection of changes loaded by the node. The sending node that created the batch is recorded, and the batch's status is either "OK" for success or "ER" for error.

2.10. Multi-Homing

Our Quick Start Tutorial is finished. We have succesfully set up and performed synchronization between two databases. However, we did want to go back and discuss one of the first steps you did in the tutorial; namely, the step where you installed two copies of SymmetricDS when doing the tutorial. Feel free to skip this section until a later time if you wish.

In the example above, we placed one properties file in the engines directory of each installed SymmetricDS application. When SymmetricDS was started in the examples above, the application initialized, and then created a "SymmetricDS engine" based on the provided property file (again, each engine serves as a SymmetricDS node and is responsible for one particular database).

In reality, though, the SymmetricDS application is capable of starting more than one engine at a time. When SymmetricDS starts, it looks in the `engines` directory for any files that end in `.properties`. It will start a SymmetricDS engine for each and every property file found. The `--engine` command line prompt is an override for this and will cause SymmetricDS to only start the one engine as specified on the command line. In cases where a single SymmetricDS application is running multiple engines, this is known as a "multi-homed" SymmetricDS application, and the feature, in general, is known as "multi-homing".
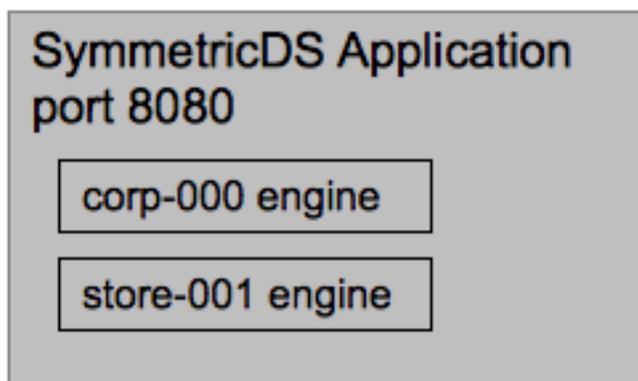
Figure 2.3. Multi-Homed version of Tutorial

So, for our tutorial above, how could we have "multi-homed" the corp and store such that we only had to install a single copy of SymmetricDS? It's fairly simple. The following changes to the above would be needed:

1. Install a single copy of the SymmetricDS software instead of two copies. You no longer need a directory to represent the two machines.

2. Instead of copying a single property file from `samples` to each separate `engines` directory, copy both files to just the one engines directory.

3. All commands in the tutorial are run from the one single `samples` directory.

4. When you start SymmetricDS, you will no longer specify a specific engine, as you want both engines to start. The command, still run from the `samples` directory, would now be:

   ../bin/sym --port 8080

   Note that we are no longer using port 9090, by the way. SymmetricDS now listens on port 8080 for traffic relevant to both the store and corp engines.

5. Other than starting the server, all other commands you executed will still have the `--engine` specification, since you are addressing the command itself to a specific node (engine) of SymmetricDS to open registration, set up the corp server to issue an initial load to store, etc.

Chapter 3. Planning

In the previous Chapter we presented a high level introduction to some basic concepts in SymmetricDS, some of the high-level features, and a tutorial demonstrating a basic, working example of SymmetricDS in action. This chapter will focus on the key considerations and decisions one must make when planning a SymmetricDS implementation. As needed, basic concepts will be reviewed or introduced throughout this Chapter. By the end of the chapter you should be able to proceed forward and implement your planned design. This Chapter will intentionally avoid discussing the underlying database tables that capture the configuration resulting from your analysis and design process. Implementation of your design, along with discussion of the tables backing each concept, is covered in Chapter 4, *Configuration* (p. 33) .

When needed, we will rely on an example of a typical use of SymmetricDS in retail situations. This example retail deployment of SymmetricDS might include many point-of-sale workstations located at stores that may have intermittent network connection to a central location. These workstations might have point-sale-software that uses a local relational database. The database is populated with items, prices and tax information from a centralized database. The point-of-sale software looks up item information from the local database and also saves sale information to the same database. The persisted sales need to be propagated back to the centralized database.

3.1. Identifying Nodes

A *node* is a single instance of SymmetricDS. It can be thought of as a proxy for a database which manages the synchronization of data to and/or from its database. For our example retail application, the following would be SymmetricDS nodes:

- Each point-of-sale workstation.

- The central office database server.

Each node of SymmetricDS can be either embedded in another application, run stand-alone, or even run in the background as a service. If desired, nodes can be clustered to help disperse load if they send and/or receive large volumes of data to or from a large number of nodes.

Individual nodes are easy to identify when planning your implementation. If a database exists in your domain that needs to send or receive data, there needs to be a corresponding SymmetricDS instance (a node) responsible for managing the synchronization for that database.

3.2. Organizing Nodes

Nodes in SymmetricDS are organized into an overall node network, with connections based on what data needs to be synchronized where. The exact organization of your nodes will be very specific to your synchronization goals. As a starting point, lay out your nodes in diagram form and draw connections between nodes to represent cases in which data is to flow in some manner. Think in terms of what data is needed at which node, what data is in common to more than one node, etc. If it is helpful, you could also show data flow into and out of external systems. As you will discover later, SymmetricDS can publish data changes from a node as well using JMS.

Our retail example, as shown in Figure 3.1 , represents a tree hierarchy with a single central office node connected by lines to one or more children nodes (the POS workstations). Information flows from the central office node to an individual register and vice versa, but never flows between registers.
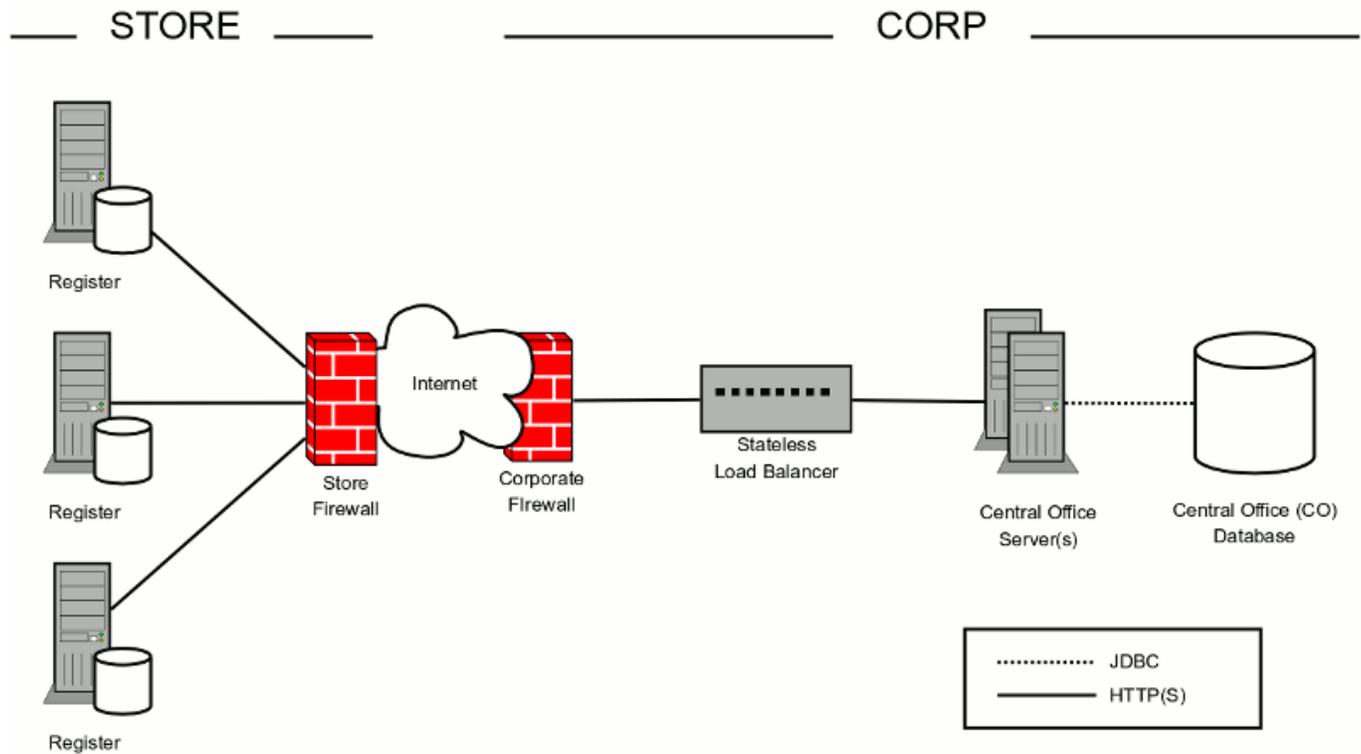
Figure 3.1. Two-Tier Retail Store Deployment Example

More complex organization can also be used. Consider, for example, if the same retail example is expanded to include store *servers* in each store to perform tasks such as opening the store for the day, reconciling registers, assigning employees, etc. One approach to this new configuration would be to create a three-tier hierarchy (see Figure 3.2). The highest tier, the centralized database, connects with each store server's database. The store servers, in turn, communicate with the individual point-of-sale workstations at the store. In this way data from each register could be accumulated at the store server, then sent on to the central office. Similarly, data from the central office can be staged in the store server and then sent on to each register, filtering the register's data based on which register it is.
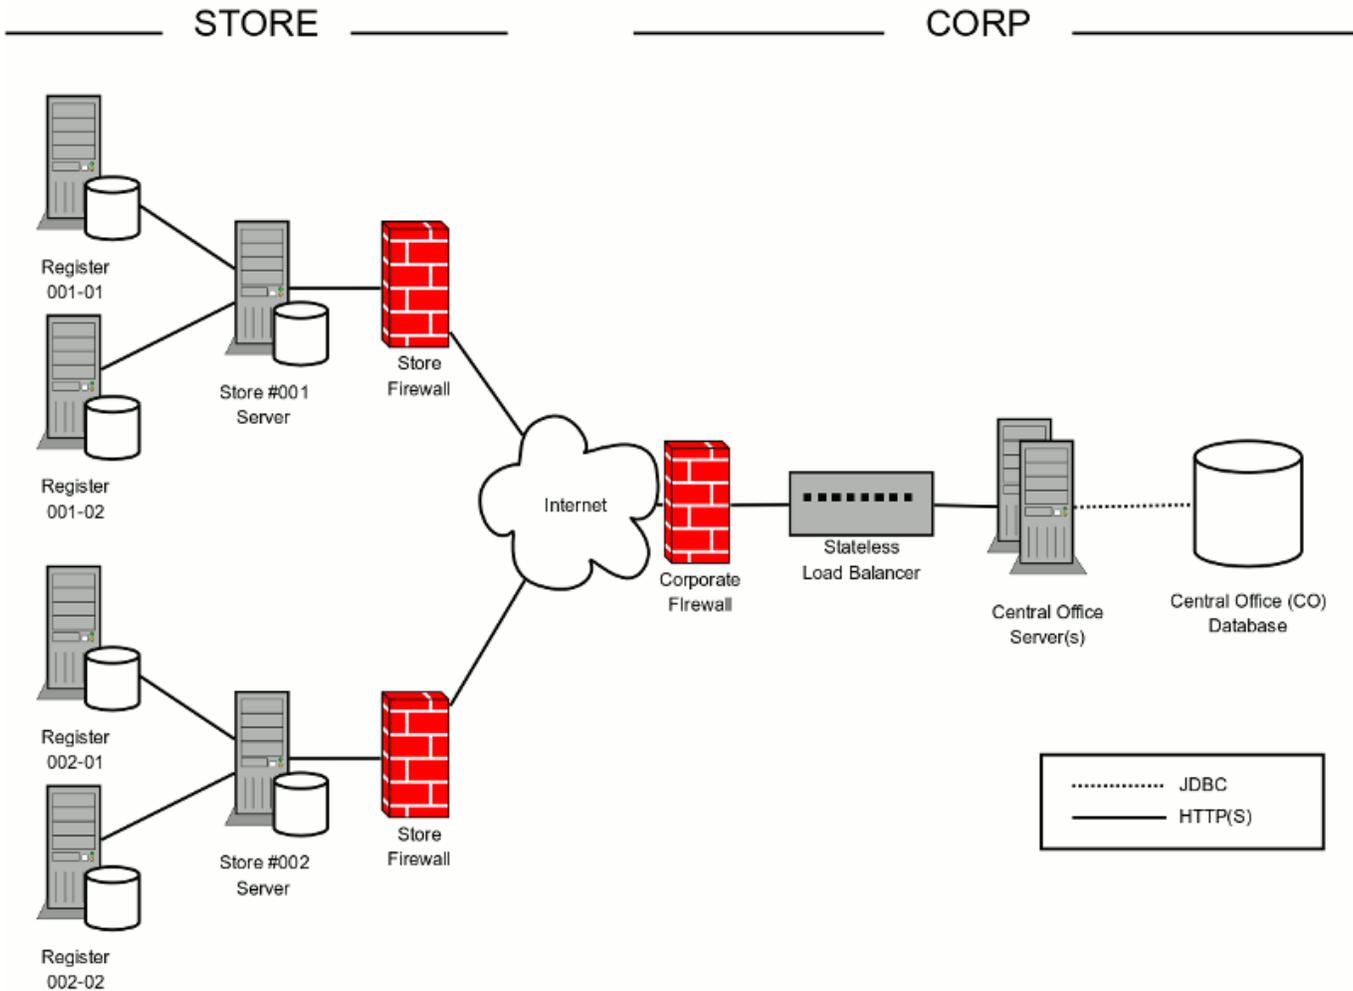
Figure 3.2. Three-Tier, In-Store Server, Retail Store Deployment Example

One final example, show in Figure 3.3 , again extending our original two-tier retail use case, would be to organize stores by "region" in the world. This three tier architecture would introduce new regional servers (and corresponding regional databases) which would consolidate information specific to stores the regional server is responsible for. The tiers in this case are therefore the central office server, regional servers, and individual store registers.
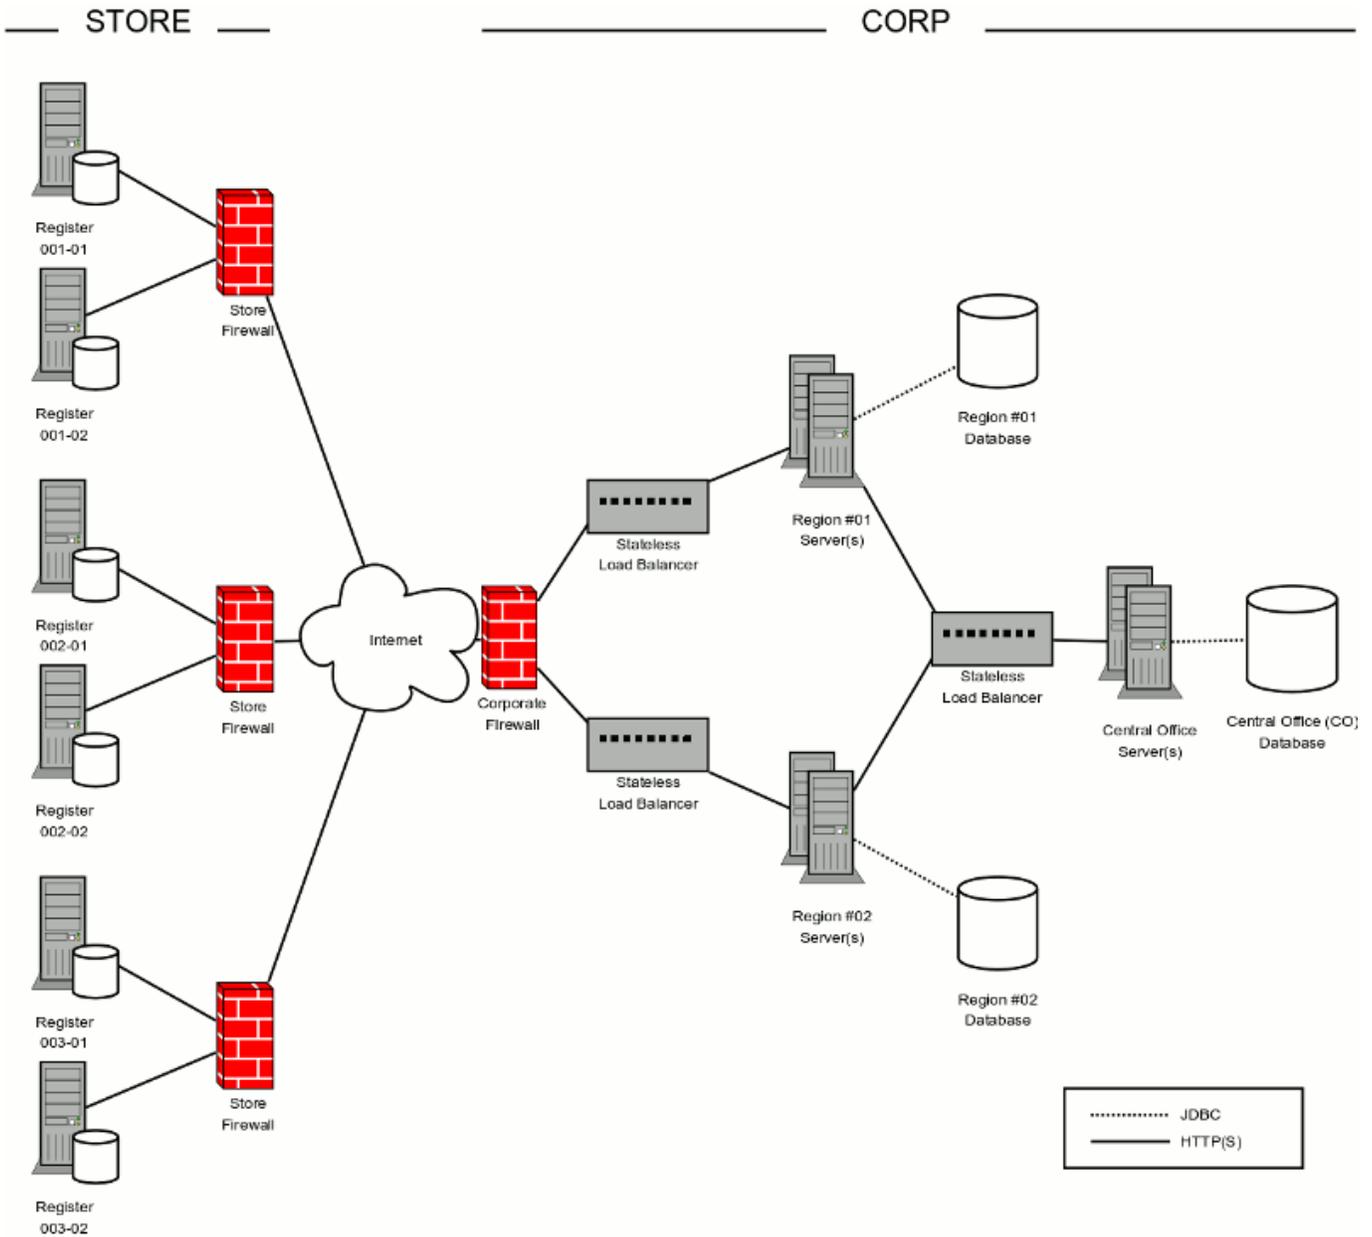
Figure 3.3. Three-Tier, Regional Server, Retail Store Deployment Example

These are just three common examples of how one might organize nodes in SymmetricDS. While the examples above were for the retail industry, the organization, they could apply to a variety of application domains.

3.3. Defining Node Groups

Once the organization of your SymmetricDS nodes has been chosen, you will need to *group* your nodes based on which nodes share common functionality. This is accomplished in SymmetricDS through the concept of a *Node Group* . Frequently, an individual tier in your network will represent one Node Group. Much of SymmetricDS' functionality is specified by Node Group and not an individual node. For

example, when it comes time to decide where to route data captured by SymmetricDS, the routing is configured by *Node Group* .

For the examples above, we might define Node Groups of:

- "workstation", to represent each point-of-sale workstation

- "corp" or "central-office" to represent the centralized node.

- "store" to represent the store server that interacts with store workstations and sends and receives data from a central office server.

- "region" to represent the a regional server that interacts with store workstations and sends and receives data from a central office server.

Considerable thought should be given to how you define the Node Groups. Groups should be created for each set of nodes that synchronize common tables in a similar manner. Also, give your Node Groups meaningful names, as they will appear in many, many places in your implementation of SymmetricDS.

Note that there are other mechanisms in SymmetricDS to route to individual nodes or smaller subsets of nodes within a Node Group, so do not choose Node Groups based on needing only subsets of data at specific nodes. For example, although you could, you would not want to create a Node Group for each store even though different tax rates need to be routed to each store. Each store needs to synchronize the same tables to the same groups, so 'store' would be a good choice for a Node Group.

3.4. Linking Nodes

Now that Node Groups have been chosen, the next step in planning is to document the individual links between Node Groups. These *Node Group Links* establish a source Node Group, a target Node Group, and a *data event action* , namely whether the data changes are *pushed* or *pulled* . The push method causes the source Node Group to connect to the target, while a pull method causes it to wait for the target to connect to it.

For our retail store example, there are two Node Group Links defined. For the first link, the "store" Node Group pushes data to the "corp" central office Node Group. The second defines a "corp" to "store" link as a pull. Thus, the store nodes will periodically pull data from the central office, but when it comes time to send data to the central office a store node will do a push.

3.5. Choosing Data Channels

When SymmetricDS captures data changes in the database, the changes are captured in the order in which they occur. In addition, that order is preserved when synchronizing the data to other nodes. Frequently, however, you will have cases where you have different "types" of data with differing priorities. Some data might, for example, need priority for synchronization despite the normal order of events. For example, in a retail environment, users may be waiting for inventory documents to update while a promotional sale event updates a large number of items.

SymmetricDS supports this by allowing tables being synchronized to be grouped together into *Channels*

of data. A number of controls to the synchronization behavior of SymmetricDS are controlled at the Channel level. For example, Channels provide a processing order when synchronizing, a limit on the amount of data that will be batched together, and isolation from errors in other channels. By categorizing data into channels and assigning them to TRIGGER s, the user gains more control and visibility into the flow of data. In addition, SymmetricDS allows for synchronization to be enabled, suspended, or scheduled by Channels as well. The frequency of synchronization can also be controlled at the channel level.

Choosing Channels is fairly straightforward and can be changed over time, if needed. Think about the differing "types" of data present in your application, the volume of data in the various types, etc. What data is considered must-have and can't be delayed due to a high volume load of another type of data? For example, you might place employee-related data, such as clocking in or out, on one channel, but sales transactions on another. We will define which tables belong to which channels in the next sections.

Important

Be sure that, when defining Channels, all tables related by foreign keys are included in the same channel.

Important

Avoid deadlocks! If client database transactions include tables that update common rows along with different rows, then concurrent synchronization can cause database deadlocks. You can avoid this by using channels to segregate those tables that cause the deadlocks.

3.6. Defining Data Changes to be Captured and Routed

At this point, you have designed the node-related aspects of your implementation, namely choosing nodes, grouping the nodes based on functionality, defining which node groups send and receive data to which others (and by what method). You have defined data Channels based on the types and priority of data being synchronized. The largest remaining task prior to starting your implementation is to define and document what data changes are to be captured (by defining SymmetricDS *Triggers*), to decide to which node(s) the data changes are to be *routed* to, and to decide which trigger applies to which router and under what conditions. We will also, in this section, discuss the concept of an *initial load* of data into a SymmetricDS node.

3.6.1. Defining Triggers

SymmetricDS uses *database triggers* to capture and record changes to be synchronized to other nodes. Based on the configuration you provide, SymmetricDS creates the needed database triggers automatically for you. There is a great deal of flexibility in terms of defining the exact conditions under which a data change is captured. SymmetricDS triggers are defined in a table named TRIGGER . Each trigger you define is for a particular table associated. Each trigger can also specify:

- whether to install a trigger for updates, inserts, and/or deletes

- conditions on which an insert, update, and/or delete fires

- a list of columns that should not be synchronized from this table

- a SQL select statement that can be used to hold data needed for routing (known as External Data)

As you define your triggers, consider which data changes are relevant to your application and which ones ar not. Consider under what special conditions you might want to route data, as well. For our retail example, we likely want to have triggers defined for updating, inserting, and deleting pricing information in the central office so that the data can be routed down to the stores. Similarly, we need triggers on sales transaction tables such that sales information can be sent back to the central office.

3.6.2. Defining Routers

The triggers that have been defined in the previous section only define *when* data changes are to be captured for synchronization. They do not define *where* the data changes are to be sent to. Routers, plus a mapping between Triggers and Routers (TRIGGER_ROUTER), define the process for determining which nodes receive the data changes.

Before we discuss Routers and Trigger Routers, we should probably take a break and discuss the process SymmetricDS uses to keep track of the changes and routing. As we stated, SymmetricDS relies on auto-created database triggers to capture and record relevant data changes into a table, the DATA table. After the data is captured, a background process chooses the nodes that the data will be synchronized to. This is called *routing* and it is performed by the Routing Job. Note that the Routing Job does not actually send any data. It just organizes and records the decisions on where to send data in a "staging" table called DATA_EVENT and OUTGOING_BATCH .

Now we are ready to discuss Routers. The router itself is what defines the configuration of where to send a data change. Each Router you define can be associated with or assigned to any number of Triggers through a join table that defines the relationship. Routers are defined the SymmetricDS table named ROUTER . For each router you define, you will need to specify:

- the target table on the destination node to route the data

- the source node group and target node group for the nodes to route the data to

- a router *type* and router *expression*

- whether to route updates, inserts, and/or deletes

For now, do not worry about the specific routing types. They will be covered later. For your design simply make notes of the information needed and decisions to determine the list of nodes to route to. You will find later that there is incredible flexibility and functionality available in routers. For example, you will find you can:

- send the changes to all nodes that belong to the target node group defined in the router.

- compare old or new column values to a constant value or the value of a node's identity.

- execute a SQL expression against the database to select nodes to route to. This SQL expression can be passed values of old and new column values.

- execute a Bean Shell expression in order to select nodes to route to. The Bean Shell expression can use the the old and new column values.

- publish data changes directly to a messaging solution instead of transmitting changes to registered nodes. (This router must be configured manually in XML as an extension point.)

3.6.3. Mapping Triggers to Routers

For each of your Triggers (which specify when a data change should be captured), you will need to decide which Router(s) to pair with the Trigger such that the change is routed to the desired target nodes. This needed mapping between Triggers and Routers, found in the table TRIGGER_ROUTER , defines configuration specific to a particular Trigger and Router combination that you need. In addition to defining which triggers map to which routers, the table also has several settings present to define various behaviors, including initial loads and ping back.

3.6.3.1. Planning Initial Loads

SymmetricDS provides the ability to "load" or "seed" a node's database with specific sets of data from its parent node. This concept is known as an *Initial Load* of data and is used to start off most synchronization scenarios. The Trigger Router mapping defines how initial loads can occur, so now is a good time to plan how your *Initial Loads* will work. Using our retail example, consider a new store being opened. Initially, you would like to pre-populate a store database with all the item, pricing, and tax data for that specific store. This is achieved through an initial load. A part of your planning, be sure to consider which tables, if any, will need to be loaded initially. SymmetricDS can also perform an initial load on a table with just a subset of data. Initial Loads are further discussed in Section 4.6.3.2, Initial Loads (p. 47).

3.6.3.2. Circular References and "Ping Back"

When routing data, SymmetricDS by default checks each data change and will not route a data change back to a node if it originated the change to begin with. This prevents the possibility of data changes resulting in an infinite loop of changes under certain circumstances. You may find that, for some reason, you need SymmetricDS to go ahead and send the data back to the originating node - a "ping back". As part of the planning process, consider whether you have a special case for needing ping back. Ping Back control is further discussed in Section 4.6.3.4, Enabling "Ping Back" (p. 50).

3.6.4. Planning for Registering Nodes

Our final step in planning an implementation of SymmetricDS involves deciding how a new node is connected to, or *registered* with a parent node for the first time.

The following are some options on ways you might register nodes:

- The tutorial uses the command line utility to register each individual node.

- A JMX interface provides the same interface that the command line utility does. JMX can be invoked programatically or via a web console.

- Both the utility and the JMX method register a node by inserting into two tables. A script can be written to directly register nodes by directly inserting into the database.

- SymmetricDS can be configured to auto register nodes. This means that any node that asks for a registration will be given one.

3.7. Planning Data Transformations

SymmetricDS also provides the abilty to *transform* synchronized data instead of simply synchronizing it. Your application might, for example require a particular column in your source data to be mapped to two different target tables with possibly different column names. Or, you might need to "merge" one or more columns of data from two indepdentent tables into one table on the target. Or, you may want to set default column values on a target table based on a particular event on the source database. All of these operations, and many more, can be accomplished using SymmetricDS' transformation capabilities.

As you plan your SymmetricDS implementation, make notes of cases where a data transformation is needed. Include details such as when the transformation might occur (is it only on an insert, or a delete?), which tables or columns play a part, etc. Complete details of all the transformation features, including how to configure a transformation, are discussed in .

3.8. Planning Conflict Detection and Resolution

As a final step to planning an implementation, consider for a moment cases in which the same data may be modified at nearly the same time at more than one node. For example, can data representing a customer be modified at both a central office and a store location? Conflict detection is the act of determining if an insert, update or delete is in "conflict" due to the target data row not being consistent with the data at the source prior to the insert/update/delete. Conflict resolution is the act of figuring out what to do when a conflict is detected. Both detection and resolution behaviour can be configured and customized in a number of ways. For example, a conflict can be "detected" based solely on a single column which has been modified to a different value, or a row can be considered in conflict if any data in the row has been changed from what was expected, even if the column that has been changed was still expected. There are also numerous ways to resolve the conflict, such as referencing a timestamp column and choosing whichever edit was "most recent" or perhaps causing the conflict to cause the channel to go into error until a manual resolution takes place. A set of conflict detection / resolution rules is configured for a given node group link, but you can set the rules to be for a given channel or for a given table in a channel.

For the purpose of planning your implementation, make a list of all tables that could have data being modified at more than one node at the same time. For each table, think through what should happen in each case if such an event occurs. If the tables on a given channel all have the same set of conflict resolution and detection rules, then you might be able to configure the rules for the channel instead of a series of table-level detections and resolutions. Complete details on how to configure conflict resolution

and detection are discussed further in .

Chapter 4. Configuration

Chapter 3 introduced numerous concepts and the analysis and design needed to create an implementation of SymmetricDS. This chapter re-visits each analysis step and documents how to turn a SymmetricDS design into reality through configuration of the various SymmetricDS tables. In addition, several advanced configuration options, not presented previously, will also be covered.

4.1. Node Properties

To get a SymmetricDS node running, it needs to be given an identity and it needs to know how to connect to the database it will be synchronizing. The preferred way to configure a SymmetricDS engine is to create a properties file in the engines directory. The SymmetricDS server will create an engine for each properties file found in the engines directory. When started up, SymmetricDS reads the synchronization configuration and state from the database. If the configuration tables are missing, they are created automatically (auto creation can be disabled). Basic configuration is described by inserting into the following tables (the complete data model is defined in Appendix A, *Data Model* (p. 91)).

- NODE_GROUP - specifies the tiers that exist in a SymmetricDS network

- NODE_GROUP_LINK - links two node groups together for synchronization

- CHANNEL - grouping and priority of synchronizations

- TRIGGER - specifies tables, channels, and conditions for which changes in the database should be captured

- ROUTER - specifies the routers defined for synchronization, along with other routing details

- TRIGGER_ROUTER - provides mappings of routers and triggers

During start up, triggers are verified against the database, and database triggers are installed on tables that require data changes to be captured. The Route, Pull and Push Jobs begin running to synchronize changes with other nodes.

Each node requires properties that allow it to connect to a database and register with a parent node. Properties are configured in a file named `xxxxx.properties` that is placed in the engines directory of the SymmetricDS install. The file is usually named according to the engine.name, but it is not a requirement.

To give a node its identity, the following properties are required. Any other properties found in `conf/symmetric.properties` can be overridden for a specific engine in an engine's properties file. If the properties are changed in `conf/symmetric.properties` they will take effect across all engines deployed to the server. Note that you can use the variable `$(hostName)` to represent the host name of the machine when defining these properties (for example, external.id=$(hostName)).

engine.name
This is an arbitrary name that is used to access a specific engine using an HTTP URL. Each node configured in the engines directory must have a unique engine name. The engine name is also used for

the domain name of registered JMX beans.

group.id
The node group that this node is a member of. Synchronization is specified between node groups, which means you only need to specify it once for multiple nodes in the same group.

external.id
The external id for this node has meaning to the user and provides integration into the system where it is deployed. For example, it might be a retail store number or a region number. The external id can be used in expressions for conditional and subset data synchronization. Behind the scenes, each node has a unique sequence number for tracking synchronization events. That makes it possible to assign the same external id to multiple nodes, if desired.

sync.url
The URL where this node can be contacted for synchronization. At startup and during each heartbeat, the node updates its entry in the database with this URL. The sync url is of the format:
`http://{hostname}:{port}/{webcontext}/sync/{engine.name}`.

The {webcontext} is blank for a standalone deployment. It will typically be the name of the war file for an application server deployment.

The {engine.name} can be left blank if there is only one engine deployed in a SymmetricDS server.

When a new node is first started, it is has no information about synchronizing. It contacts the registration server in order to join the network and receive its configuration. The configuration for all nodes is stored on the registration server, and the URL must be specified in the following property:

registration.url
The URL where this node can connect for registration to receive its configuration. The registration server is part of SymmetricDS and is enabled as part of the deployment. This is typically equal to the value of the sync.url of the registration server.

> ⚠️ **Important**
>
> Note that a *registration server node* is defined as one whose `registration.url` is either (a) blank, or (b) identical to its `sync.url`.

For a deployment where the database connection pool should be created using a JDBC driver, set the following properties:

db.driver
The class name of the JDBC driver.

db.url
The JDBC URL used to connect to the database.

db.user
The database username, which is used to login, create, and update SymmetricDS tables.

db.password
The password for the database user.

4.2. Node

A *node*, a single instance of SymmetricDS, is defined in the NODE table. Two other tables play a direct role in defining a node, as well The first is NODE_IDENTITY. The *only* row in this table is inserted in the database when the node first *registers* with a parent node. In the case of a root node, the row is entered by the user. The row is used by a node instance to determine its node identity.

The following SQL statements set up a top-level registration server as a node identified as "00000" in the "corp" node group.

```
insert into SYM_NODE
  (node_id, node_group_id, external_id, sync_enabled)
values
  ('00000', 'corp', '00000', 1);

insert into SYM_NODE_IDENTITY values ('00000');
```

The second table, NODE_SECURITY has rows created for each *child* node that registers with the node, assuming auto-registration is enabled. If auto registration is not enabled, you must create a row in NODE and NODE_SECURITY for the node to be able to register. You can also, with this table, manually cause a node to re-register or do a re-initial load by setting the corresponding columns in the table itself. Registration is discussed in more detail in .

4.3. Node Group

Node Groups are straightforward to configure and are defined in the NODE_GROUP table. The following SQL statements would create node groups for "corp" and "store" based on our retail store example.

```
insert into SYM_NODE_GROUP
  (node_group_id, description)
values
  ('store', 'A retail store node');

insert into SYM_NODE_GROUP
  (node_group_id, description)
values
  ('corp', 'A corporate node');
```

4.4. Node Group Link

Similarly, Node Group links are established using a data event action of 'P' for Push and 'W' for Pull ("wait"). The following SQL statements links the "corp" and "store" node groups for synchronization. It configures the "store" nodes to push their data changes to the "corp" nodes, and the "corp" nodes to send changes to "store" nodes by waiting for a pull.

```
insert into SYM_NODE_GROUP_LINK
  (source_node_group, target_node_group, data_event_action)
values
  ('store', 'corp', 'P');

insert into SYM_NODE_GROUP_LINK
  (source_node_group, target_node_group, data_event_action)
values
  ('corp', 'store', 'W');
```

A node group link can be configured to use the same node group as the source and the target. This configuration allows a node group to sync with every other node in its group.

A third type of link action of 'R' for 'Route Only' exists if you want to associate a router with a link that will not move the data. This action type might be useful when using an XML publishing router or an audit table changes router.

4.5. Channel

By categorizing data into channels and assigning them to TRIGGERs, the user gains more control and visibility into the flow of data. In addition, SymmetricDS allows for synchronization to be enabled, suspended, or scheduled by channels as well. The frequency of synchronization and order that data gets synchronized is also controlled at the channel level.

The following SQL statements setup channels for a retail store. An "item" channel includes data for items and their prices, while a "sale_transaction" channel includes data for ringing sales at a register.

```
insert into SYM_CHANNEL
  (channel_id, processing_order, max_batch_size, max_batch_to_send,
   extract_period_millis, batch_algorithm, enabled, description)
values
  ('item', 10, 1000, 10,  0, 'default', 1, 'Item and pricing data');

insert into SYM_CHANNEL
  (channel_id, processing_order, max_batch_size, max_batch_to_send,
   extract_period_millis, batch_algorithm, enabled, description)
values
  ('sale_transaction', 1, 1000, 10,  60000, 'transactional', 1,
   'retail sale transactions from register');
```

Batching is the grouping of data, by channel, to be transferred and committed at the client together. There are three different out-of-the-box batching algorithms which may be configured in the batch_algorithm column on channel.

default
All changes that happen in a transaction are guaranteed to be batched together. Multiple transactions will be batched and committed together until there is no more data to be sent or the max_batch_size is reached.

transactional
Batches will map directly to database transactions. If there are many small database transactions, then there will be many batches. The max_batch_size column has no effect.

nontransactional
Multiple transactions will be batched and committed together until there is no more data to be sent or the max_batch_size is reached. The batch will be cut off at the max_batch_size regardless of whether it is in the middle of a transaction.

If a channel contains *only* tables that will be synchronized in one direction and and data is routed to all the nodes in the target node groups, then batching on the channel can be optimized to share batches across nodes. This is an important feature when data needs to be routed to thousands of nodes. When this mode is detected, you will see batches created in OUTGOING_BATCH with the common_flag set to 1.

There are also several size-related parameters that can be set by channel. They include:

max_batch_size
Specifies the maximum number of data events to process within a batch for this channel.

max_batch_to_send
Specifies the maximum number of batches to send for a given channel during a 'synchronization' between two nodes. A 'synchronization' is equivalent to a push or a pull. For example, if there are 12 batches ready to be sent for a channel and max_batch_to_send is equal to 10, then only the first 10 batches will be sent even though 12 batches are ready.

max_data_to_route
Specifices the maximum number of data rows to route for a channel at a time.

Based on your particular synchronization requirements, you can also specify whether old, new, and primary key data should be read and included during routing for a given channel. These are controlled by the columns use_old_data_to_route, use_row_data_to_route, and use_pk_data_to_route, respectively. By default, they are all 1 (true).

Finally, if data on a particular channel contains big lobs, you can set the column contains_big_lob to 1 (true) to provide SymmetricDS the hint that the channel contains big lobs. Some databases have shortcuts that SymmetricDS can take advantage of if it knows that the lob columns in DATA aren't going to contain large lobs. The definition of how large a 'big' lob is varies from database to database.

4.6. Triggers, Routers, and Trigger / Routers Mappings

In order to synchronize data, you must define at least one trigger, at least one router, and provide at least one link between the two (known as a trigger-router).

4.6.1. Trigger

SymmetricDS captures synchronization data using database triggers. SymmetricDS' Triggers are defined in the TRIGGER table. Each record is used by SymmetricDS when generating database triggers. Database triggers are only generated when a trigger is associated with a ROUTER whose `source_node_group_id` matches the node group id of the current node.

The `source_table_name` may contain the asterisk ('*') wildcard character so that one TRIGGER table entry can define synchronization for many tables. System tables and any tables that start with the SymmetricDS table prefix will be excluded. A list of wildcard tokens can also be supplied. If there are multiple tokens, they should be delimited with a comma. A wildcard token can also start with a bang ('!') to indicate an exclusive match. Tokens are always evaluated from left to right. When a table match is made, the table is either added to or removed from the list of tables. If another trigger already exists for a table, then that table is not included in the wildcard match (the explictly defined trigger entry take precendence).

When determining whether a data change has occurred or not, by defalt the triggers will record a change even if the data was updated to the same value(s) they were originally. For example, a data change will be captured if an update of one column in a row updated the value to the same value it already was. There is a global property, `trigger.update.capture.changed.data.only.enabled` (false by default), that allows you to override this behavior. When set to true, SymmetricDS will only capture a change if the data has truly changed (i.e., when the new column data is not equal to the old column data).

> **Important**
>
> The property `trigger.update.capture.changed.data.only.enabled` is currently only supported in the MySQL, DB2 and Oracle dialects.

The following SQL statement defines a trigger that will capture data for a table named "item" whenever data is inserted, updated, or deleted. The trigger is assigned to a channel also called 'item'.

```
insert into SYM_TRIGGER
    (trigger_id,source_table_name,channel_id,last_update_time,create_time)
  values
    ('item', 'item', 'item', current_timestamp, current_timestamp);
```

> **Important**
>
> Note that many databases allow for multiple triggers of the same type to be defined. Each database defines the order in which the triggers fire differently. If you have additional triggers beyond those SymmetricDS installs on your table, please consult your database documentation to determine if there will be issues with the ordering of the triggers.

4.6.1.1. Large Objects

Two lobs-related settings are also available on TRIGGER:

use_stream_lobs

Specifies whether to capture lob data as the trigger is firing or to stream lob columns from the source tables using callbacks during extraction. A value of 1 indicates to stream from the source via callback; a value of 0, lob data is captured by the trigger.

use_capture_lobs

Provides a hint as to whether this trigger will capture big lobs data. If set to 1 every effort will be made during data capture in trigger and during data selection for initial load to use lob facilities to extract and store data in the database.

4.6.1.2. External Select

Occasionally, you may find that you need to capture and save away a piece of data present in another table when a trigger is firing. This data is typically needed for the purposes of determining where to 'route' the data to once routing takes place. Each trigger definition contains an optional `external_select` field which can be used to specify the data to be captured. Once captured, this data is available during routing in DATA's `external_data` field. For these cases, place a SQL select statement which returns the data item you need for routing in `external_select`. An example of the use of external select can be found in .

4.6.2. Router

Routers provided in the base implementation currently include:

- Default Router - a router that sends all data to all nodes that belong to the target node group defined in the router.

- Column Match Router - a router that compares old or new column values to a constant value or the value of a node's external_id or node_id.

- Lookup Router - a router which can be configured to determine routing based on an existing or ancillary table specifically for the purpose of routing data.

- Subselect Router - a router that executes a SQL expression against the database to select nodes to route to. This SQL expression can be passed values of old and new column values.

- Scripted Router - a router that executes a Bean Shell script expression in order to select nodes to route to. The script can use the the old and new column values.

- Xml Publishing Router - a router the publishes data changes directly to a messaging solution instead of transmitting changes to registered nodes. This router must be configured manually in XML as an extension point.

- Audit Table Router - a router that inserts into an automatically created audit table. It records captured changes to tables that it is linked to.

The mapping between the set of triggers and set of routers is many-to-many. This means that one trigger can capture changes and route to multiple locations. It also means that one router can be defined an

associated with many different triggers.

4.6.2.1. Default Router

The simplest router is a router that sends all the data that is captured by its associated triggers to all the nodes that belong to the target node group defined in the router. A router is defined as a row in the ROUTER table. It is then linked to triggers in the TRIGGER_ROUTER table.

The following SQL statement defines a router that will send data from the 'corp' group to the 'store' group.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id,
    create_time, last_update_time)
values
  ('corp-2-store','corp', 'store', current_timestamp, current_timestamp);
```

The following SQL statement maps the 'corp-2-store' router to the item trigger.

```
insert into SYM_TRIGGER_ROUTER
  (trigger_id, router_id, initial_load_order,  create_time, last_update_time)
values
  ('item', 'corp-2-store', 1, current_timestamp, current_timestamp);
```

4.6.2.2. Column Match Router

Sometimes requirements may exist that require data to be routed based on the current value or the old value of a column in the table that is being routed. Column routers are configured by setting the router_type column on the ROUTER table to column and setting the router_expression column to an equality expression that represents the expected value of the column.

The first part of the expression is always the column name. The column name should always be defined in upper case. The upper case column name prefixed by OLD_ can be used for a comparison being done with the old column data value.

The second part of the expression can be a constant value, a token that represents another column, or a token that represents some other SymmetricDS concept. Token values always begin with a colon (:).

Consider a table that needs to be routed to all nodes in the target group only when a status column is set to 'READY TO SEND.' The following SQL statement will insert a column router to accomplish that.

```
insert into SYM_ROUTER
(router_id, source_node_group_id, target_node_group_id, router_type,
 router_expression, create_time, last_update_time)
values
('corp-2-store-ok','corp', 'store', 'column',
  'STATUS=READY TO SEND', current_timestamp, current_timestamp);
```

Consider a table that needs to be routed to all nodes in the target group only when a status column changes values. The following SQL statement will insert a column router to accomplish that. Note the use of OLD_STATUS, where the OLD_ prefix gives access to the old column value.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-status','corp', 'store', 'column',
    'STATUS!=:OLD_STATUS', current_timestamp, current_timestamp);
```

Consider a table that needs to be routed to only nodes in the target group whose STORE_ID column matches the external id of a node. The following SQL statement will insert a column router to accomplish that.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-id','corp', 'store', 'column',
    'STORE_ID=:EXTERNAL_ID', current_timestamp, current_timestamp);
```

Attributes on a NODE that can be referenced with tokens include:

- :NODE_ID

- :EXTERNAL_ID

- :NODE_GROUP_ID

Captured EXTERNAL_DATA is also available for routing as a virtual column.

Consider a table that needs to be routed to a redirect node defined by its external id in the REGISTRATION_REDIRECT table. The following SQL statement will insert a column router to accomplish that.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-redirect','corp', 'store', 'column',
    'STORE_ID=:REDIRECT_NODE', current_timestamp, current_timestamp);
```

More than one column may be configured in a router_expression. When more than one column is configured, all matches are added to the list of nodes to route to. The following is an example where the STORE_ID column may contain the STORE_ID to route to or the constant of ALL which indicates that all nodes should receive the update.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-multiple-matches','corp', 'store', 'column',
    'STORE_ID=ALL or STORE_ID=:EXTERNAL_ID', current_timestamp, current_timestamp);
```

The NULL keyword may be used to check if a column is null. If the column is null, then data will be routed to all nodes who qualify for the update. This following is an example where the STORE_ID column is used to route to a set of nodes who have a STORE_ID equal to their EXTERNAL_ID, or to all nodes if the STORE_ID is null.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-multiple-matches','corp', 'store', 'column',
    'STORE_ID=NULL or STORE_ID=:EXTERNAL_ID', current_timestamp, current_timestamp);
```

4.6.2.3. Lookup Table Router

A lookup table may contain the id of the node where data needs to be routed. This could be an existing table or an ancillary table that is added specifically for the purpose of routing data. Lookup table routers are configured by setting the `router_type` column on the ROUTER table to `lookuptable` and setting a list of configuration parameters in the `router_expression` column.

Each of the following configuration parameters are required.

LOOKUP_TABLE
This is the name of the lookup table.

KEY_COLUMN
This is the name of the column on the table that is being routed. It will be used as a key into the lookup table.

LOOKUP_KEY_COLUMN
This is the name of the column that is the key on the lookup table.

EXTERNAL_ID_COLUMN
This is the name of the column that contains the external_id of the node to route to on the lookup table.

Note that the lookup table will be read into memory and cached for the duration of a routing pass for a single channel.

Consider a table that needs to be routed to a specific store, but the data in the changing table only contains brand information. In this case, the STORE table may be used as a lookup table.

```
insert into SYM_ROUTER
(router_id, source_node_group_id, target_node_group_id, router_type,
 router_expression, create_time, last_update_time)
values
('corp-2-store-ok','corp', 'store', 'lookuptable',
 'LOOKUP_TABLE=STORE
KEY_COLUMN=BRAND_ID
LOOKUP_KEY_COLUMN=BRAND_ID
EXTERNAL_ID_COLUMN=STORE_ID', current_timestamp, current_timestamp);
```

4.6.2.4. Subselect Router

Sometimes routing decisions need to be made based on data that is not in the current row being synchronized. A 'subselect' router can be used in these cases. A 'subselect' is configured with a `router_expression` that is a SQL select statement which returns a result set of the node ids that need routed to. Column tokens can be used in the SQL expression and will be replaced with row column data. The overhead of using this router type is high because the 'subselect' statement runs for each row that is routed. It should not be used for tables that have a lot of rows that are updated. It also has the disadvantage that if the data being relied on to determine the node id has been deleted before routing takes place, then no results would be returned and routing would not happen.

The `router_expression` you specify is appended to the following SQL statement in order to select the node ids:

```
select c.node_id from sym_node c where
   c.node_group_id=:NODE_GROUP_ID and c.sync_enabled=1 and ...
```

As you can see, you have access to information about the node currently under consideration for routing through the 'c' alias, for example `c.external_id`. There are two node-related tokens you can use in your expression:

- :NODE_GROUP_ID

- :EXTERNAL_DATA

Column names representing data for the row in question are prefixed with a colon as well, for example: `:EMPLOYEE_ID`, or `:OLD_EMPLOYEE_ID`. Here, the OLD_ prefix indicates the value before the change in cases where the old data has been captured.

For an example, consider the case where an Order table and a OrderLineItem table need to be routed to a specific store. The Order table has a column named order_id and STORE_ID. A store node has an external_id that is equal to the STORE_ID on the Order table. OrderLineItem, however, only has a foreign key to its Order of order_id. To route OrderLineItems to the same nodes that the Order will be routed to, we need to reference the master Order record.

There are two possible ways to solve this in SymmetricDS. One is to configure a 'subselect' router_type on the ROUTER table, shown below (The other possible approach is to use an `external_select` to capture the data via a trigger for use in a column match router, demonstrated in Section 4.6.2.7, Utilizing External

Our solution utilizing subselect compares the external id of the current node with the store id from the Order table where the order id matches the order id of the current row being routed:

```
insert into SYM_ROUTER
   (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
   ('corp-2-store','corp', 'store', 'subselect',
    'c.external_id in (select STORE_ID from order where order_id=:ORDER_ID)',
    current_timestamp, current_timestamp);
```

As a final note, please note in this example that the parent row in Order must still exist at the moment of routing for the child rows (OrderLineItem) to route, since the select statement is run when routing is occurring, not when the change data is first captured.

4.6.2.5. Scripted Router

When more flexibility is needed in the logic to choose the nodes to route to, then the a scripted router may be used. The currently available scripting language is Bean Shell. Bean Shell is a Java-like scripting language. Documentation for the Bean Shell scripting language can be found at http://www.beanshell.org.

The router_type for a Bean Shell scripted router is 'bsh'. The router_expression is a valid Bean Shell script that:

- adds node ids to the `targetNodes` collection which is bound to the script

- returns a new collection of node ids

- returns a single node id

- returns true to indicate that all nodes should be routed or returns false to indicate that no nodes should be routed

Also bound to the script evaluation is a list of `nodes`. The list of `nodes` is a list of eligible `org.jumpmind.symmetric.model.Node` objects. The current data column values and the old data column values are bound to the script evaluation as Java object representations of the column data. The columns are bound using the uppercase names of the columns. Old values are bound to uppercase representations that are prefixed with 'OLD_'.

If you need access to any of the SymmetricDS services, then the instance of `org.jumpmind.symmetric.ISymmetricEngine` is accessible via the bound `engine` variable.

In the following example, the node_id is a combination of STORE_ID and WORKSTATION_NUMBER, both of which are columns on the table that is being routed.

```
insert into SYM_ROUTER
   (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
```

```
values
  ('corp-2-store-bsh','corp', 'store', 'bsh',
    'targetNodes.add(STORE_ID + "-" + WORKSTATION_NUMBER);',
    current_timestamp, current_timestamp);
```

The same could also be accomplished by simply returning the node id. The last line of a bsh script is always the return value.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-bsh','corp', 'store', 'bsh',
    'STORE_ID + "-" + WORKSTATION_NUMBER',
    current_timestamp, current_timestamp);
```

The following example will synchronize to all nodes if the FLAG column has changed, otherwise no nodes will be synchronized. Note that here we make use of OLD_, which provides access to the old column value.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-flag-changed','corp', 'store', 'bsh',
    'FLAG != null && !FLAG.equals(OLD_FLAG)',
    current_timestamp, current_timestamp);
```

The next example shows a script that iterates over each eligible node and checks to see if the trimmed value of the column named STATION equals the external_id.

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
values
  ('corp-2-store-trimmed-station','corp', 'store', 'bsh',
    'for (org.jumpmind.symmetric.model.Node node : nodes) {
        if (STATION != null && node.getExternalId().equals(STATION.trim())) {
            targetNodes.add(node.getNodeId());
        }
      }',
    current_timestamp, current_timestamp);
```

4.6.2.6. Audit Table Router

This router audits captured data by recording the change in an audit table that the router creates and keeps

up to date (as long as `auto.config.database` is set to true.) The router creates a table named the same as the table for which data was captured with the suffix of _AUDIT. It will contain all of the same columns as the original table with the same data types only each column is nullable with no default values.

Three extra "AUDIT" columns are added to the table:

- AUDIT_ID - the primary key of the table.

- AUDIT_TIME - the time at which the change occurred.

- AUDIT_EVENT - the DML type that happened to the row.

The following is an example of an audit router

```
insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    create_time, last_update_time)
values
  ('audit_at_corp','corp', 'local', 'audit', current_timestamp, current_timestamp);
```

Because the audit router isn't capturing data for a specific node in the system, but it still has to be associated with a node_group_link a new link action of type 'R' has been introduced. The 'R' stands for 'only routes to'

4.6.2.7. Utilizing External Select when Routing

There may be times when you wish to route based on a piece of data that exists in a table other than the one being routed. The approach, first discussed in Section 4.6.2.4, Subselect Router (p. 43) , is to utlize an `external_select` to save away data in `external_data`, which can then be referenced during routing.

Reconsider subselect's Order / OrderLineItem example (found in Section 4.6.2.4, Subselect Router (p. 43)), where routing for the line item is accomplished by linking to the "header" Order row. As an alternate way of solving the problem, we will now use External Select combined with a column match router.

In this version of the solution, the STORE_ID is captured from the Order table in the EXTERNAL_DATA column when the trigger fires. The router is configured to route based on the captured EXTERNAL_DATA to all nodes whose external id matches the captured external data.

```
insert into SYM_TRIGGER
  (trigger_id,source_table_name,channel_id,external_select,
    last_update_time,create_time)
values
  ('orderlineitem', 'orderlineitem', 'orderlineitem','select STORE_ID
    from order where order_id=$(curTriggerValue).$(curColumnPrefix)order_id',
    current_timestamp, current_timestamp);

insert into SYM_ROUTER
  (router_id, source_node_group_id, target_node_group_id, router_type,
    router_expression, create_time, last_update_time)
```

```
values
  ('corp-2-store-ext','corp', 'store', 'column',
    'EXTERNAL_DATA=:EXTERNAL_ID', current_timestamp, current_timestamp);
```

Note the syntax $(curTriggerValue).$(curColumnPrefix). This translates into "OLD_" or "NEW_" based on the DML type being run. In the case of Insert or Update, it's NEW_. For Delete, it's OLD_ (since there is no new data). In this way, you can access the DML-appropriate value for your select statement.

The advantage of this approach over the 'subselect' approach is that it guards against the (somewhat unlikely) possibility that the master Order table row might have been deleted before routing has taken place. This external select solution also is a bit more efficient than the 'subselect' approach, although the triggers produced do run the extra external_select SQL inline with application database updates.

4.6.3. Trigger / Router Mappings

The TRIGGER_ROUTER table is used to define which specific combinations of triggers and routers are needed for your configuration. The relationship between triggers and routers is many-to-many, so this table serves as the join table to define which combinations are valid, as well as to define settings available at the trigger-router level of granularity.

Three important controls can be configured for a specific Trigger / Router combination: Enabled, Initial Loads and Ping Back. The parameters for these can be found in the Trigger / Router mapping table,TRIGGER_ROUTER .

4.6.3.1. Enable / disable trigger router

Each individual trigger-router combination can be disabled or enabled if needed. By default, a trigger router is enabled, but if you have a reason you wish to define a trigger router combination prior to it being active, you can set the `enabled` flag to 0. This will cause the trigger-router mapping to be sent to all nodes, but the trigger-router mapping will not be considered active or enabled for the purposes of capturing data changes or routing.

4.6.3.2. Initial Loads

An initial load is the process of seeding tables at a target node with data from its parent node. When a node connects and data is extracted, after it is registered and if an initial load was requested, each table that is configured to synchronize to the target node group will be given a reload event in the order defined by the end user. A SQL statement is run against each table to get the data load that will be streamed to the target node. The selected data is filtered through the configured router for the table being loaded. If the data set is going to be large, then SQL criteria can optionally be provided to pair down the data that is selected out of the database.

An initial load can not occur until after a node is registered. An initial load is requested by setting the `initial_load_enabled` column on NODE_SECURITY to *1* on the row for the target node in the parent node's database. You can configure SymmetricDS to automatically perform an initial load when a node registers by setting the parameter `auto.reload` to true. Regardless of how the initial load is initiated, the next time the source node routes data, reload batches will be inserted. At the same time reload batches are

inserted, all previously pending batches for the node are marked as successfully sent.

> ⚠️ **Important**
>
> Note that if the parent node that a node is registering with is *not* a registration server node (as can happen with a registration redirect or certain non-tree structure node configurations) the parent node's NODE_SECURITY entry must exist at the parent node and have a non-null value for column `initial_load_time`. Nodes can't be registered to non-registration-server nodes without this value being set one way or another (i.e., manually, or as a result of an initial load occuring at the parent node).

SymmetricDS recognizes that an initial load has completed when the `initial_load_time` column on the target node is set to a non-null value.

An initial load is accomplished by inserting reload batches in a defined order according to the `initial_load_order` column on TRIGGER_ROUTER. If the `initial_load_order` column contains a negative value the associated table will *NOT* be loaded. If the `initial_load_order` column contains the same value for multiple tables, SymmetricDS will attempt to order the tables according to foreign key constraints. If there are cyclical constraints, then foreign keys might need to be turned off or the initial load will need to be manually configured based on knowledge of how the data is structured.

Initial load data is always queried from the source database table. All data is passed through the configured router to filter out data that might not be targeted at a node.

4.6.3.2.1. Target table prep for initial load

There are several parameters that can be used to specify what, if anything, should be done to the table on the target database just prior to loading the data. Note that the parameters below specify the desired behavior for all tables in the initial load, not just one.

- `initial.load.delete.first / initial.load.delete.first.sql`

  By default, an initial load will not delete existing rows from a target table before loading the data. If a delete is desired, the parameter `initial.load.delete.first` can be set to true. If true, the command found in `initial.load.delete.first.sql` will be run on each table prior to loading the data. Thd default value for `initial.load.delete.first.sql` is `delete from %s`, but could be changed if needed. Note that additional reload batches are created, in the correct order, to achieve the delete.

- `initial.load.create.first`

  By default, an initial load will not create the table on the target if it doesn't aleady exist. If the desired behavior is to create the table on the target if it is not present, set the parameter `intial.load.create.first` to true. SymmetricDS will attempt to create the table and indexes on the target database before doing the initial load. (Additional batches are created to represent the table schema).

4.6.3.2.2. Loading subsets of data

An efficient way to select a subset of data from a table for an initial load is to provide an

`initial_load_select` clause on [TRIGGER_ROUTER](). This clause, if present, is applied as a `where` clause to the SQL used to select the data to be loaded. The clause may use "t" as an alias for the table being loaded, if needed. The `$(externalId)` token can be used for subsetting the data in the where clause.

In cases where routing is done using a feature like [Section 4.6.2.4, Subselect Router (p. 43)](), an `initial_load_select` clause matching the subselect's criteria would be a more efficient approach. Some routers will check to see if the `initial_load_select` clause is provided, and they will *not* execute assuming that the more optimal path is using the `initial_load_select` statement.

One example of the use of an initial load select would be if you wished to only load data created more recently than the start of year 2011. Say, for example, the column `created_time` contains the creation date. Your `initial_load_select` would read `created_time > ts {'2011-01-01 00:00:00.0000'}` (using whatever timestamp format works for your database). This then gets applied as a `where` clause when selecting data from the table.

> ⚠️ **Important**
>
> When providing an `initial_load_select` be sure to test out the criteria against production data in a query browser. Do an explain plan to make sure you are properly using indexes.

4.6.3.2.3. Reverse Initial Loads

The default behavior for initial loads is to load data from the registration server or parent node, to a client node. Occasionally, there may be need to do a one-time intial load of data in the opposite or "reverse" direction, namely from a client node to the registration node. To achieve this, set the parameter `auto.reload.reverse` to be true, *but only for the specific node group representing the client nodes*. This will cause a one time reverse load of data, for tables configured with non-negative initial load orders, to be batched at the point when registration of the client node is occurring. These batches are then sent to the parent or registration node. This capability might be needed, for example, if there is data already present in the client that doesn't exist in the parent but needs to.

4.6.3.3. Dead Triggers

Occasionally the decision of what data to load initially results in additional triggers. These triggers, known as *Dead Triggers*, are configured such that they do not capture any data changes. A "dead" Trigger is one that does not capture data changes. In other words, the `sync_on_insert`, `sync_on_update`, and `sync_on_delete` properties for the Trigger are all set to false. However, since the Trigger is specified, it *will* be included in the initial load of data for target Nodes.

Why might you need a Dead Trigger? A dead Trigger might be used to load a read-only lookup table, for example. It could also be used to load a table that needs populated with example or default data. Another use is a recovery load of data for tables that have a single direction of synchronization. For example, a retail store records sales transaction that synchronize in one direction by trickling back to the central office. If the retail store needs to recover all the sales transactions from the central office, they can be sent are part of an initial load from the central office by setting up dead Triggers that "sync" in that direction.

The following SQL statement sets up a non-syncing dead Trigger that sends the `sale_transaction` table to the "store" Node Group from the "corp" Node Group during an initial load.

```
insert into sym_trigger (TRIGGER_ID,SOURCE_CATALOG_NAME,
   SOURCE_SCHEMA_NAME,SOURCE_TABLE_NAME,CHANNEL_ID,
   SYNC_ON_UPDATE,SYNC_ON_INSERT,SYNC_ON_DELETE,
   SYNC_ON_INCOMING_BATCH,NAME_FOR_UPDATE_TRIGGER,
   NAME_FOR_INSERT_TRIGGER,NAME_FOR_DELETE_TRIGGER,
   SYNC_ON_UPDATE_CONDITION,SYNC_ON_INSERT_CONDITION,
   SYNC_ON_DELETE_CONDITION,EXTERNAL_SELECT,
   TX_ID_EXPRESSION,EXCLUDED_COLUMN_NAMES,
   CREATE_TIME,LAST_UPDATE_BY,LAST_UPDATE_TIME)
   values ('SALE_TRANSACTION_DEAD',null,null,
   'SALE_TRANSACTION','transaction',
   0,0,0,0,null,null,null,null,null,null,null,null,null,
   current_timestamp,'demo',current_timestamp);

insert into sym_router (ROUTER_ID,TARGET_CATALOG_NAME,TARGET_SCHEMA_NAME,
   TARGET_TABLE_NAME,SOURCE_NODE_GROUP_ID,TARGET_NODE_GROUP_ID,ROUTER_TYPE,
   ROUTER_EXPRESSION,SYNC_ON_UPDATE,SYNC_ON_INSERT,SYNC_ON_DELETE,
   CREATE_TIME,LAST_UPDATE_BY,LAST_UPDATE_TIME)
   values ('CORP_2_STORE',null,null,null,
   'corp','store',null,null,1,1,1,
   current_timestamp,'demo',current_timestamp);

insert into sym_trigger_router (TRIGGER_ID,ROUTER_ID,INITIAL_LOAD_ORDER,
   INITIAL_LOAD_SELECT,CREATE_TIME,LAST_UPDATE_BY,LAST_UPDATE_TIME)
   values ('SALE_TRANSACTION_DEAD','CORP_2_REGION',100,null,
    current_timestamp,'demo',current_timestamp);
```

4.6.3.4. Enabling "Ping Back"

As discussed in Section 3.6.3.2, Circular References and "Ping Back" (p. 30) SymmetricDS, by default, avoids circular data changes. When a trigger fires as a result of SymmetricDS itself (such as the case when sync on incoming batch is set), it records the originating source node of the data change in `source_node_id`. During routing, if routing results in sending the data back to the originating source node, the data is not routed by default. If instead you wish to route the data back to the originating node, you can set the `ping_back_enabled` column for the needed particular trigger / router combination. This will cause the router to "ping" the data back to the originating node when it usually would not.

4.7. Opening Registration

Node registration is the act of setting up a new NODE and NODE_SECURITY so that when the new node is brought online it is allowed to join the system. Nodes are only allowed to register if rows exist for the node and the `registration_enabled` flag is set to 1. If the `auto.registration` SymmetricDS property is set to true, then when a node attempts to register, if registration has not already occurred, the node will automatically be registered.

SymmetricDS allows you to have multiple nodes with the same `external_id`. Out of the box, openRegistration will open a new registration if a registration already exists for a node with the same external_id. A new registration means a new node with a new `node_id` and the same `external_id` will be created. If you want to re-register the same node you can use the `reOpenRegistration()` JMX method

which takes a `node_id` as an argument.

4.8. Transforming Data

New as of SymmetricDS 2.4, SymmetricDS is now able to transform synchronized data by way of configuration (previously, for most cases a custom data loader would need to have been written). This transformation can take place on a source node or on a target node, as the data is being loaded or extracted. With this new feature you can, for example:

- Copy a column from a source table to two (or more) target table columns,

- Merge columns from two or more source tables into a single row in a target table,

- Insert constants in columns in target tables based on source data synchronizations,

- Insert multiple rows of data into a single target table based on one change in a source table,

- Apply a Bean Shell script to achieve a custom transform when loading into the target database.

These transformations can take place either on the target or on the source, and as data is either being extracted or loaded. In either case, the transformation is initiated due to existence of a source synchronization trigger. The source trigger creates the synchronization data, while the transformation configuration decides what to do with the sychronization data as it is either being extracted from the source or loaded into the target. You have the flexibility of defining different transformation behavior depending on whether the source change that triggered the synchronization was an Insert, Update, or Delete. In the case of Delete, you even have options on what exactly to do on the target side, be it a delete of a row, setting columns to specific values, or absolutely nothing at all.

A few key concepts are important to keep in mind to understand how SymmetricDS performs transformations. The first concept is that of the "source operation" or "source DML type", which is the type of operation that occurred to generate the synchronization data in the first place (i.e., an insert, a delete, or an update). Your transformations can be configured to act differently based on the source DML type, if desired. When transforming, by default the DML action taken on the target matches that of the action taken on the row in the source (although this behavior can be altered through configuration if needed). If the source DML type is an Insert, for example, the resulting transformation DML(s) will be Insert(s).

Another important concept is the way in which transforms are applied. Each source operation may map to one or more transforms and result in one or more operations on the target tables. Each of these target operations are performed as independent operations in sequence and must be "complete" from a SQL perspective. In other words, you must define columns for the transformation that are sufficient to fill in any primary key or other required data in the target table if the source operation was an Insert, for example.

Finally, please note that the tranformation engine relies on a source trigger / router existing to supply the source data for the transformation. The transform configuration will never be used if the source table and target node group does not have a defined trigger / router combination for that source table and target node group.

4.8.1. Transform Configuration Tables

SymmetricDS stores its transformation configuration in two configuration tables, TRANSFORM_TABLE and TRANSFORM_COLUMN. Defining a transformation involves configuration in both tables, with the first table defining which source and destination tables are involved, and the second defining the columns involved in the transformation and the behavior of the data for those columns. We will explain the various options available in both tables and the various pre-defined transformation types.

To define a transformation, you will first define the source table and target table that applies to a particular transformation. The source and target tables, along with a unique identifier (the transform_id column) are defined in TRANSFORM_TABLE. In addition, you will specify the source_node_group_id and target_node_group_id to which the transform will apply, along with whether the transform should occur on the Extract step or the Load step (transform_point). All of these values are required.

Three additional configuration settings are also defined at the source-target table level: the order of the transformations, the behavior when deleting, and whether an update should always be attempted first. More specifically,

- transform_order: For a single source operation that is mapped to a transformation, there could be more than one target operation that takes place. You may control the order in which the target operations are applied through a configuration parameter defined for each source-target table combination. This might be important, for example, if the foreign key relationships on the target tables require you to execute the transformations in a particular order.

- column_policy: Indicates whether unspecified columns are passed thru or if all columns must be explicitly defined. The options include:

  - SPECIFIED - Indicates that only the transform columns that are defined will be the ones that end up as part of the transformation.

  - IMPLIED - Indicates that if not specified, then columns from the source are passed through to the target. This is useful if you just want to map a table from one name to anther or from one schema to another. It is also useful if you want to transform a table, but also want to pass it through. You would define and implied transform from the source to the target and would not have to configure each column.

- delete_action: When a source operation of Delete takes place, there are three possible ways to handle the transformation at the target. The options include:

  - NONE - The delete results in no target changes.

  - DEL_ROW - The delete results in a delete of the row as specified by the pk columns defined in the transformation configuration.

  - UPDATE_COL - The delete results in an Update operation on the target which updates the specific rows and columns based on the defined transformation.

- update_first: This option overrides the default behavior for an Insert operation. Instead of

attempting the Insert first, SymmetricDS will always perform an Update first and then fall back to an Insert if that fails. Note that, by default, fall back logic *always* applies for Insert and Updates. Here, all you a specifying is whether to always do an Update first, which can have performance benefits under certain situations you may run into.

For each transformation defined in TRANSFORM_TABLE, the columns to be transformed (and how they are transformed) are defined in TRANSFORM_COLUMN. This column-level table typically has several rows for each transformation id, each of which defines the source column name, the target column name, as well as the following details:

- include_on: Defines whether this entry applies to source operations of Insert (I), Update (U), or Delete (D), or any source operation.

- pk: Indicates that this mapping is used to define the "primary key" for identifying the target row(s) (which may or may not be the true primary key of the target table). This is used to define the "where" clause when an Update or Delete on the target is occurring. At least one row marked as a pk should be present for each transform_id.

- transform_type, transform_expression: Specifies how the data is modified, if at all. The available transform types are discussed below, and the default is 'copy', which just copies the data from source to target.

- transform_order: In the event there are more than one columns to transform, this defines the relative order in which the transformations are applied.

4.8.2. Transformation Types

There are several pre-defined transform types available in SymmetricDS. Additional ones can be defined by creating and configuring an extension point which implements the `IColumnTransform` interface. The pre-defined transform types include the following (the transform_type entry is shown in parentheses):

- Copy Column Transform ('copy'): This transformation type copies the source column value to the target column. This is the default behavior.

- Constant Transform ('const'): This transformation type allows you to map a constant value to the given target column. The constant itself is placed in transform_expression.

- Variable Transform ('variable'): This transformation type allows you to map a built-in dynamic variable to the given target column. The variable name is placed in transform_expression. The following variables are available: `Ssystem_date` is the current system date, `system_timestamp` is the current system date and time, `source_node_id` is the node id of the source, `target_node_id` is the node id of the target, and `null` is a null value.

- Additive Transform ('additive'): This transformation type is used for numeric data. It computes the change between the old and new values on the source and then adds the change to the existing value in the target column. That is, target = target + multiplier (source_new - source_old), where multiplier is a constant found in the transform_expression (default is 1 if not specified). For example, if the source column changed from a 2 to a 4, the target column is currently 10, and the

multiplier is 3, the effect of the transform will be to change the target column to a value of 16 (10+3*(4-2) => 16). Note that, in the case of deletes, the new column value is considered 0 for the purposes of the calculation.

- Substring Transform ('substr'): This transformation computes a substring of the source column data and uses the substring as the target column value. The transform_expression can be a single integer (`n`, the beginning index), or a pair of comma-separated integers (`n,m` - the beginning and ending index). The transform behaves as the Java substring function would using the specified values in transform_expression.

- Multiplier Transform ('multiply'): This transformation allows for the creation of multiple rows in the target table based on the transform_expression. This transform type can only be used on a primary key column. The transform_expression is a SQL statement that returns the list to be used to create the multiple targets.

- Lookup Transform ('lookup'): This transformation determines the target column value by using a query, contained in transform_expression to lookup the value in another table. The query must return a single row, and the first column of the query is used as the value. Your query references source column names by prefixing with a colon (e.g., :MY_COLUMN).

- Shell Script Transform ('bsh'): This transformation allows you to provide a Bean Shell script in transform_expression and executes the script at the time of transformation. Some variables are provided to the script: `COLUMN_NAME` is a variable for a source column in the row, where the variable name is the column name in uppercase; `currentValue` is the value of the current source column; `oldValue` is the old value of the source column for an updated row; `sqlTemplate` is a `org.jumpmind.db.sql.ISqlTemplate` object for querying or updating the database; `channelId` is a reference to the channel on which the transformation is happening; `sourceNode` is a `org.jumpmind.symmetric.model.Node` object that represents the node from where the data came; `targetNode` is a `org.jumpmind.symmetric.model.Node` object that represents the node where the data is being loaded.

- Identity Transform ('identity'): This transformation allows you to insert into an identity column by computing a new identity, not copying the actual identity value from the source.

4.9. Data Load Filters

New as of SymmetricDS 3.1, SymmetricDS is now capable of taking actions upon the load of certain data via configurable load filters. This new configurable option is in additon to the already existing option of writing a class that implements IDatabaseWriterFilter. A configurable load filter watches for specific data that is being loaded and then takes action based on the load of that data.

Specifying which data to action is done by specifying a souce and target node group (data extracted from this node group, and loaded into that node group), and a target catalog, schema and table name. You can decide to take action on rows that are inserted, updated and/or deleted, and can also further delineate which rows of the target table to take action on by specifying additional criteria in the bean shell script that is executed in response to the loaded data. As an example, old and new values for the row of data being loaded are available in the bean shell script, so you can action rows with a certain column value in

old or new data.

The action taken is based on a bean shell script that you can provide as part of the configuration. Actions can be taken at different points in the load process including before write, after write, at batch complete, at batch commit and/or at batch rollback.

4.9.1. Load Filter Configuration Table

SymmetricDS stores its load filter configuration in a single table called LOAD_FILTER. The load filter table allows you to specify the following:

- Load Filter Type ('load_filter_type'): The type of load filter. Today only beanshell is support ('BSH'), but SQL scripts may be added in a future release.

- Source Node Group ('source_node_group_id'): The source node group for which you would like to watch for changes.

- Target Node Group ('target_node_group_id'): The target node group for which you would like to watch for changes. The source and target not groups are used together to identify the node group link for which you would like to watch for changes (i.e. When the Server node group sends data to a Client node group).

- Target Catalog ('target_catalog_name'): The name of the target catalog for which you would like to watch for changes.

- Target Schema ('target_schema_name'): The name of the target schema for which you would like to watch for changes.

- Target Table ('target_table_name'): The name of the target table for which you would like to watch for changes. The target catalog, target schema and target table name are used together to fully qualify the table for which you would like to watch for changes.

- Filter on Update ('filter_on_update'): Determines whether the load filter takes action (executes) on a database update statement.

- Filter on Insert ('filter_on_insert'): Determines whether the load filter takes action (executes) on a database insert statement.

- Filter on Delete ('filter_on_delete'): Determines whether the load filter takes action (executes) on a database delete statement.

- Before Write Script ('before_write_script'): The script to execute before the database write occurs.

- After Write Script ('after_write_script'): The script to execute after the database write occurs.

- Batch Complete Script ('batch_complete_script'): The script to execute after the entire batch completes.

- Batch Commit Script ('batch_commit_script'): The script to execute after the entire batch is committed.

- Batch Rollback Script ('batch_rollback_script'): The script to execute if the batch rolls back.

- Handle Error Script ('handle_error_script'): A script to execute if data cannot be processed.

- Load Filter Order ('load_filter_order'): The order in which load filters should execute if there are multiple scripts pertaining to the same source and target data.

4.9.2. Variables available to Data Load Filters

As part of the bean shell load filters, SymmetricDS provides certain variables for use in the bean shell script. Those variables include:

- Symmetric Engine ('ENGINE'): The Symmetric engine object.

- Source Values ('<COLUMN_NAME>'): The source values for the row being inserted, updated or deleted.

- Old Values ('OLD_<COLUMN_NAME>'): The old values for the row being inserted, updated or deleted.

- Data Context ('CONTEXT'): The data context object for the data being inserted, updated or deleted. .

- Table Data ('TABLE'): The table object for the table being inserted, updated or deleted.

4.9.3. Data Load Filter Example

The following is an example of a load filter that watches a table named TABLE_TO_WATCH being loaded from the Server Node Group to the Client Node Group for inserts or updates, and performs an initial load on a table named "TABLE_TO_RELOAD" for KEY_FIELD on the reload table equal to a column named KEY_FIELD on the TABLE_TO_WATCH table.

```
insert into sym_load_filter
    (LOAD_FILTER_ID, LOAD_FILTER_TYPE, SOURCE_NODE_GROUP_ID, TARGET_NODE_GROUP_ID,
     TARGET_CATALOG_NAME, TARGET_SCHEMA_NAME, TARGET_TABLE_NAME,
     FILTER_ON_UPDATE, FILTER_ON_INSERT, FILTER_ON_DELETE,
     BEFORE_WRITE_SCRIPT, AFTER_WRITE_SCRIPT, BATCH_COMPLETE_SCRIPT,
     BATCH_COMMIT_SCRIPT, BATCH_ROLLBACK_SCRIPT, HANDLE_ERROR_SCRIPT,
     CREATE_TIME, LAST_UPDATE_BY, LAST_UPDATE_TIME,
     LOAD_FILTER_ORDER, FAIL_ON_ERROR)
values ('TABLE_TO_RELOAD','BSH','Client','Server',NULL,NULL,
     'TABLE_TO_WATCH',1,1,0,null,
     'engine.getDataService().reloadTable(context.getBatch().getSourceNodeId(),
     table.getCatalog(), table.getSchema(),
     "TABLE_TO_RELOAD","KEY_FIELD=''" + KEY_FIELD + "''");'
     ,null,null,null,null,sysdate,'userid',sysdate,1,1);
```

4.10. Conflict Detection and Resolution

Conflict detection and resolution is new as of SymmetricDS 3.0. Conflict detection is the act of determining if an insert, update or delete is in "conflict" due to the target data row not being consistent with the data at the source prior to the insert/update/delete. Conflict resolution is the act of figuring out what to do when a conflict is detected.

Conflict detection and resolution strategies are configured in the CONFLICT table. They are configured at minimum for a specific NODE_GROUP_LINK . The configuration can also be specific to a CHANNEL and/or table.

Conflict detection is configured in the `detect_type` and `detect_expression` columns of CONFLICT . The value for `detect_expression` depends on the `detect_type` . Conflicts are detected while data is being loaded into a target system.

USE_PK_DATA
Indicates that only the primary key is used to detect a conflict. If a row exists with the same primary key, then no conflict is detected during an update or a delete. Updates and deletes rows are resolved using only the primary key columns. If a row already exists during an insert then a conflict has been detected.

USE_OLD_DATA
Indicates that all of the old data values are used to detect a conflict. Old data is the data values of the row on the source system prior to the change. If a row exists with the same old values on the target system as they were on the source system, then no conflict is detected during an update or a delete. If a row already exists during an insert then a conflict has been detected.

Note that some platforms do not support comparisons of binary columns. Conflicts in binary column values will not be detected on the following platforms: DB2, DERBY, ORACLE, and SQLSERVER.

USE_CHANGED_DATA
Indicates that the primary key plus any data that has changed on the source system will be used to detect a conflict. If a row exists with the same old values on the target system as they were on the source system for the columns that have changed on the source system, then no conflict is detected during an update or a delete. If a row already exists during an insert then a conflict has been detected.

Note that some platforms do not support comparisons of binary columns. Conflicts in binary column values will not be detected on the following platforms: DB2, DERBY, ORACLE, and SQLSERVER.

USE_TIMESTAMP
Indicates that the primary key plus a timestamp column (as configured in `detect_expression`) will indicate whether a conflict has occurred. If the target timestamp column is not equal to the old source timestamp column, then a conflict has been detected. If a row already exists during an insert then a conflict has been detected.

USE_VERSION
Indicates that the primary key plus a version column (as configured in `detect_expression`) will indicate whether a conflict has occurred. If the target version column is not equal to the old source version column, then a conflict has been detected. If a row already exists during an insert then a

conflict has been detected.

⚠️ **Important**

Be aware that conflict detection will *not* detect changes to binary columns in the case where `use_stream_lobs` is true in the trigger for the table. In addition, some databases do not allow comparisons of binary columns whether `use_stream_lobs` is true or not.

The choice of how to resolve a detected conflict is configured via the `resolve_type` column. Depending on the setting, two additional boolean settings may also be configured, namely `resolve_row_only` and `resolve_changes_only`, as discussed in the resolution settings below.

FALLBACK
Indicates that when a conflict is detected the system should automatically apply the changes anyways. If the source operation was an insert, then an update will be attempted. If the source operation was an update and the row does not exist, then an insert will be attempted. If the source operation was a delete and the row does not exist, then the delete will be ignored. The `resolve_changes_only` flag controls whether all columns will be updated or only columns that have changed will be updated during a fallback operation.

IGNORE
Indicates that when a conflict is detected the system should automatically ignore the incoming change. The `resolve_row_only` column controls whether the entire batch should be ignore or just the row in conflict.

MANUAL
Indicates that when a conflict is detected the batch will remain in error until manual intervention occurs. A row in error is inserted into the INCOMING_ERROR table. The conflict detection id that detected the conflict is recorded (i.e., the `conflict_id` value from CONFLICT), along with the old data, new data, and the "current data" (by current data, we mean the unexpected data at the target which doesn't match the old data as expected) in columns `old_data`, `new_data`, and `cur_data`. In order to resolve, the `resolve_data` column can be manually filled out which will be used on the next load attempt instead of the original source data. The `resolve_ignore` flag can also be used to indicate that the row should be ignored on the next load attempt.

NEWER_WINS
Indicates that when a conflict is detected by USE_TIMESTAMP or USE_VERSION that the either the source or the target will win based on the which side has the newer timestamp or higher version number.

For each configured conflict, you also have the ability to control if and how much "resolved" data is sent back to the node who's data change is in conflict. This "ping back" behavior is specified by the setting of the `ping_back` column and can be one of the following values:

OFF
No data is sent back to the originating node, even if the resolved data doesn't match the data the node sent.

SINGLE_ROW

The resolved data of the single row in the batch that caused the conflict is sent back to the originating node.

REMAINING_ROWS.

The resolved data of the single row in the batch in conflict, along with the entire remainder of the batch, is sent back to the originating node.

Chapter 5. Advanced Topics

This chapter focuses on a variety of topics, including deployment options, jobs, clustering, encryptions, synchronization control, and configuration of SymmetricDS.

5.1. Advanced Synchronization

5.1.1. Bi-Directional Synchronization

SymmetricDS allows tables to be synchronized bi-directionally. Note that an outgoing synchronization does not process changes during an incoming synchronization on the same node unless the trigger was created with the `sync_on_incoming_batch` flag set. If the `sync_on_incoming_batch` flag is set, then update loops are prevented by a feature that is available in most database dialects. More specifically, during an incoming synchronization the source `node_id` is put into a database session variable that is available to the database trigger. Data events are not generated if the target `node_id` on an outgoing synchronization is equal to the source `node_id`.

By default, only the columns that changed will be updated in the target system.

Conflict resolution strategies can be configured for specific links and/or sets of tables.

5.1.2. Multi-Tiered Synchronization

As shown in Section 3.2, Organizing Nodes (p. 23) , there may be scenarios where data needs to flow through multiple tiers of nodes that are organized in a tree-like network with each tier requiring a different subset of data. For example, you may have a system where the lowest tier may by a computer or device located in a store. Those devices may connect to a server located physically at that store. Then the store server may communicate with a corporate server for example. In this case, the three tiers would be device, store, and corporate. Each tier is typically represented by a node group. Each node in the tier would belong to the node group representing that tier.

A node will always push and pull data to other node groups according to the node group link configuration. A node can only pull and push data to other nodes that are represented `node` table in its database and having `sync_enabled = 1`. Because of this, a tree-like hierarchy of nodes can be created by having only a subset of nodes belonging to the same node group represented at the different branches of the tree.

If auto registration is turned *off*, then this setup must occur manually by opening registration for the desired nodes at the desired parent node and by configuring each node's `registration.url` to be the parent node's URL. The parent node is always tracked by the setting of the parent's `node_id` in the `created_at_node_id` column of the new node. When a node registers and downloads its configuration it is always provided the configuration for nodes that might register with the node itself based on the Node Group Links defined in the parent node.

5.1.2.1. Registration Redirect

When deploying a multi-tiered system it may be advantageous to have only one registration server, even though the parent node of a registering node could be any of a number of nodes in the system. In SymmetricDS the parent node is always the node that a child registers with. The REGISTRATION_REDIRECT table allows a single node, usually the root server in the network, to redirect registering nodes to their true parents. It does so based on a mapping found in the table of the external id (`registrant_external_id`) to the parent's node id (`registration_node_id`).

For example, if it is desired to have a series of regional servers that workstations at retail stores get assigned to based on their `external_id`, the store number, then you might insert into REGISTRATION_REDIRECT the store number as the `registrant_external_id` and the `node_id` of the assigned region as the `registration_node_id`. When a workstation at the store registers, the root server send an HTTP redirect to the `sync_url` of the node that matches the `registration_node_id`.

> **Important**
>
> Please see Section 4.6.3.2, Initial Loads (p. 47) for important details around initial loads and registration when using registration redirect.

5.2. Jobs

The SymmetricDS software allows for outgoing and incoming changes to be synchronized to/from other databases. The node that initiates a synchronization connection is the client, and the node receiving a connection is the host. Because synchronization is configurable to push or pull in either direction, the same node can act as either a client or a host in different circumstances.

The SymmetricDS software consists of a series of background jobs, managers, Servlets, and services wired together via dependency injection using the Spring Framework.

As a client, the node runs the router job, push job and pull job on a timer thread. The router job uses services to create batches that are targeted at certain nodes. The push job uses services to extract and stream data to another node (that is, it pushes data). The response from a push is a list of batch acknowlegements to indicate that data was loaded. The pull job uses services to load data that is streamed from another node (*i.e.*, it pulls data). After loading data, a second connection is made to send a list of batch acknowlegements.

As a host, the node waits for incoming connections that pull, push, or acknowledge data changes. The push Servlet uses services to load data that is pushed from a client node. After loading data, it responds with a list of batch acknowledgements. The pull Servlet uses services to extract, and stream data back to the client node. The ack Servlet uses services to update the status of data that was loaded at a client node. The router job batches and routes data.

By default, data is extracted from the source database into memory until a threshold size is reached. If the threshold size is reached, data is streamed to a temporary file in the JVM's default temporary directory. Next, the data is streamed to the target node across the transport layer. The receiving node will cache the data in memory until the threshold size is reached, writing to a temporary file if necessary. At last, the data is loaded into the target database by the data loader. This step by step approach allows for extract time, transport time, and load time to all be measured independently. It also allows database resources to

be used most optimally.

The transport manager handles the incoming and outgoing streams of data between nodes. The default transport is based on a simple implementation over HTTP. An internal transport is also provided. It is possible to add other implementations, such as a socket-based transport manager.

Node communication over HTTP is represented in the following figure.
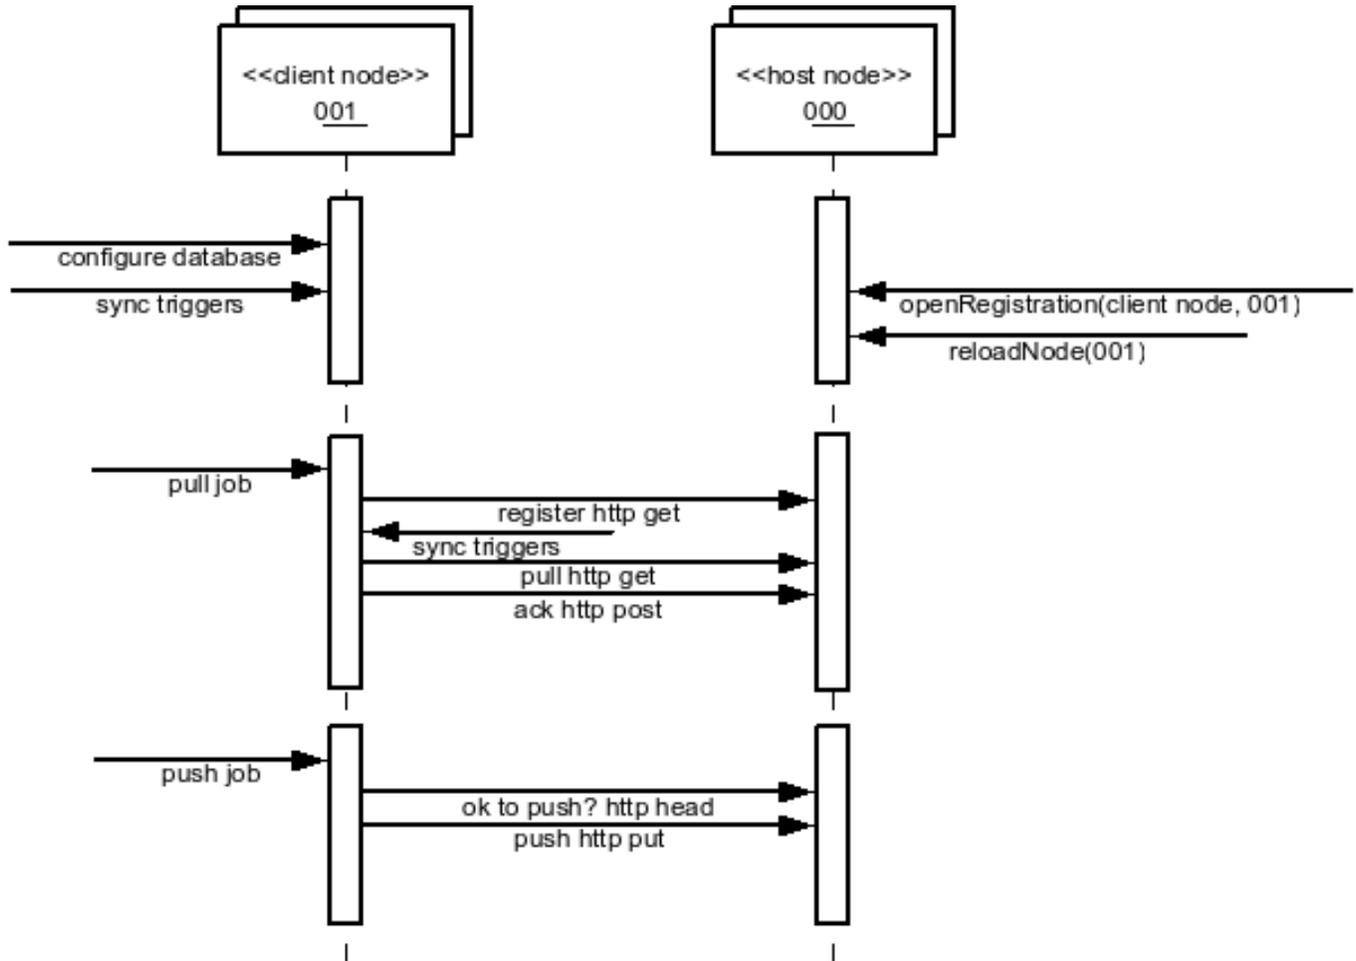
Figure 5.1. Node Communication

The `StandaloneSymmetricEngine` is wrapper API that can be used to directly start the client services only. The `SymmetricWebServer` is a wrapper API that can be used to directly start *both* the client and host services inside a Jetty web container. The `SymmetricLauncher` provides command line tools to work with and start SymmetricDS.

5.2.1. Route Job

5.2.1.1. Overview

The SymmetricDS-created database triggers cause data to be capture in the DATA table. The next step in the synchronization process is to process the change data to determine which nodes, if any, the data should be routed to. This step is performed by the *Route Job*. In addition to determining which nodes data will be sent to, the Route Job is also responsible for determing how much data will be batched together for transport. It is a single background task that inserts into DATA_EVENT and OUTGOING_BATCH.

At a high level, the Route Job is straightforward. It collects a list of data ids from DATA which haven't yet been routed (see Section 5.2.1.2, Data Gaps (p. 64) for much more detail about this step), one channel at a time, up to a limit specified by the channel configuration (`max_data_to_route`, on CHANNEL). The data is then batched based on the `batch_algorithm` defined for the channel and as documented in Section 4.5, Channel (p. 36) . Note that, for the default batching algorithm, there may actually be more than `max_data_to_route` included depending on the transaction boundaries. The mapping of data to specific nodes, organized into batches, is then recorded in OUTGOING_BATCH with a status of "RT" in each case (representing the fact that the Route Job is still running). Once the routing algorithms and batching are completed, the batches are organized with their corresponding data ids and saved in DATA_EVENT. Once DATA_EVENT is updated, the rows in OUTGOING_BATCH are updated to a status of New "NE".

5.2.1.2. Data Gaps

On the surface, the first Route Job step of collecting unrouted data ids seems simple: assign sequential data ids for each data row as it's inserted and keep track of which data id was last routed and start from there. The difficulty arises, however, due to the fact that there can be multiple transactions inserting into DATA simultaneously. As such, a given section of rows in the DATA table may actually contain "gaps" in the data ids when the Route Job is executing. Most of these gaps are only temporarily and fill in at some point after routing and need to be picked up with the next run of the Route Job. Thus, the Route Job needs to remember to route the filled-in gaps. Worse yet, some of these gaps are actually permanent and result from a transaction that is rolled back for some reason. In this case, the Route Job must continue to watch for the gap to fill in and, at some point, eventually gives up and assumes the gap is permanent and can be skipped. All of this must be done in some fashion that guarantees that gaps are routed when they fill in while also keeping routing as efficient as possible.

SymmetricDS handles the issue of data gaps by making use of a table, DATA_GAP, to record gaps found in the data ids. In fact, this table completely defines the entire range of data tha can be routed at any point in time. For a brand new instance of SymmetricDS, this table is empty and SymmetricDS creates a gap starting from data id of zero and ending with a very large number (defined by `routing.largest.gap.size`). At the start of a Route Job, the list of valid gaps (gaps with status of 'GP') is collected, and each gap is evaluated in turn. If a gap is sufficiently old (as defined by `routing.stale.dataid.gap.time.ms`, the gap is marked as skipped (status of 'SK') and will no longer be evaluated in future Route Jobs (note that the 'last' gap (the one with the highest starting data id) is never skipped). If not skipped, then DATA_EVENT is searched for data ids present in the gap. If one or more data ids is found in DATA_EVENT, then the current gap is marked with a status of OK, and new gap(s) are created to represent the data ids still missing in the gap's range. This process is done for all gaps. If the very last gap contained data, a new gap starting from the highest data id and ending at (highest data id + `routing.largest.gap.size`) is then created. This process has resulted in an updated list of gaps which may contain new data to be routed.

5.2.2. Synchronization Frequency

The frequency of data synchronization is controlled by the coordination of a series of asynchronous events.

The route job determines which nodes data will be sent to, and batches it together for transport. When the `start.route.job` SymmetricDS property is set to `true`, the frequency that routing occurs is controlled by the `job.routing.period.time.ms`.

After data is routed, it awaits transport to the target nodes. Transport can occur when a client node is configured to pull data or when the host node is configured to push data. These events are controlled by the *push* and the *pull jobs*. When the `start.pull.job` SymmetricDS property is set to `true`, the frequency that data is pulled is controlled by the `job.pull.period.time.ms`. When the `start.push.job` SymmetricDS property is set to `true`, the frequency that data is pushed is controlled by the `job.push.period.time.ms`. Data is extracted by channel from the source database's DATA table at an interval controlled by the `extract_period_millis` column on the CHANNEL table. The `last_extract_time` is always recorded, by channel, on the NODE_CHANNEL_CTL table for the host node's id. When the Pull and Push Job run, if the extract period has not passed according to the last extract time, then the channel will be skipped for this run. If the `extract_period_millis` is set to zero, data extraction will happen every time the jobs run.

Both the push and pull jobs can be configured to push and pull from multiple nodes in parallel. In order to take advantage of this the `pull.thread.per.server.count` or `push.thread.per.server.count` should be adjusted (from their default value of 10) to the number to the number of concurrent push/pulls you want to occur per period on each SymmetricDS instance. Push and pull activity is recorded in the NODE_COMMUNICATION table. This table is also used to lock push and pull activity across multiple servers in a cluster.

SymmetricDS also provides the ability to configure windows of time when synchronization is allowed. This is done using the NODE_GROUP_CHANNEL_WINDOW table. A list of allowed time windows can be specified for a node group and a channel. If one or more windows exist, then data will only be extracted and transported if the time of day falls within the window of time specified. The configured times are always for the target node's local time. If the `start_time` is greater than the `end_time`, then the window crosses over to the next day.

All data loading may be disabled by setting the `dataloader.enable` property to false. This has the effect of not allowing incoming synchronizations, while allowing outgoing synchronizations. All data extractions may be disabled by setting the `dataextractor.enable` property to false. These properties can be controlled by inserting into the root server's PARAMETER table. These properties affect every channel with the exception of the 'config' channel.

5.2.3. Sync Triggers Job

SymmetricDS examines the current configuration, corresponding database triggers, and the underlying tables to determine if database triggers need created or updated. The change activity is recorded on the TRIGGER_HIST table with a reason for the change. The following reasons for a change are possible:

- N - New trigger that has not been created before

- S - Schema changes in the table were detected

- C - Configuration changes in Trigger

- T - Trigger was missing

A configuration entry in Trigger without any history in Trigger Hist results in a new trigger being created (N). The Trigger Hist stores a hash of the underlying table, so any alteration to the table causes the trigger to be rebuilt (S). When the `last_update_time` is changed on the Trigger entry, the configuration change causes the trigger to be rebuilt (C). If an entry in Trigger Hist is missing the corresponding database trigger, the trigger is created (T).

The process of examining triggers and rebuilding them is automatically run during startup and each night by the SyncTriggersJob. The user can also manually run the process at any time by invoking the `syncTriggers()` method over JMX. The SyncTriggersJob is enabled by default to run at 15 minutes past midnight. If SymmetricDS is being run from a collection of servers (multiple instances of the same Node running against the same database), then locking should be enable to prevent database contention. The following runtime properties control the behavior of the process.

start.synctriggers.job
Whether the sync triggers job is enabled for this node. [Default: true]

job.synctriggers.aftermidnight.minutes
If scheduled, the sync triggers job will run nightly. This is how long after midnight that job will run. [Default: 15]

cluster.lock.during.sync.triggers
Indicate if the sync triggers job is clustered and requires a lock before running. [Default: false]

5.3. JMS Publishing

With the proper configuration SymmetricDS can publish XML messages of captured data changes to JMS during routing or transactionally while data loading synchronized data into a target database. The following explains how to publish to JMS during synchronization to the target database.

The XmlPublisherDatabaseWriterFilter is a IDatabaseWriterFilter that may be configured to publish specific tables as an XML message to a JMS provider. See Section 5.10, Extension Points (p. 76) for information on how to configure an extension point. If the publish to JMS fails, the batch will be marked in error, the loaded data for the batch will be rolled back and the batch will be retried during the next synchronization run.

The following is an example extension point configuration that will publish four tables in XML with a root tag of *'sale'*. Each XML message will be grouped by the batch and the column names identified by the groupByColumnNames property which have the same values.

```xml
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:context="http://www.springframework.org/schema/context"
    xsi:schemaLocation="http://www.springframework.org/schema/beans
           http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
           http://www.springframework.org/schema/context
           http://www.springframework.org/schema/context/spring-context-3.0.xsd">
```

```
    <bean id="configuration-publishingFilter"
      class="org.jumpmind.symmetric.integrate.XmlPublisherDatabaseWriterFilter">
        <property name="xmlTagNameToUseForGroup" value="sale"/>
        <property name="tableNamesToPublishAsGroup">
            <list>
                <value>SALE_TX</value>
                <value>SALE_LINE_ITEM</value>
                <value>SALE_TAX</value>
                <value>SALE_TOTAL</value>
            </list>
        </property>
        <property name="groupByColumnNames">
            <list>
                <value>STORE_ID</value>
                <value>BUSINESS_DAY</value>
                <value>WORKSTATION_ID</value>
                <value>TRANSACTION_ID</value>
            </list>
        </property>
        <property name="publisher">
            <bean class="org.jumpmind.symmetric.integrate.SimpleJmsPublisher">
                <property name="jmsTemplate" ref="definedSpringJmsTemplate"/>
            </bean>
        </property>
    </bean>
</beans>
```

The publisher property on the XmlPublisherDatabaseWriterFilter takes an interface of type IPublisher. The implementation demonstrated here is an implementation that publishes to JMS using Spring's JMS template. Other implementations of IPublisher could easily publish the XML to other targets like an HTTP server, the file system or secure copy it to another server.

The above configuration will publish XML similiar to the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<sale xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  id="0012010-01-220031234" nodeid="00001" time="1264187704155">
  <row entity="SALE_TX" dml="I">
    <data key="STORE_ID">001</data>
    <data key="BUSINESS_DAY">2010-01-22</data>
    <data key="WORKSTATION_ID">003</data>
    <data key="TRANSACTION_ID">1234</data>
    <data key="CASHIER_ID">010110</data>
  </row>
  <row entity="SALE_LINE_ITEM" dml="I">
    <data key="STORE_ID">001</data>
    <data key="BUSINESS_DAY">2010-01-22</data>
    <data key="WORKSTATION_ID">003</data>
    <data key="TRANSACTION_ID">1234</data>
    <data key="SKU">9999999</data>
    <data key="PRICE">10.00</data>
    <data key="DESC" xsi:nil="true"/>
  </row>
  <row entity="SALE_LINE_ITEM" dml="I">
    <data key="STORE_ID">001</data>
    <data key="BUSINESS_DAY">2010-01-22</data>
    <data key="WORKSTATION_ID">003</data>
    <data key="TRANSACTION_ID">1234</data>
    <data key="SKU">9999999</data>
```

```
      <data key="PRICE">10.00</data>
      <data key="DESC" xsi:nil="true"/>
    </row>
    <row entity="SALE_TAX" dml="I">
      <data key="STORE_ID">001</data>
      <data key="BUSINESS_DAY">2010-01-22</data>
      <data key="WORKSTATION_ID">003</data>
      <data key="TRANSACTION_ID">1234</data>
      <data key="AMOUNT">1.33</data>
    </row>
    <row entity="SALE_TOTAL" dml="I">
      <data key="STORE_ID">001</data>
      <data key="BUSINESS_DAY">2010-01-22</data>
      <data key="WORKSTATION_ID">003</data>
      <data key="TRANSACTION_ID">1234</data>
      <data key="AMOUNT">21.33</data>
    </row>
  </sale>
```

To publish JMS messages during routing the same pattern is valid, with the exception that the extension point would be the XmlPublisherDataRouter and the router would be configured by setting the router_type of a ROUTER to the Spring bean name of the registered extension point. Of course, the router would need to be linked through TRIGGER_ROUTERs to each TRIGGER table that needs published.

5.4. Deployment Options

An instance of SymmetricDS can be deployed in several ways:

- Web application archive (WAR) deployed to an application server

  This option means packaging a WAR file and deploying to your favorite web server, like Apache Tomcat. It's a little more work, but you can configure the web server to do whatever you need. SymmetricDS can also be embedded in an existing web application, if desired.

- Standalone service that embeds Jetty web server

  This option means running the *sym* command line, which launches the built-in Jetty web server. This is a simple option because it is already provided, but you lose the flexibility to configure the web server any further.

- Embedded as a Java library in an application

  This option means you must write a wrapper Java program that runs SymmetricDS. You would probably use Jetty web server, which is also embeddable. You could bring up an embedded database like Derby or H2. You could configure the web server, database, or SymmetricDS to do whatever you needed, but it's also the most work of the three options discussed thus far.

The deployment model you choose depends on how much flexibility you need versus how easy you want it to be. Both Jetty and Tomcat are excellent, scalable web servers that compete with each other and have great performance. Most people choose either the *Standalone* or *Web Archive* with Tomcat 5.5 or 6.

Deploying to Tomcat is a good middle-of-the-road decision that requires a little more work for more flexibility.

Next, we will go into a little more detail on the first three deployment options listed above.

5.4.1. Web Archive (WAR)

As a web application archive, a WAR is deployed to an application server, such as Tomcat, Jetty, or JBoss. The structure of the archive will have a `web.xml` file in the `WEB-INF` folder, an appropriately configured `symmetric.properties` file in the `WEB-INF/classes` folder, and the required JAR files in the `WEB-INF/lib` folder.
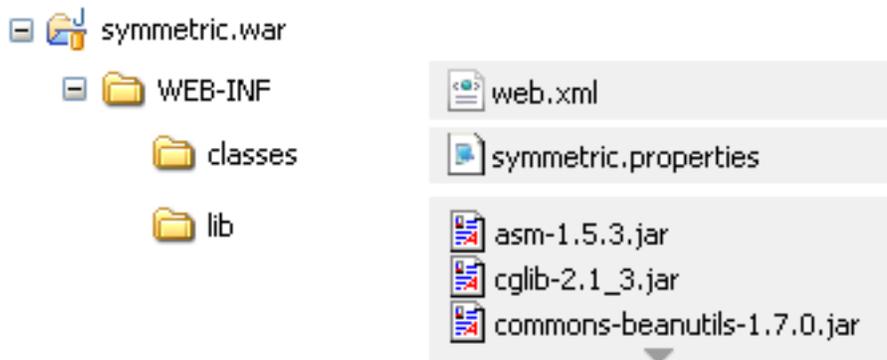
A war file can be generated using the standalone installation's `symadmin` utility and the `create-war` subcommand. The command requires the name of the war file to generate. It essentially packages up the web directory, the conf directory and includes an optional properties file. Note that if a properties file is included, it will be copied to WEB-INF/classes/symmetric.properties. This is the same location conf/symmetric.properties would have been copied to. The generated war distribution uses the same web.xml as the standalone deployment.

../bin/symadmin -p my-symmetric-ds.properties create-war /some/path/to/symmetric-ds.war

The `web.base.servlet.path` property in `symmetric.properties` can be set if the SymmetricServlet needs to coexist with other Servlets. By default, the value is blank. If you set it to, say, `web.base.servlet.path=sync` for exmaple, `registration.url` would be `http://server:port/sync`.

5.4.2. Standalone

A standalone service can use the `sym` command line options to start a server. An embedded instance of Jetty is used to service web requests for all the servlets.

```
/symmetric/bin/sym --properties root.properties --port 8080 --server
```

This example starts the SymmetricDS server on port 8080 with the startup properties found in the `root.properties` file.

5.4.3. Embedded

A Java application with the SymmetricDS Java Archive (JAR) library on its classpath can use the `SymmetricWebServer` to start the server.

```
import org.jumpmind.symmetric.SymmetricWebServer;

public class StartSymmetricEngine {

    public static void main(String[] args) throws Exception {

        SymmetricWebServer node = new SymmetricWebServer(
                             "classpath://my-application.properties", "conf/web_dir");

        // this will create the database, sync triggers, start jobs running
        node.start(8080);

        // this will stop the node
        node.stop();
    }

}
```

This example starts the SymmetricDS server on port 8080. The configuration properties file, `my-application.properties`, is packaged in the application to provide properties that override the SymmetricDS default values. The second parameter to the constructor points to the web directory. The default location is `../web`. In this example the web directory is located at conf/web_dir. The web.xml is expected to be found at conf/web_dir/WEB-INF/web.xml.

5.5. Running SymmetricDS as a Service

SymmetricDS can be configured to start and run as a service in both Windows and *nix platforms.

5.5.1. Running as a Windows Service

SymmetricDS uses the [Java Service Wrapper](#) product from Tanuki Software to run in the background as a Windows system service. The Java Service Wrapper executable is named `sym_service.exe` so it can be easily identified from a list of running processes. To install the service, use the provided script:

```
bin\install_service.bat
```

The service configuration is found in `conf/sym_service.conf`. Edit this file if you want to change the default port number (8080), initial memory size (256 MB), log file size (10 MB), or other settings. When started, the server will look in the `conf` directory for the `symmetric.properties` file and the `log4j.xml` file. Logging for standard out, error, and application are written to the `logs` directory.

Most configuration changes do not require the service to be re-installed. To un-install the service, use the provided script:

```
bin\uninstall_service.bat
```

Use the **net** command to start and stop the service:

```
net start symmetricds
net stop symmetricds
```

5.5.2. Running as a *nix Service

SymmetricDS uses the [Java Service Wrapper](#) product from Tanuki Software to run in the background as a Unix system service. The Java Service Wrapper executable is named `sym_service` so it can be easily identified from a list of running processes. The service configuration is found in `conf/sym_service.conf`. Edit this file if you want to change the default port number (8080), initial memory size (256 MB), log file size (10 MB), or other settings.

An init script is provided to work with standard Unix run configuration levels. The `sym_service.initd` file follows the Linux Standard Base specification, which should work on many systems, including Fedora and Debian-based distributions. To install the script, copy it into the system init directory:

```
cp bin/sym_service.initd /etc/init.d/sym_service
```

Edit the init script to set the SYM_HOME variable to the directory where SymmetricDS is located. The init script calls the `sym_service` executable.

Enabling the service varies based on the version of Linux in use. Three possible approaches are listed below:

1. Using `chkconfig` command:

   To enable the service to run automatically when the system is started:

```
/sbin/chkconfig --add sym_service
```

   To disable the service from running automatically:

```
/sbin/chkconfig --del sym_service
```

2. Using `install_initd` command (Suse Linux):

   On Suse Linux install the service by calling:

```
/usr/lib/lsb/install_initd sym_service
```

   Remove the service by calling:

```
/usr/lib/lsb/remove_initd sym_service
```

3. Using `sysv-rc-conf` command (Ubuntu Linux):

   On Ubuntu Linux, you might need to use sysv-rc-conf instead of chkconfig. Try running sys-rc-conf as a super user (consider utilizing apt-get to install sysv-rc-conf if it is not present: `sudo apt-get install sysv-rc-conf`). Run sysv-rc-conf with the following command:

```
sudo sysv-rc-conf
```

   You should see a list of the scripts residing in your /etc/init.d folder. Use control-N to navigate through the list to locate sym_service, then activate the service for the desired run-levels (most likely 2-5).

Finally, you can use the **service** command to start, stop, and query the status of the service:

```
/sbin/service sym_service start
/sbin/service sym_service stop
/sbin/service sym_service status
```

Alternatively, call the init.d script directly:

```
/etc/init.d/sym_service start
/etc/init.d/sym_service stop
/etc/init.d/sym_service status
```

5.6. Clustering

A single SymmetricDS node may be clustered across a series of instances, creating a web farm. A node might be clustered to provide load balancing and failover, for example.

When clustered, a hardware load balancer is typically used to round robin client requests to the cluster. The load balancer should be configured for stateless connections. Also, the `sync.url` (discussed in Section 4.1, Node Properties (p. 33)) SymmetricDS property should be set to the URL of the load balancer.

If the cluster will be running any of the SymmetricDS jobs, then the `cluster.lock.enabled` property should be set to `true`. By setting this property to true, SymmetricDS will use a row in the LOCK table as a semaphore to make sure that only one instance at a time runs a job. When a lock is acquired, a row is updated in the lock table with the time of the lock and the server id of the locking job. The lock time is set back to null when the job is finished running. Another instance of SymmetricDS cannot aquire a lock until the locking instance (according to the server id) releases the lock. If an instance is terminated while the lock is still held, an instance with the same server id is allowed to reaquire the lock. If the locking instance remains down, the lock can be broken after a period of time, specified by the `cluster.lock.timeout.ms` property, has expired. Note that if the job is still running and the lock expires, two jobs could be running at the same time which could cause database deadlocks.

By default, the locking server id is the hostname of the server. If two clustered instances are running on the same server, then the `cluster.server.id` property may be set to indicate the name that the instance should use for its server id.

When deploying SymmetricDS to an application server like Tomcat or JBoss, no special session clustering needs to be configured for the application server.

5.7. Encrypted Passwords

The `db.user` and `db.password` properties will accept encrypted text, which protects against casual observation. The text is prefixed with `enc:` to indicate that it is encrypted. To encrypt text, use the following command:

```
sym -e secret
```

The text is encrypted using a secret key named "sym.secret" that is retrieved from a keystore file. By default, the keystore is located in `security/keystore`. The location and filename of the keystore can be overridden by setting the "sym.keystore.file" system property. If the secret key is not found, the system will generate and install a secret key for use with Triple DES cipher.

Generate a new secret key for encryption using the `keytool` command that comes with the JRE. If there is an existing key in the keystore, first remove it:

```
keytool -keystore keystore -storepass changeit -storetype jceks \
    -alias sym.secret -delete
```

Then generate a secret key, specifying a cipher algorithm and key size. Commonly used algorithms that are supported include aes, blowfish, desede, and rc4.

```
keytool -keystore keystore -storepass changeit -storetype jceks \
    -alias sym.secret -genseckey -keyalg aes -keysize 128
```

If using an alternative provider, place the provider JAR file in the SymmetricDS `lib` folder. The provider class name should be installed in the JRE security properties or specified on the command line. To install in the JRE, edit the JRE `lib/security/java.security` file and set a `security.provider.i` property for the provider class name. Or, the provider can be specified on the command line instead. Both `keytool` and `sym` accept command line arguments for the provider class name. For example, using the Bouncy Castle provider, the command line options would look like:

```
keytool -keystore keystore -storepass changeit -storetype jceks \
    -alias sym.secret -genseckey -keyalg idea -keysize 56 \
    -providerClass org.bouncycastle.jce.provider.BouncyCastleProvider \
    -providerPath ..\lib\bcprov-ext.jar
```

```
symadmin -providerClass org.bouncycastle.jce.provider.BouncyCastleProvider -e secret
```

To customize the encryption, write a Java class that implements the ISecurityService or extends the default SecurityService, and place the class on the classpath in either `lib` or `web/WEB-INF/lib` folders.

Then, in the `symmetric.properties` specify your class name for the security service.

```
security.service.class.name=org.jumpmind.security.SecurityService
```

Remember to specify your properties file when encrypting passwords, so it will use your custom ISecurityService.

```
symadmin -p symmetric.properties -e secret
```

5.8. Secure Transport

By specifying the "https" protocol for a URL, SymmetricDS will communicate over Secure Sockets Layer (SSL) for an encrypted transport. The following properties need to be set with "https" in the URL:

sync.url
This is the URL of the current node, so if you want to force other nodes to communicate over SSL with this node, you specify "https" in the URL.

registration.url
This is the URL where the node will connect for registration when it first starts up. To protect the registration with SSL, you specify "https" in the URL.

For incoming HTTPS connections, SymmetricDS depends on the webserver where it is deployed, so the webserver must be configured for HTTPS. As a standalone deployment, the "sym" launcher command provides options for enabling HTTPS support.

5.8.1. Sym Launcher

The "sym" launch command uses Jetty as an embedded web server. Using command line options, the web server can be told to listen for HTTP, HTTPS, or both.

sym --port 8080 --server

sym --secure-port 8443 --secure-server

sym --port 8080 --secure-port 8443 --mixed-server

5.8.2. Tomcat

If you deploy SymmetricDS to Apache Tomcat, it can be secured by editing the `TOMCAT_HOME/conf/server.xml` configuration file. There is already a line that can be uncommented and changed to the following:

```
<Connector port="8443" protocol="HTTP/1.1" SSLEnabled="true"
  maxThreads="150" scheme="https" secure="true"
  clientAuth="false" sslProtocol="TLS"
  keystoreFile="/symmetric-ds-1.x.x/security/keystore" />
```

5.8.3. Keystores

When SymmetricDS connects to a URL with HTTPS, Java checks the validity of the certificate using the built-in trusted keystore located at `JRE_HOME/lib/security/cacerts`. The "sym" launcher command overrides the trusted keystore to use its own trusted keystore instead, which is located at `security/cacerts`. This keystore contains the certificate aliased as "sym" for use in testing and easing deployments. The trusted keystore can be overridden by specifying the `javax.net.ssl.trustStore` system property.

When SymmetricDS is run as a secure server with the "sym" launcher, it accepts incoming requests using the key installed in the keystore located at `security/keystore`. The default key is provided for convenience of testing, but should be re-generated for security.

5.8.4. Generating Keys

To generate new keys and install a server certificate, use the following steps:

1. Open a command prompt and navigate to the `security` subdirectory of your SymmetricDS installation on the server to which communication will be secured (typically the "root" or "central office" server).

2. Delete the old key pair and certificate.

    keytool -keystore keystore -delete -alias sym

    keytool -keystore cacerts -delete -alias sym

```
Enter keystore password:   changeit
```

3. Generate a new key pair. Note that the first name/last name (the "CN") must match the fully qualified hostname the client will be using to communcate to the server.

    keytool -keystore keystore -alias sym -genkey -keyalg RSA -validity 10950

```
Enter keystore password:   changeit
What is your first and last name?
  [Unknown]:   localhost
What is the name of your organizational unit?
  [Unknown]:   SymmetricDS
What is the name of your organization?
  [Unknown]:   JumpMind
What is the name of your City or Locality?
  [Unknown]:
What is the name of your State or Province?
  [Unknown]:
What is the two-letter country code for this unit?
  [Unknown]:
Is CN=localhost, OU=SymmetricDS, O=JumpMind, L=Unknown, ST=Unknown, C=Unknown
correct?
  [no]:   yes

Enter key password for <sym>
```

```
         (RETURN if same as keystore password):
```

4. Export the certificate from the private keystore.

   keytool -keystore keystore -export -alias sym -rfc -file sym.cer

5. Install the certificate in the trusted keystore.

   keytool -keystore cacerts -import -alias sym -file sym.cer

6. Copy the cacerts file that is generated by this process to the `security` directory of each client's SymmetricDS installation.

5.9. Basic Authentication

SymmetricDS supports basic authentication for client and server nodes.

To configure a client node to use basic authentication when communicating with a server node, specify the following startup parameters:

http.basic.auth.username
username for client node basic authentication. [Default:]

http.basic.auth.password
password for client node basic authentication. [Default:]

The SymmetricDS Standalone Web Server also supports Basic Authentication. It can be enabled by passing the following arguments to the startup program

--http-basic-auth-user
username for basic authentication [Default:]

--http-basic-auth-password
password for basic authentication [Default:]

If the server node is deployed to Tomcat or another application server as a WAR or EAR file, then basic authentication is setup with the standard configuration in the WEB.xml file.

5.10. Extension Points

SymmetricDS has a pluggable architecture that can be extended. A Java class that implements the appropriate extension point interface, can implement custom logic and change the behavior of SymmetricDS to suit special needs. All supported extension points extend the `IExtensionPoint` interface.

The available extension points are documented in the following sections.

When SymmetricDS starts up, the `ExtensionPointManager` searches a [Spring Framework](#) context for classes that implement the `IExtensionPoint` interface, then creates and registers the class with the appropriate SymmetricDS component.

If an extension point needs access to SymmetricDS services or needs to connect to the database it may implement the `ISymmetricEngineAware` interface in order to get a handle to the `ISymmetricEngine`.

The `INodeGroupExtensionPoint` interface may be optionally implemented to indicate that a registered extension point should only be registered with specific node groups.

```
/**
 * Only apply this extension point to the 'root' node group.
 */
public String[] getNodeGroupIdsToApplyTo() {
    return new String[] { "root" };
}
```

Extensions are configured in the `conf/symmetric-extensions.xml` file.

5.10.1. IParameterFilter

Parameter values can be specified in code using a parameter filter. Note that there can be only one parameter filter per engine instance. The IParameterFilter replaces the deprecated IRuntimeConfig from prior releases.

```
public class MyParameterFilter
    implements IParameterFilter, INodeGroupExtensionPoint {

    /**
     * Only apply this filter to stores
     */
    public String[] getNodeGroupIdsToApplyTo() {
        return new String[] { "store" };
    }

    public String filterParameter(String key, String value) {
        // look up a store number from an already existing properties file.
        if (key.equals(ParameterConstants.EXTERNAL_ID)) {
            return StoreProperties.getStoreProperties().
                getProperty(StoreProperties.STORE_NUMBER);
        }
        return value;
    }

    public boolean isAutoRegister() {
        return true;
    }

}
```

5.10.2. IDatabaseWriterFilter

Data can be filtered or manipulated before it is loaded into the target database. A filter can change the data in a column, save it somewhere else or do something else with the data entirely. It can also specify by the return value of the function call that the data loader should continue on and load the data (by returning true) or ignore it (by returning false). One possible use of the filter, for example, might be to route credit card data to a secure database and blank it out as it loads into a less-restricted reporting database.

A `DataContext` is passed to each of the callback methods. A new context is created for each synchronization. The context provides a mechanism to share data during the load of a batch between different rows of data that are committed in a single database transaction.

The filter also provide callback methods for the batch lifecycle. The `DatabaseWriterFilterAdapter` may be used if not all methods are required.

A class implementing the IDatabaseWriterFilter interface is injected onto the DataLoaderService in order to receive callbacks when data is inserted, updated, or deleted.

```java
public class MyFilter extends DatabaseWriterFilterAdapter {

    @Override
    public boolean beforeWrite(DataContext context, Table table, CsvData data) {
        if (table.getName().equalsIgnoreCase("CREDIT_CARD_TENDER")
                && data.getDataEventType().equals(DataEventType.INSERT)) {
            String[] parsedData = data.getParsedData(CsvData.ROW_DATA);
            // blank out credit card number
            parsedData[table.getColumnIndex("CREDIT_CARD_NUMBER")] = null;
        }
        return true;
    }
}
```

The filter class should be specified in `conf/symmetric-extensions.xml` as follows.

```xml
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:context="http://www.springframework.org/schema/context"
    xsi:schemaLocation="http://www.springframework.org/schema/beans
            http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
            http://www.springframework.org/schema/context
            http://www.springframework.org/schema/context/spring-context-3.0.xsd">

    <bean id="myFilter" class="com.mydomain.MyFilter"/>

</beans>
```

5.10.3. IDatabaseWriterErrorHandler

Implement this extension point to override how errors are handled. You can use this extension point to ignore rows that produce foreign key errors.

5.10.4. IDataLoaderFactory

Implement this extension point to provide a different implementation of the `org.jumpmind.symmetric.io.data.IDataWriter` that is used by the SymmetricDS data loader. Data loaders are configured for a channel. After this extension point is registered it can be activated for a CHANNEL by indicating the data loader name in the `data_loader_type` column.

One possible use of this extension point is to route data to a NOSQL data sink.

5.10.5. IAcknowledgeEventListener

Implement this extension point to receive callback events when a batch is acknowledged. The callback for this listener happens at the point of extraction.

5.10.6. IReloadListener

Implement this extension point to listen in and take action before or after a reload is requested for a Node. The callback for this listener happens at the point of extraction.

5.10.7. ISyncUrlExtension

This extension point is used to select an appropriate URL based on the URI provided in the `sync_url` column of `sym_node`.

To use this extension point configure the sync_url for a node with the protocol of ext://beanName. The beanName is the name you give the extension point in the extension xml file.

5.10.8. IColumnTransform

This extension point allows custom column transformations to be created. There are a handful of out-of-the-box implementations. If any of these do not meet the column transformation needs of the application, then a custom transform can be created and registered. It can be activated by referencing the column transform's name `transform_type` column of TRANSFORM_COLUMN

5.10.9. INodeIdCreator

This extension point allows SymmetricDS users to implement their own algorithms for how node ids and passwords are generated or selected during the registration process. There may be only one node creator per SymmetricDS instance (Please note that the node creator extension has replaced the node generator extension).

5.10.10. ITriggerCreationListener

Implement this extension point to get status callbacks during trigger creation.

5.10.11. IBatchAlgorithm

Implement this extension point and set the name of the Spring bean on the batch_algorithm column of the Channel table to use. This extension point gives fine grained control over how a channel is batched.

5.10.12. IDataRouter

Implement this extension point and set the name of the Spring bean on the router_type column of the Router table to use. This extension point gives the ability to programatically decide which nodes data should be routed to.

5.10.13. IHeartbeatListener

Implement this extension point to get callbacks during the heartbeat job.

5.10.14. IOfflineClientListener

Implement this extension point to get callbacks for offline events on client nodes.

5.10.15. IOfflineServerListener

Implement this extension point to get callbacks for offline events detected on a server node during monitoring of client nodes.

5.10.16. INodePasswordFilter

Implement this extension point to intercept the saving and rendering of the node password.

Chapter 6. Administration

6.1. Solving Synchronization Issues

By design, whenever SymmetricDS encounters an issue with a synchronization, the batch containing the error is marked as being in an error state, and all subsequent batches *for that particular channel to that particular node* are held and not synchronized until the error batch is resolved. SymmetricDS will retry the batch in error until the situation creating the error is resolved (or the data for the batch itself is changed).

Analyzing and resolving issues can take place on the outgoing or incoming side. The techniques for analysis are slightly different in the two cases, however, due to the fact that the node with outgoing batch data also has the data and data events associated with the batch in the database. On the incoming node, however, all that is available is the incoming batch header and data present in an incoming error table.

6.1.1. Analyzing the Issue - Outgoing Batches

The first step in analyzing the cause of a failed batch is to locate information about the data in the batch, starting with OUTGOING_BATCH To locate batches in error, use:

```
select * from sym_outgoing_batch where error_flag=1;
```

Several useful pieces of information are available from this query:

- The batch number of the failed batch, available in column BATCH_ID.

- The node to which the batch is being sent, available in column NODE_ID.

- The channel to which the batch belongs, available in column CHANNEL_ID. All subsequent batches on this channel to this node will be held until the error condition is resolved.

- The specific data id in the batch which is causing the failure, available in column FAILED_DATA_ID.

- Any SQL message, SQL State, and SQL Codes being returned during the synchronization attempt, available in columns SQL_MESSAGE, SQL_STATE, and SQL_CODE, respectively.

> **Note**
> Using the error_flag on the batch table, as shown above, is more reliable than using the status column. The status column can change from 'ER' to a different status temporarily as the batch is retried.

> **Note**
> The query above will also show you any recent batches that were originally in error and were changed to be manually skipped. See the end of Section 6.1.3, Resolving the Issue - Outgoing Batches (p. 83) for more details.

To get a full picture of the batch, you can query for information representing the complete list of all data changes associated with the failed batch by joining DATA and DATA_EVENT, such as:

```
select * from sym_data where data_id in
       (select data_id from sym_data_event where batch_id='XXXXXX');
```

where XXXXXX is the batch id of the failing batch.

This query returns a wealth of information about each data change in a batch, including:

- The table involved in each data change, available in column TABLE_NAME,

- The event type (Update [U], Insert [I], or Delete [D]), available in column EVENT_TYPE,

- A comma separated list of the new data and (optionally) the old data, available in columns ROW_DATA and OLD_DATA, respectively.

- The primary key data, available in column PK_DATA

- The channel id, trigger history information, transaction id if available, and other information.

More importantly, if you narrow your query to just the failed data id you can determine the exact data change that is causing the failure:

```
select * from sym_data where data_id in
       (select failed_data_id from sym_outgoing_batch where batch_id='XXXXX'
       and node_id='YYYYY');
```

where XXXXXX is the batch id and YYYYY is the node id of the batch that is failing.

The queries above usually yield enough information to be able to determine why a particular batch is failing. Common reasons a batch might be failing include:

- The schema at the destination has a column that is not nullable yet the source has the column defined as nullable and a data change was sent with the column as null.

- A foreign key constraint at the destination is preventing an insertion or update, which could be caused from data being deleted at the destination or the foreign key constraint is not in place at the source.

- The data size of a column on the destination is smaller than the data size in the source, and data that is too large for the destination has been synced.

6.1.2. Analyzing the Issue - Incoming Batches

Analysis using an incoming batch is different than that of outgoing batches. For incoming batches, you will rely on two tables, INCOMING_BATCH and INCOMING_ERROR. The first step in analyzing the cause of an incoming failed batch is to locate information about the batch, starting with INCOMING_BATCH To locate batches in error, use:

```
select * from sym_incoming_batch where error_flag=1;
```

Several useful pieces of information are available from this query:

- The batch number of the failed batch, available in column BATCH_ID. Note that this is the batch number of the outgoing batch on the outgoing node.

- The node the batch is being sent from, available in column NODE_ID.

- The channel to which the batch belongs, available in column CHANNEL_ID. All subsequent batches on this channel from this node will be held until the error condition is resolved.

- The data_id that was being processed when the batch failed, available in column FAILED_DATA_ID.

- Any SQL message, SQL State, and SQL Codes being returned during the synchronization attempt, available in columns SQL_MESSAGE, SQL_STATE, and SQL_CODE, respectively.

For incoming batches, we do not have data and data event entries in the database we can query. We do, however, have a table, INCOMING_ERROR, which provides some information about the batch.

```
select * from sym_incoming_error
          where batch_id='XXXXXX' and node_id='YYYYY';
```

where XXXXXX is the batch id and YYYYY is the node id of the failing batch.

This query returns a wealth of information about each data change in a batch, including:

- The table involved in each data change, available in column TARGET_TABLE_NAME,

- The event type (Update [U], Insert [I], or Delete [D]), available in column EVENT_TYPE,

- A comma separated list of the new data and (optionally) the old data, available in columns ROW_DATA and OLD_DATA, respectively,

- The column names of the table, available in column COLUMN_NAMES,

- The primary key column names of the table, available in column PK_COLUMN_NAMES,

6.1.3. Resolving the Issue - Outgoing Batches

Once you have decided upon the cause of the issue, you'll have to decide the best course of action to fix the issue. If, for example, the problem is due to a database schema mismatch, one possible solution would be to alter the destination database in such a way that the SQL error no longer occurs. Whatever approach you take to remedy the issue, once you have made the change, on the next push or pull SymmetricDS will retry the batch and the channel's data will start flowing again.

If you have instead decided that the batch itself is wrong, or does not need synchronized, or you wish to remove a particular data change from a batch, you do have the option of changing the data associated with the batch directly.

⚠ **Warning**

Be cautious when using the following two approaches to resolve synchronization issues. By far, the best approach to solving a synchronization error is to resolve what is truly causing the error at the destination database. Skipping a batch or removing a data id as discussed below should be your solution of last resort, since doing so results in differences between the source and destination databases.

Now that you've read the warning, if you *still* want to change the batch data itself, you do have several options, including:

- Causing SymmetricDS to skip the batch completely. This is accomplished by setting the batch's status to 'OK', as in:

```
update sym_outgoing_batch set status='OK' where batch_id='XXXXXX'
```

where XXXXXX is the failing batch. On the next pull or push, SymmetricDS will skip this batch since it now thinks the batch has already been synchronized. Note that you can still distinguish between successful batches and ones that you've artificially marked as 'OK', since the `error_flag` column on the failed batch will still be set to '1' (in error).

- Removing the failing data id from the batch by deleting the corresponding row in DATA_EVENT. Eliminating the data id from the list of data ids in the batch will cause future synchronization attempts of the batch to no longer include that particular data change as part of the batch. For example:

```
delete from sym_data_event where batch_id='XXXXXX' and data_id='YYYYYY'
```

where XXXXXX is the failing batch and YYYYYY is the data id to longer be included in the batch.

6.1.4. Resolving the Issue - Incoming Batches

For batches in error, from the incoming side you'll also have to decide the best course of action to fix the issue. Incoming batch errors *that are in conflict* can by fixed by taking advantage of two columns in INCOMING_ERROR which are examined each time batches are processed. The first column, `resolve_data` if filled in will be used in place of `row_data`. The second column, `resolve_ignore` if set will cause this particular data item to be ignored and batch processing to continue. This is the same two columns used when a manual conflict resolution strategy is chosen, as discussed in Section 4.10, Conflict Detection and Resolution (p. 57) .

6.2. Changing Triggers

A trigger row may be updated using SQL to change a synchronization definition. SymmetricDS will look for changes each night or whenever the Sync Triggers Job is run (see below). For example, a change to place the table `price_changes` into the price channel would be accomplished with the following statement:

```
update SYM_TRIGGER
set channel_id = 'price',
    last_update_by = 'jsmith',
    last_update_time = current_timestamp
where source_table_name = 'price_changes';
```

All configuration should be managed centrally at the registration node. If enabled, configuration changes will be synchronized out to client nodes. When trigger changes reach the client nodes the Sync Triggers Job will run automatically.

Centrally, the trigger changes will not take effect until the Sync Triggers Job runs. Instead of waiting for the Sync Triggers Job to run overnight after making a Trigger change, you can invoke the syncTriggers() method over JMX or simply restart the SymmetricDS server. A complete record of trigger changes is kept in the table TRIGGER_HIST, which was discussed in Section 5.2.3, Sync Triggers Job (p. 65) .

6.3. Maintaining multiple sychronization configurations through Grouplets

As you probably know by now, SymmetricDS stores its single configuration centrally and distributes it to all nodes. By default, a trigger-router is in effect for all nodes in the source node group or target node group. Triggers will be established on each node that is a member of the source node, and changes will be routed to all relevant nodes that are members of the target node group. If, for example, the router routes to "all" nodes, "all" means every node that is in the target node group. This is the default behavior of SymmetricDS.

Once in production, however, you will likely find you need or want to make configuration changes to triggers and routers as new features are rolled out to your network of SymmetricDS nodes. You may, for example, wish to "pilot" a new configuration, containing new synchronizations, only on specific nodes initially, and then increase the size of the pilot over time. SymmetricDS' does provide the ability to specify that only particular trigger-router combinations are applicableto particular nodes for this purpose. It does this by allowing you to define an arbitray collection of nodes, called a "grouplet", and then choosing which trigger-routers apply to the normal set of nodes (the default behavior) and which apply just to nodes in one or more "grouplets". This allows you, essentially, to filter the list of nodes that would otherwise be included as source nodes and/or target nodes. Through the use of grouplets, you can, for example, specify a subset of nodes on which a given trigger would be created. It also allows you to specify a subset of the normal set of nodes a change would be routed to. This behaviour is in addition to, and occurs before, any subsetting or filtering the router might otherwise do.

In its simplest form, a grouplet is just an arbitrary collection of nodes. To define a grouplet, you start by creating a grouplet with a unique id, a description, and a link policy, as defined in GROUPLET. To defined which nodes are members of (or are not members of) a grouplet, you provide a list of external ids of the nodes in GROUPLET_LINK. How those external ids are used varies based on the grouplet link policy. The `grouplet_link_policy` can be either I or E, representing an "inclusive" list of nodes or an "exclusive" list of nodes, respectively. In the case of "inclusive", you'll be listing each external id to be included in the grouplet. In the case of exclusive, all nodes will be included in the grouplet *except* ones which have an external id in the list of external ids.

Once you have defined your grouplet and which nodes are members of a grouplet, you can tie a grouplet to a given trigger-router through the use of TRIGGER_ROUTER_GROUPLET. If a particular trigger-router does not appear in this table, SymmetricDS behaves as normal. If, however, an entry for a particular trigger-router appears in this table, the default behavior is overridden based on the grouplet_id and applies_when settings. The grouplet id provides the node list, and the applies_when indicates whether the grouplet nodes are to be used to filter the source node list, the target node list, or both (settings are "S", "T", and "B", respectively). Nodes that survive the filtering process on as a source will have a trigger defined, and nodes that survive the filtering process as a target are eligible nodes that can be routed to.

6.3.1. Grouplet Example

At this point, an example would probably be useful. Picture the case where you have 100 retail stores (each containing one database, and each a member of the "store" node group) and a central office database (external id of corp, and a member of the "corp" node group). You wish to pilot two new trigger and routers for a new feature on your point-of-sale software (one which moves data from corp to store, and one which moves data from store to corp), but you only want the triggers to be installed on 10 specific stores that represent your "pilot" stores. In this case, the simplest approach would be to define a grouplet with, say, a grouplet id of "pilot". We'd use a grouplet link policy of "inclusive", and list each of the 10 external ids in the GROUPLET_LINK table.

For the trigger-router meant to send data from corp to store, we'd create an entry in TRIGGER_ROUTER_GROUPLET for our grouplet id of "pilot", and we'd specify "T" (target) as the applies-when setting. In this way, the source node list is not filtered, but the target node list used during routing will filter the potential target nodes to just our pilot stores. For the trigger-router meant to send data from a pilot store back to corp, we would have the grouplet apply when the node is in the source node list (i.e., applies_when will be "S"). This will cause the trigger to only be created for stores in the pilot list and not other stores.

An important thing to mention in this example: Since your grouplet only included the store nodes, you can't simply specify "both" for the applies when setting. For the corp-to-store trigger, for example, if you had said "both", no trigger would have been installed in corp since the grouplet nodes represent all possible source nodes as well as target nodes, and "corp" is not in the list! The same is true for the store to corp trigger-router as well. You could, however, use "both" as the applies when if you had included the "corp" external id in with the list of the 10 pilot store external ids.

6.4. Re-synchronizing Data

There may be times where you find you need to re-send or re-synchronize data when the change itself was not captured. This could be needed, for example, if the data changes occurred prior to SymmetricDS placing triggers on the data tables themselves, or if the data at the destination was accidentally deleted, or for some other reason. Two approaches are commonly taken to re-send the data, both of which are discussed below.

> **Important**
>
> Be careful when re-sending data using either of these two techniques. Be sure you are only

sending the rows you intend to send and, more importantly, be sure to re-send the data in a way that won't cause foreign key constraint issues at the destination. In other words, if more than one table is involved, be sure to send any tables which are referred to by other tables by foreign keys first. Otherwise, the channel's synchronization will block because SymmetricDS is unable to insert or update the row because the foreign key relationship refers to a non-existent row in the destination!

One possible approach would be to "touch" the rows in individual tables that need re-sent. By "touch", we mean to alter the row data in such a way that SymmetricDS detects a data change and therefore includes the data change in the batching and synchronizing steps. Note that you have to change the data in some meaningful way (e.g., update a time stamp); setting a column to its current value is not sufficient (by default, if there's not an actual data value change SymmetricDS won't treat the change as something which needs synched.

A second approach would be to take advantage of SymmetricDS built-in functionality by simulating a partial "initial load" of the data. The approach is to manually create "reload" events in DATA for the necessary tables, thereby resending the desired rows for the given tables. Again, foreign key constraints must be kept in mind when creating these reload events. These reload events are created in the source database itself, and the necessary table, trigger-router combination, and channel are included to indicate the direction of synchronization.

To create a reload event, you create a DATA row, using:

- data_id: null

- table_name: name of table to be sent

- event_type: 'R', for reload

- row_data: a "where" clause (minus the word 'where') which defines the subset of rows from the table to be sent. To send all rows, one can use 1=1 for this value.

- pk_data: null

- old_data: null

- trigger_hist_id: use the id of the most recent entry (i.e., max(trigger_hist_id)) in TRIGGER_HIST for the trigger-router combination for your table and router.

- channel_id: the channel in which the table is routed

- transaction_id: pick a value, for example '1'

- source_node_id: null

- external_data: null

- create_time: current_timestamp

By way of example, take our retail hands-on tutorial covered in . Let's say we need to re-send a particular sales transaction from the store to corp over again because we lost the data in corp due to an overzealous delete. For the tutorial, all transaction-related tables start with `sale_`, use the `sale_transaction` channel, and are routed using the `store_corp_identity` router. In addition, the trigger-routers have been set up with an initial load order based on the necessary foreign key relationships (i.e., transaction tables which are "parents" have a lower initial load order than those of their "children"). An insert statement that would create the necessary "reload" events (three in this case, one for each table) would be as follows (where MISSING_ID is changed to the needed transaciton id):

```
insert into sym_data (
    select null, t.source_table_name, 'R', 'tran_id=''MISSING-ID''', null, null,
            h.trigger_hist_id, t.channel_id, '1', null, null, current_timestamp
        from sym_trigger t inner join sym_trigger_router tr on
            t.trigger_id=tr.trigger_id inner join sym_trigger_hist h on
            h.trigger_hist_id=(select max(trigger_hist_id) from sym_trigger_hist
                where trigger_id=t.trigger_id)
    where channel_id='sale_transaction' and
        tr.router_id like 'store_corp_identity' and
        (t.source_table_name like 'sale_%')
    order by tr.initial_load_order asc);
```

This insert statement generates three rows, one for each configured sale table. It uses the most recent trigger history id for the corresponding table. Finally, it takes advantage of the initial load order for each trigger-router to create the three rows in the correct order (the order corresponding to the order in which the tables would have been initial loaded).

6.5. Changing Configuration

The configuration of your system as defined in the `sym_*` tables may be modified at runtime. By default, any changes made to the `sym_*` tables (with the exception of `sym_node`) should be made at the registration server. The changes will be synchronized out to the leaf nodes by SymmetricDS triggers that are automatically created on the tables.

If this behavior is not desired, the feature can be turned off using a parameter. Custom triggers may be added to the `sym_*` tables when the auto syncing feature is disabled.

6.6. Logging Configuration

The standalone SymmetricDS installation uses [Log4J](#) for logging. The configuration file is `conf/log4j.xml`. The `log4j.xml` file has hints as to what logging can be enabled for useful, finer-grained logging.

There is a command line option to turn on preconfigured debugging levels. When the `--debug` option is used the `conf/debug-log4j.xml` is used instead of log4j.xml.

SymmetricDS proxies all of its logging through [SLF4J](#). When deploying to an application server or if

Log4J is not being leveraged, then the general rules for for SLF4J logging apply.

6.7. Java Management Extensions

Monitoring and administrative operations can be performed using Java Management Extensions (JMX). SymmetricDS uses MX4J to expose JMX attributes and operations that can be accessed from the built-in web console, Java's jconsole, or an application server. By default, the web management console can be opened from the following address:

```
http://localhost:31416/
```

In order to use JConsole, you must enable the JVM. You can edit the startup scripts to set the following system parameters.

```
        -Dcom.sun.management.jmxremote.port=31417
        -Dcom.sun.management.jmxremote.authenticate=false
        -Dcom.sun.management.jmxremote.ssl=false
```

More details about enabling JMX for JConsole can be found here.

Using the Java jconsole command, SymmetricDS is listed as a local process named SymmetricLauncher. In jconsole, SymmetricDS appears under the MBeans tab under then name defined by the `engine.name` property. The default value is SymmetricDS.

The management interfaces under SymmetricDS are organized as follows:

- Node - administrative operations

- Parameters - access to properties set through the parameter service

6.8. Temporary Files

SymmetricDS creates temporary extraction and data load files with the CSV payload of a synchronization when the value of the `stream.to.file.threshold.bytes` SymmetricDS property has been reached. Before reaching the threshold, files are streamed to/from memory. The default threshold value is 32,767 bytes. This feature may be turned off by setting the `stream.to.file.enabled` property to false.

SymmetricDS creates these temporary files in the directory specified by the `java.io.tmpdir` Java System property.

The location of the temporary directory may be changed by setting the Java System property passed into the Java program at startup. For example,

```
  -Djava.io.tmpdir=/home/.symmetricds/tmp
```

6.9. Database Purging

Purging is the act of cleaning up captured data that is no longer needed in SymmetricDS's runtime tables. Data is purged through delete statements by the *Purge Job*. Only data that has been successfully synchronized will be purged. Purged tables include:

- DATA

- DATA_EVENT

- OUTGOING_BATCH

- INCOMING_BATCH

- DATA_GAP

- NODE_HOST_STATS

- NODE_HOST_CHANNEL_STATS

- NODE_HOST_JOB_STATS

The purge job is enabled by the `start.purge.job` SymmetricDS property. The timing of the three purge jobs (incoming, outgoing, and data gaps) is controlled by a cron expression as specified by the following properties: `job.purge.outgoing.cron`, `job.purge.incoming.cron`, and `job.purge.datagaps.cron`. The default is `0 0 0 * * *`, or once per day at midnight.

> ⚠️ **Important**
>
> SymmetricDS utilizes Spring's CRON support, which includes seconds as the first parameter. This differs from the typical Unix-based implementation, where the first parameter is usually minutes. For example, `*/15 * * * * *` means every 15 seconds, not every 15 minutes. See Spring's documentation for more details.

Two retention period properties indicate how much history SymmetricDS will retain before purging. The `purge.retention.minutes` property indicates the period of history to keep for synchronization tables. The default value is 5 days. The `statistic.retention.minutes` property indicates the period of history to keep for statistics. The default value is also 5 days.

The purge properties should be adjusted according to how much data is flowing through the system and the amount of storage space the database has. For an initial deployment it is recommended that the purge properties be kept at the defaults, since it is often helpful to be able to look at the captured data in order to triage problems and profile the synchronization patterns. When scaling up to more nodes, it is recomended that the purge parameters be scaled back to 24 hours or less.

Appendix A. Data Model

What follows is the complete SymmetricDS data model. Note that all tables are prepended with a configurable prefix so that multiple instances of SymmetricDS may coexist in the same database. The default prefix is *sym_*.

SymmetricDS configuration is entered by the user into the data model to control the behavior of what data is synchronized to which nodes.
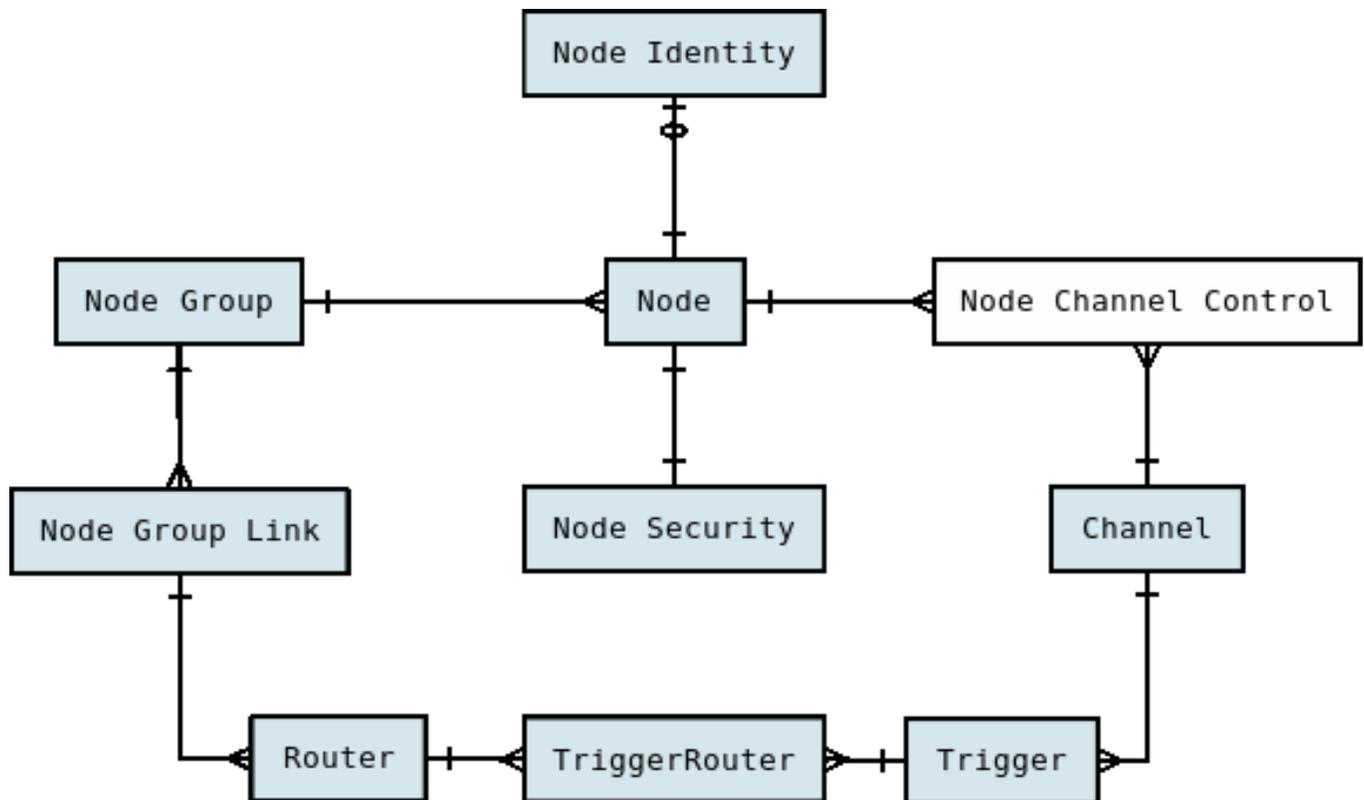
Figure A.1. Configuration Data Model

At runtime, the configuration is used to capture data changes and route them to nodes. The data changes are placed together in a single unit called a batch that can be loaded by another node. Outgoing batches are delivered to nodes and acknowledged. Incoming batches are received and loaded. History is recorded for batch status changes and statistics.
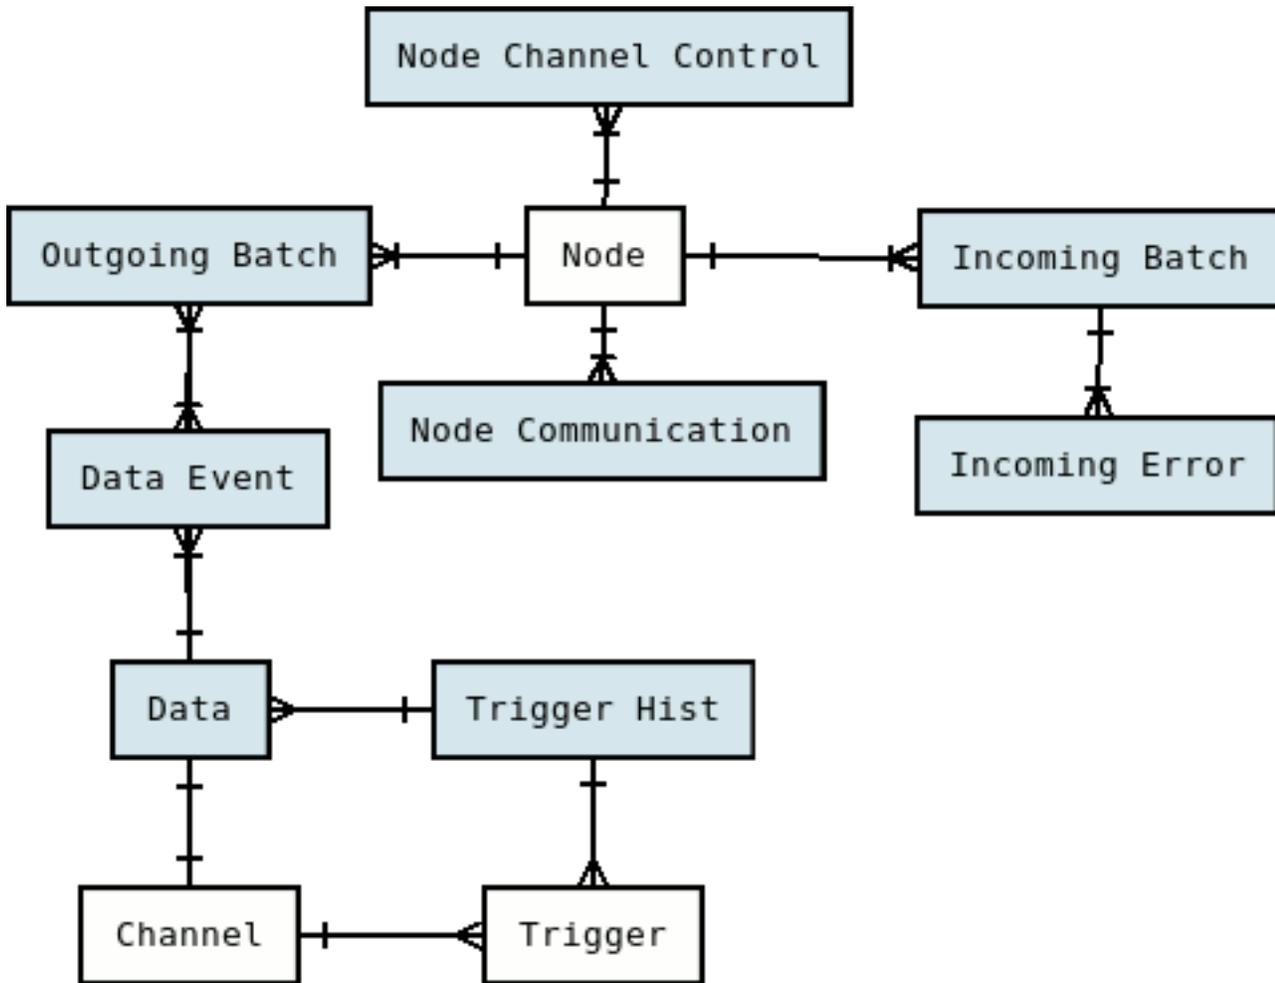
Figure A.2. Runtime Data Model

A.1. NODE

Representation of an instance of SymmetricDS that synchronizes data with one or more additional nodes. Each node has a unique identifier (nodeId) that is used when communicating, as well as a domain-specific identifier (externalId) that provides context within the local system.

Table A.1. NODE

Name	Type / Size	Default	PK FK	not null	Description
NODE_ID	VARCHAR (50)		PK	X	A unique identifier for a node.
NODE_GROUP_ID	VARCHAR (50)			X	The node group that this node belongs to, such as 'store'.
EXTERNAL_ID	VARCHAR (50)			X	A domain-specific identifier for context within the local system. For example, the retail store number.

Name	Type / Size	Default	PK FK	not null	Description
SYNC_ENABLED	INTEGER (1)	0			Indicates whether this node should be sent synchronization. Disabled nodes are ignored by the triggers, so no entries are made in data_event for the node.
SYNC_URL	VARCHAR (255)				The URL to contact the node for synchronization.
SCHEMA_VERSION	VARCHAR (50)				The version of the database schema this node manages. Useful for specifying synchronization by version.
SYMMETRIC_VERSION	VARCHAR (50)				The version of SymmetricDS running at this node.
DATABASE_TYPE	VARCHAR (50)				The database product name at this node as reported by JDBC.
DATABASE_VERSION	VARCHAR (50)				The database product version at this node as reported by JDBC.
HEARTBEAT_TIME	TIMESTAMP				Deprecated. Use node_host.heartbeat_time instead.
TIMEZONE_OFFSET	VARCHAR (6)				Deprecated. Use node_host.timezone_offset instead.
BATCH_TO_SEND_COUNT	INTEGER	0			The number of outgoing batches that have not yet been sent. This field is updated as part of the heartbeat job.
BATCH_IN_ERROR_COUNT	INTEGER	0			The number of outgoing batches that are in error at this node. This field is updated as part of the heartbeat job.
CREATED_AT_NODE_ID	VARCHAR (50)				The node_id of the node where this node was created. This is typically filled automatically with the node_id found in node_identity where registration was opened for the node.
DEPLOYMENT_TYPE	VARCHAR (50)				An indicator as to the type of SymmetricDS software that is running. Possible values are, but not limited to: engine, standalone, war, professional, mobile

A.2. NODE_SECURITY

Security features like node passwords and open registration flag are stored in the node_security table.

Table A.2. NODE_SECURITY

Name	Type / Size	Default	PK FK	not null	Description
node_id	VARCHAR		PK	X	Unique identifier for a node.

Name	Type / Size	Default	PK FK	not null	Description
	(50)		FK		
NODE_PASSWORD	VARCHAR (50)			X	The password used by the node to prove its identity during synchronization.
REGISTRATION_ENABLED	INTEGER (1)	0			Indicates whether registration is open for this node. Re-registration may be forced for a node if this is set back to '1' in a parent database for the node_id that should be re-registred.
REGISTRATION_TIME	TIMESTAMP				The timestamp when this node was last registered.
INITIAL_LOAD_ENABLED	INTEGER (1)	0			Indicates whether an initial load will be sent to this node.
INITIAL_LOAD_TIME	TIMESTAMP				The timestamp when an initial load was started for this node.
REV_INITIAL_LOAD_ENABLED	INTEGER (1)	0			Indicates that this node should send a reverse initial load.
REV_INITIAL_LOAD_TIME	TIMESTAMP				The timestamp when this node last sent an initial load.
CREATED_AT_NODE_ID	VARCHAR (50)			X	The node_id of the node where this node was created. This is typically filled automatically with the node_id found in node_identity where registration was opened for the node.

A.3. NODE_IDENTITY

After registration, this table will have one row representing the identity of the node. For a root node, the row is entered by the user.

Table A.3. NODE_IDENTITY

Name	Type / Size	Default	PK FK	not null	Description
node_id	VARCHAR (50)		PK FK	X	Unique identifier for a node.

A.4. NODE_GROUP

A category of Nodes that synchronizes data with one or more NodeGroups. A common use of NodeGroup is to describe a level in a hierarchy of data synchronization.

Table A.4. NODE_GROUP

Name	Type / Size	Default	PK FK	not null	Description
NODE_GROUP_ID	VARCHAR (50)		PK	X	Unique identifier for a node group, usually named something meaningful, like 'store' or 'warehouse'.
DESCRIPTION	VARCHAR (255)				A description of this node group.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a user last updated this entry.

A.5. NODE_GROUP_LINK

A source node_group sends its data updates to a target NodeGroup using a pull, push, or custom technique.

Table A.5. NODE_GROUP_LINK

Name	Type / Size	Default	PK FK	not null	Description
source_node_group_id	VARCHAR (50)		PK FK	X	The node group where data changes should be captured.
target_node_group_id	VARCHAR (50)		PK FK	X	The node group where data changes will be sent.
DATA_EVENT_ACTION	CHAR (1)	W		X	The notification scheme used to send data changes to the target node group. (P = Push, W = Wait for Pull,)
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a user last updated this entry.

A.6. NODE_HOST

Representation of an physical workstation or server that is hosting the SymmetricDS software. In a clustered environment there may be more than one entry per node in this table.

Table A.6. NODE_HOST

Name	Type / Size	Default	PK FK	not null	Description
node_id	VARCHAR (50)		PK FK	X	A unique identifier for a node.
HOST_NAME	VARCHAR (60)		PK	X	The host name of a workstation or server. If more than one instance of SymmetricDS runs on the same server, then this value can be a 'server id' specified by -Druntime.symmetric.cluster.server.id
IP_ADDRESS	VARCHAR (50)				The ip address for the host.
OS_USER	VARCHAR (50)				The user SymmetricDS is running under
OS_NAME	VARCHAR (50)				The name of the OS
OS_ARCH	VARCHAR (50)				The hardware architecture of the OS
OS_VERSION	VARCHAR (50)				The version of the OS
AVAILABLE_PROCESSORS	INTEGER	0			The number of processors available to use.
FREE_MEMORY_BYTES	BIGINT	0			The amount of free memory available to the JVM.
TOTAL_MEMORY_BYTES	BIGINT	0			The amount of total memory available to the JVM.
MAX_MEMORY_BYTES	BIGINT	0			The max amount of memory available to the JVM.
JAVA_VERSION	VARCHAR (50)				The version of java that SymmetricDS is running as.
JAVA_VENDOR	VARCHAR (255)				The vendor of java that SymmetricDS is running as.
SYMMETRIC_VERSION	VARCHAR (50)				The version of SymmetricDS running at this node.
TIMEZONE_OFFSET	VARCHAR (6)				The timezone offset in RFC822 format at the time of the last heartbeat.
HEARTBEAT_TIME	TIMESTAMP				The last timestamp when the node sent a heartbeat, which is attempted every ten minutes by default.
LAST_RESTART_TIME	TIMESTAMP			X	Timestamp when this instance was last restarted.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.

A.7. NODE_HOST_CHANNEL_STATS

Table A.7. NODE_HOST_CHANNEL_STATS

Name	Type / Size	Default	PK FK	not null	Description
NODE_ID	VARCHAR (50)		PK	X	A unique identifier for a node.
HOST_NAME	VARCHAR (60)		PK	X	The host name of a workstation or server. If more than one instance of SymmetricDS runs on the same server, then this value can be a 'server id' specified by -Druntime.symmetric.cluster.server.id
CHANNEL_ID	VARCHAR (20)		PK	X	The channel_id of the channel that data changes will flow through.
START_TIME	TIMESTAMP		PK	X	The start time for the period which this row represents.
END_TIME	TIMESTAMP		PK	X	The end time for the period which this row represents.
DATA_ROUTED	BIGINT	0			Indicate the number of data rows that have been routed during this period.
DATA_UNROUTED	BIGINT	0			The amount of data that has not yet been routed at the time this stats row was recorded.
DATA_EVENT_INSERTED	BIGINT	0			Indicate the number of data rows that have been routed during this period.
DATA_EXTRACTED	BIGINT	0			The number of data rows that were extracted during this time period.
DATA_BYTES_EXTRACTED	BIGINT	0			The number of bytes that were extracted during this time period.
DATA_EXTRACTED_ERRORS	BIGINT	0			The number of errors that occured during extraction during this time period.
DATA_BYTES_SENT	BIGINT	0			The number of bytes that were sent during this time period.
DATA_SENT	BIGINT	0			The number of rows that were sent during this time period.
DATA_SENT_ERRORS	BIGINT	0			The number of errors that occurred while sending during this time period.
DATA_LOADED	BIGINT	0			The number of rows that were loaded during this time period.
DATA_BYTES_LOADED	BIGINT	0			The number of bytes that were loaded during this time period.
DATA_LOADED_ERRORS	BIGINT	0			The number of errors that occurred while loading during this time period.

A.8. NODE_HOST_STATS

Table A.8. NODE_HOST_STATS

Name	Type / Size	Default	PK FK	not null	Description
NODE_ID	VARCHAR (50)		PK	X	A unique identifier for a node.
HOST_NAME	VARCHAR (60)		PK	X	The host name of a workstation or server. If more than one instance of SymmetricDS runs on the same server, then this value can be a 'server id' specified by -Druntime.symmetric.cluster.server.id
START_TIME	TIMESTAMP		PK	X	The end time for the period which this row represents.
END_TIME	TIMESTAMP		PK	X	
RESTARTED	BIGINT	0		X	Indicate that a restart occurred during this period.
NODES_PULLED	BIGINT	0			
TOTAL_NODES_PULL_TIME	BIGINT	0			
NODES_PUSHED	BIGINT	0			
TOTAL_NODES_PUSH_TIME	BIGINT	0			
NODES_REJECTED	BIGINT	0			
NODES_REGISTERED	BIGINT	0			
NODES_LOADED	BIGINT	0			
NODES_DISABLED	BIGINT	0			
PURGED_DATA_ROWS	BIGINT	0			
PURGED_DATA_EVENT_ROWS	BIGINT	0			
PURGED_BATCH_OUTGOING_ROWS	BIGINT	0			
PURGED_BATCH_INCOMING_ROWS	BIGINT	0			
TRIGGERS_CREATED_COUNT	BIGINT				
TRIGGERS_REBUILT_COUNT	BIGINT				
TRIGGERS_REMOVED_COUNT	BIGINT				

A.9. NODE_HOST_JOB_STATS

Table A.9. NODE_HOST_JOB_STATS

Name	Type / Size	Default	PK FK	not null	Description
NODE_ID	VARCHAR (50)		PK	X	A unique identifier for a node.
HOST_NAME	VARCHAR (60)		PK	X	The host name of a workstation or server. If more than one instance of SymmetricDS runs on the same server, then this value can be a 'server id' specified by -Druntime.symmetric.cluster.server.id
JOB_NAME	VARCHAR (50)		PK	X	The name of the job.
START_TIME	TIMESTAMP		PK	X	The start time for the period which this row represents.
END_TIME	TIMESTAMP		PK	X	The end time for the period which this row represents.
PROCESSED_COUNT	BIGINT	0			The number of items that were processed during the job run.

A.10. CHANNEL

This table represents a category of data that can be synchronized independently of other channels. Channels allow control over the type of data flowing and prevents one type of synchronization from contending with another.

Table A.10. CHANNEL

Name	Type / Size	Default	PK FK	not null	Description
CHANNEL_ID	VARCHAR (20)		PK	X	A unique identifer, usually named something meaningful, like 'sales' or 'inventory'.
PROCESSING_ORDER	INTEGER	1		X	Order of sequence to process channel data.
MAX_BATCH_SIZE	INTEGER	1000		X	The maximum number of Data Events to process within a batch for this channel.
MAX_BATCH_TO_SEND	INTEGER	60		X	The maximum number of batches to send during a 'synchronization' between two nodes. A 'synchronization' is equivalent to a push or a pull. If there are 12 batches ready to be sent for a channel and max_batch_to_send is equal to 10, then only the first 10 batches will be sent.
MAX_DATA_TO_ROUTE	INTEGER	100000		X	The maximum number of data rows to route for a channel at a time.
EXTRACT_PERIOD_MILLIS	INTEGER	0		X	The minimum number of milliseconds allowed

Name	Type / Size	Default	PK FK	not null	Description
					between attempts to extract data for targeted at a node_id.
ENABLED	INTEGER (1)	1		X	Indicates whether channel is enabled or not.
USE_OLD_DATA_TO_ROUTE	INTEGER (1)	1		X	Indicates whether to read the old data during routing.
USE_ROW_DATA_TO_ROUTE	INTEGER (1)	1		X	Indicates whether to read the row data during routing.
USE_PK_DATA_TO_ROUTE	INTEGER (1)	1		X	Indicates whether to read the pk data during routing.
CONTAINS_BIG_LOB	INTEGER (1)	0		X	Provides SymmetricDS a hint on how to treat captured data. Currently only supported by Oracle. If set to '0', then selects for routing and data extraction will be more efficient and lobs will be truncated at 4k in the trigger text. When it is set to '0' there is a 4k limit on the total size of a row and on the size of a LOB column. Note, when switching this value back and forth triggers need to be forced to regenerate.
BATCH_ALGORITHM	VARCHAR (50)	default		X	The algorithm to use when batching data on this channel. Possible values are: 'default', 'transactional', and 'nontransactional'
DATA_LOADER_TYPE	VARCHAR (50)	default		X	Identify the type of data loader this channel should use. Allows for the default dataloader to be swapped out via configuration for more efficient platform specific data loaders.
DESCRIPTION	VARCHAR (255)				Description on the type of data carried in this channel.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a user last updated this entry.

A.11. NODE_COMMUNICATION

This table is used to coordinate communication with other nodes.

Table A.11. NODE_COMMUNICATION

Name	Type / Size	Default	PK FK	not null	Description
NODE_ID	VARCHAR (50)		PK	X	Unique identifier for a node.

Name	Type / Size	Default	PK FK	not null	Description
COMMUNICATION_TYPE	VARCHAR (10)		PK	X	The type of communication that is taking place with this node. Valid values are: PULL, PUSH
LOCK_TIME	TIMESTAMP				The timestamp when this node was locked
LOCKING_SERVER_ID	VARCHAR (255)				The name of the server that currently has a pull lock for the node. This is typically a host name, but it can be overridden using the -Druntime.symmetric.cluster.server.id=name System property.
LAST_LOCK_TIME	TIMESTAMP				The timestamp when this node was last locked
LAST_LOCK_MILLIS	BIGINT	0			The amount of time the last communication took.
SUCCESS_COUNT	BIGINT	0			The number of successive successful communication attempts.
FAIL_COUNT	BIGINT	0			The number of successive failed communication attempts.
TOTAL_SUCCESS_COUNT	BIGINT	0			The total number of successful communication attempts with the node.
TOTAL_FAIL_COUNT	BIGINT	0			The total number of failed communication attempts with the node.
TOTAL_SUCCESS_MILLIS	BIGINT	0			The total amount of time spent during successful communication attempts with the node.
TOTAL_FAIL_MILLIS	BIGINT	0			The total amount of time spent during failed communication attempts with the node.

A.12. NODE_CHANNEL_CTL

Used to ignore or suspend a channel. A channel that is ignored will have its data_events batched and they will immediately be marked as 'OK' without sending them. A channel that is suspended is skipped when batching data_events.

Table A.12. NODE_CHANNEL_CTL

Name	Type / Size	Default	PK FK	not null	Description
NODE_ID	VARCHAR (50)		PK	X	Unique identifier for a node.
CHANNEL_ID	VARCHAR (20)		PK	X	The name of the channel_id that is being controlled.
SUSPEND_ENABLED	INTEGER (1)	0			Indicates if this channel is suspended, which prevents its Data Events from being batched.

Name	Type / Size	Default	PK FK	not null	Description
IGNORE_ENABLED	INTEGER (1)	0			Indicates if this channel is ignored, which marks its Data Events as if they were actually processed.
LAST_EXTRACT_TIME	TIMESTAMP				Record the last time data was extract for a node and a channel.

A.13. NODE_GROUP_CHANNEL_WINDOW

An optional window of time for which a node group and channel will extract and send data.

Table A.13. NODE_GROUP_CHANNEL_WINDOW

Name	Type / Size	Default	PK FK	not null	Description
NODE_GROUP_ID	VARCHAR (50)		PK	X	The node_group_id that this window applies to.
CHANNEL_ID	VARCHAR (20)		PK	X	The channel_id that this window applies to.
START_TIME	TIME		PK	X	The start time for the active window.
END_TIME	TIME		PK	X	The end time for the active window. Note that if the end_time is less than the start_time then the window crosses a day boundary.
ENABLED	INTEGER (1)	0		X	Enable this window. If this is set to '0' then this window is ignored.

A.14. TRIGGER

Configures database triggers that capture changes in the database. Configuration of which triggers are generated for which tables is stored here. Triggers are created in a node's database if the source_node_group_id of a router is mapped to a row in this table.

Table A.14. TRIGGER

Name	Type / Size	Default	PK FK	not null	Description
TRIGGER_ID	VARCHAR (50)		PK	X	Unique identifier for a trigger.
SOURCE_CATALOG_NAME	VARCHAR (255)				Optional name for the catalog the configured table is in.
SOURCE_SCHEMA_NAME	VARCHAR				Optional name for the schema a configured

Name	Type / Size	Default	PK FK	not null	Description
	(255)				table is in.
SOURCE_TABLE_NAME	VARCHAR (255)			X	The name of the source table that will have a trigger installed to watch for data changes.
channel_id	VARCHAR (20)		FK	X	The channel_id of the channel that data changes will flow through.
SYNC_ON_UPDATE	INTEGER (1)	1		X	Whether or not to install an update trigger.
SYNC_ON_INSERT	INTEGER (1)	1		X	Whether or not to install an insert trigger.
SYNC_ON_DELETE	INTEGER (1)	1		X	Whether or not to install an delete trigger.
SYNC_ON_INCOMING_BATCH	INTEGER (1)	0		X	Whether or not an incoming batch that loads data into this table should cause the triggers to capture data_events. Be careful turning this on, because an update loop is possible.
NAME_FOR_UPDATE_TRIGGER	VARCHAR (255)				Override the default generated name for the update trigger.
NAME_FOR_INSERT_TRIGGER	VARCHAR (255)				Override the default generated name for the insert trigger.
NAME_FOR_DELETE_TRIGGER	VARCHAR (255)				Override the default generated name for the delete trigger.
SYNC_ON_UPDATE_CONDITION	LONGVARCHAR				Specify a condition for the update trigger firing using an expression specific to the database.
SYNC_ON_INSERT_CONDITION	LONGVARCHAR				Specify a condition for the insert trigger firing using an expression specific to the database.
SYNC_ON_DELETE_CONDITION	LONGVARCHAR				Specify a condition for the delete trigger firing using an expression specific to the database.
EXTERNAL_SELECT	LONGVARCHAR				Specify a SQL select statement that returns a single result. It will be used in the generated database trigger to populate the EXTERNAL_DATA field on the data table.
TX_ID_EXPRESSION	LONGVARCHAR				Override the default expression for the transaction identifier that groups the data changes that were committed together.
EXCLUDED_COLUMN_NAMES	LONGVARCHAR				Specify a comma-delimited list of columns that should not be synchronized from this table. Note that if a primary key is found in this list, it will be ignored.
SYNC_KEY_NAMES	LONGVARCHAR				Specify a comma-delimited list of columns that should be used as the key for synchronization operations. By default, if not specified, then the primary key of the table will be used.
USE_STREAM_LOBS	INTEGER (1)	0		X	Specifies whether to capture lob data as the trigger is firing or to stream lob columns from the source tables using callbacks during extraction. A value of 1 indicates to stream from the source via callback; a value of 0, lob

Name	Type / Size	Default	PK FK	not null	Description
					data is captured by the trigger.
USE_CAPTURE_LOBS	INTEGER (1)	0		X	Provides a hint as to whether this trigger will capture big lobs data. If set to 1 every effort will be made during data capture in trigger and during data selection for initial load to use lob facilities to extract and store data in the database. On Oracle, this may need to be set to 1 to get around 4k concatenation errors during data capture and during initial load.
USE_CAPTURE_OLD_DATA	INTEGER (1)	1		X	Indicates whether this trigger should capture and send the old data (previous state of the row before the change).
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.15. ROUTER

Configure a type of router from one node group to another. Note that routers are mapped to triggers through trigger_routers.

Table A.15. ROUTER

Name	Type / Size	Default	PK FK	not null	Description
ROUTER_ID	VARCHAR (50)		PK	X	Unique description of a specific router
TARGET_CATALOG_NAME	VARCHAR (255)				Optional name for the catalog a target table is in. Only use this if the target table is not in the default catalog.
TARGET_SCHEMA_NAME	VARCHAR (255)				Optional name of the schema a target table is in. On use this if the target table is not in the default schema.
TARGET_TABLE_NAME	VARCHAR (255)				Optional name for a target table. Only use this if the target table name is different than the source.
source_node_group_id	VARCHAR (50)		FK	X	Routers with this node_group_id will install triggers that are mapped to this router.
target_node_group_id	VARCHAR (50)		FK	X	The node_group_id for nodes to route data to. Note that routing can be further narrowed down by the configured router_type and router_expression.

Name	Type / Size	Default	PK FK	not null	Description
ROUTER_TYPE	VARCHAR (50)				The name of a specific type of router. Out of the box routers are 'default','column','bsh', 'subselect' and 'audit.' Custom routers can be configured as extension points.
ROUTER_EXPRESSION	LONGVARCHAR				An expression that is specific to the type of router that is configured in router_type. See the documentation for each router for more details.
SYNC_ON_UPDATE	INTEGER (1)	1		X	Flag that indicates that this router should route updates.
SYNC_ON_INSERT	INTEGER (1)	1		X	Flag that indicates that this router should route inserts.
SYNC_ON_DELETE	INTEGER (1)	1		X	Flag that indicates that this router should route deletes.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.16. TRIGGER_ROUTER

Map a trigger to a router.

Table A.16. TRIGGER_ROUTER

Name	Type / Size	Default	PK FK	not null	Description
trigger_id	VARCHAR (50)		PK FK	X	The id of a trigger.
router_id	VARCHAR (50)		PK FK	X	The id of a router.
ENABLED	INTEGER (1)	1		X	Indicates whether this trigger router is enabled or not.
INITIAL_LOAD_ORDER	INTEGER	1		X	Order sequence of this table when an initial load is sent to a node. If this value is the same for multiple tables, then SymmetricDS will attempt to order the tables according to FK constraints. If this value is set to a negative number, then the table will be excluded from an initial load.
INITIAL_LOAD_SELECT	LONGVARCHAR				Optional expression that can be used to pair down the data selected from a table during the initial load process.

Name	Type / Size	Default	PK FK	not null	Description
INITIAL_LOAD_DELETE_STMT	LONGVARCHAR				The expression that is used to delete data when an initial load occurs. If this field is empty, no delete will occur before the initial load. If this field is not empty, the text will be used as a sql statement and executed for the initial load delete.
PING_BACK_ENABLED	INTEGER (1)	0		X	When enabled, the node will route data that originated from a node back to that node. This attribute is only effective if sync_on_incoming_batch is set to 1.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.17. PARAMETER

Provides a way to manage most SymmetricDS settings in the database.

Table A.17. PARAMETER

Name	Type / Size	Default	PK FK	not null	Description
EXTERNAL_ID	VARCHAR (50)		PK	X	Target the parameter at a specific external id. To target all nodes, use the value of 'ALL.'
NODE_GROUP_ID	VARCHAR (50)		PK	X	Target the parameter at a specific node group id. To target all groups, use the value of 'ALL.'
PARAM_KEY	VARCHAR (80)		PK	X	The name of the parameter.
PARAM_VALUE	LONGVARCHAR				The value of the parameter.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a user last updated this entry.

A.18. REGISTRATION_REDIRECT

Provides a way for a centralized registration server to redirect registering nodes to their prospective parent node in a multi-tiered deployment.

Table A.18. REGISTRATION_REDIRECT

Name	Type / Size	Default	PK FK	not null	Description
REGISTRANT_EXTERNAL_ID	VARCHAR (50)		PK	X	Maps the external id of a registration request to a different parent node.
REGISTRATION_NODE_ID	VARCHAR (50)			X	The node_id of the node that a registration request should be redirected to.

A.19. REGISTRATION_REQUEST

Audits when a node registers or attempts to register.

Table A.19. REGISTRATION_REQUEST

Name	Type / Size	Default	PK FK	not null	Description
NODE_GROUP_ID	VARCHAR (50)		PK	X	The node group that this node belongs to, such as 'store'.
EXTERNAL_ID	VARCHAR (50)		PK	X	A domain-specific identifier for context within the local system. For example, the retail store number.
STATUS	CHAR (2)			X	The current status of the registration attempt. Valid statuses are NR (not registered), IG (ignored), OK (sucessful)
HOST_NAME	VARCHAR (60)			X	The host name of a workstation or server. If more than one instance of SymmetricDS runs on the same server, then this value can be a 'server id' specified by -Druntime.symmetric.cluster.server.id
IP_ADDRESS	VARCHAR (50)			X	The ip address for the host.
ATTEMPT_COUNT	INTEGER	0			The number of registration attempts.
REGISTERED_NODE_ID	VARCHAR (50)				A unique identifier for a node.
CREATE_TIME	TIMESTAMP		PK	X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.20. TRIGGER_HIST

A history of a table's definition and the trigger used to capture data from the table. When a database trigger captures a data change, it references a trigger_hist entry so it is possible to know which columns the data represents. trigger_hist entries are made during the sync trigger process, which runs at each startup, each night in the syncTriggersJob, or any time the syncTriggers() JMX method is manually invoked. A new entry is made when a table definition or a trigger definition is changed, which causes a database trigger to be created or rebuilt.

Table A.20. TRIGGER_HIST

Name	Type / Size	Default	PK FK	not null	Description
TRIGGER_HIST_ID	INTEGER		PK	X	Unique identifier for a trigger_hist entry
TRIGGER_ID	VARCHAR (50)			X	Unique identifier for a trigger
SOURCE_TABLE_NAME	VARCHAR (255)			X	The name of the source table that will have a trigger installed to watch for data changes.
SOURCE_CATALOG_NAME	VARCHAR (255)				The catalog name where the source table resides.
SOURCE_SCHEMA_NAME	VARCHAR (255)				The schema name where the source table resides.
NAME_FOR_UPDATE_TRIGGER	VARCHAR (255)			X	The name used when the insert trigger was created.
NAME_FOR_INSERT_TRIGGER	VARCHAR (255)			X	The name used when the update trigger was created.
NAME_FOR_DELETE_TRIGGER	VARCHAR (255)			X	The name used when the delete trigger was created.
TABLE_HASH	BIGINT			X	A hash of the table definition, used to detect changes in the definition.
TRIGGER_ROW_HASH	BIGINT			X	A hash of the trigger definition. If changes are detected to the values that affect a trigger definition, then the trigger will be regenerated.
COLUMN_NAMES	LONGVARCHAR			X	The column names defined on the table. The column names are stored in comma-separated values (CSV) format.
PK_COLUMN_NAMES	LONGVARCHAR			X	The primary key column names defined on the table. The column names are stored in comma-separated values (CSV) format.
LAST_TRIGGER_BUILD_REASON	CHAR (1)			X	The following reasons for a change are possible: New trigger that has not been created before (N); Schema changes in the table were detected (S); Configuration changes in Trigger (C); Trigger was missing (T).
ERROR_MESSAGE	LONGVARCHAR				Record any errors or warnings that occurred when attempting to build the trigger.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.

Name	Type / Size	Default	PK FK	not null	Description
INACTIVE_TIME	TIMESTAMP				The date and time when a trigger was inactivated.

A.21. DATA

The captured data change that occurred to a row in the database. Entries in data are created by database triggers.

Table A.21. DATA

Name	Type / Size	Default	PK FK	not null	Description
DATA_ID	BIGINT		PK	X	Unique identifier for a data.
TABLE_NAME	VARCHAR (255)			X	The name of the table in which a change occurred that this entry records.
EVENT_TYPE	CHAR (1)			X	The type of event captured by this entry. For triggers, this is the change that occurred, which is 'I' for insert, 'U' for update, or 'D' for delete. Other events include: 'R' for reloading the entire table (or subset of the table) to the node; 'S' for running dynamic SQL at the node, which is used for adhoc administration.
ROW_DATA	LONGVARCHAR				The captured data change from the synchronized table. The column values are stored in comma-separated values (CSV) format.
PK_DATA	LONGVARCHAR				The primary key values of the captured data change from the synchronized table. This data is captured for updates and deletes. The primary key values are stored in comma-separated values (CSV) format.
OLD_DATA	LONGVARCHAR				The captured data values prior to the update. The column values are stored in CSV format.
TRIGGER_HIST_ID	INTEGER			X	The foreign key to the trigger_hist entry that contains the primary key and column names for the table being synchronized.
CHANNEL_ID	VARCHAR (20)				The channel that this data belongs to, such as 'prices'
TRANSACTION_ID	VARCHAR (255)				An optional transaction identifier that links multiple data changes together as the same transaction.
SOURCE_NODE_ID	VARCHAR (50)				If the data was inserted by a SymmetricDS data loader, then the id of the source node is record so that data is not re-routed back to it.

Name	Type / Size	Default	PK FK	not null	Description
EXTERNAL_DATA	VARCHAR (50)				A field that can be populated by a trigger that uses the EXTERNAL_SELECT
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.

A.22. DATA_GAP

Used only when routing.data.reader.type is set to 'gap.' Table that tracks gaps in the data table so that they may be processed efficiently, if data shows up. Gaps can show up in the data table if a database transaction is rolled back.

Table A.22. DATA_GAP

Name	Type / Size	Default	PK FK	not null	Description
START_ID	BIGINT		PK	X	The first missing data_id from the data table where a gap is detected. This could be the last data_id inserted plus one.
END_ID	BIGINT		PK	X	The last missing data_id from the data table where a gap is detected. If the start_id is the last data_id inserted plus one, then this field is filled in with a -1.
STATUS	CHAR (2)				GP, SK, or FL. GP means there is a detected gap. FL means that the gap has been filled. SK means that the gap has been skipped either because the gap expired or because no database transaction was detected which means that no data will be committed to fill in the gap.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_HOSTNAME	VARCHAR (255)				The host who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.23. DATA_EVENT

Represents routing of a data row to one or more nodes. Entries in data_event are created by database triggers.

Table A.23. DATA_EVENT

Name	Type / Size	Default	PK FK	not null	Description
DATA_ID	BIGINT		PK	X	Id of the data to be routed.
BATCH_ID	BIGINT		PK	X	The node_id of the node that is to receive the data.
ROUTER_ID	VARCHAR (50)		PK	X	The router_id of the router that routed this data_event.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.

A.24. OUTGOING_BATCH

Used for tracking the sending a collection of data to a node in the system. A new outgoing_batch is created and given a status of 'NE'. After sending the outgoing_batch to its target node, the status becomes 'SE'. The node responds with either a success status of 'OK' or an error status of 'ER'. An error while sending to the node also results in an error status of 'ER' regardless of whether the node sends that acknowledgement.

Table A.24. OUTGOING_BATCH

Name	Type / Size	Default	PK FK	not null	Description
BATCH_ID	BIGINT		PK	X	A unique id for the batch.
NODE_ID	VARCHAR (50)		PK	X	The node that this batch is targeted at.
CHANNEL_ID	VARCHAR (20)				The channel that this batch is part of.
STATUS	CHAR (2)				The current status of a batch can be routing (RT), newly created and ready for replication (NE), being queried from the database (QY), sent to a Node (SE), ready to be loaded (LD) and acknowledged as successful (OK), ignored (IG) or in error (ER).
LOAD_FLAG	INTEGER (1)	0			A flag that indicates that this batch is part of an initial load.
ERROR_FLAG	INTEGER (1)	0			A flag that indicates that this batch was in error during the last synchornization attempt.
COMMON_FLAG	INTEGER (1)	0			A flag that indicates that the data in this batch is shared by other nodes (they will have the same batch_id). Shared batches will be extracted to a common location.
IGNORE_COUNT	BIGINT	0		X	The number of times a batch was ignored.
BYTE_COUNT	BIGINT	0		X	The number of bytes that were sent as part of this batch.
EXTRACT_COUNT	BIGINT	0		X	The number of times this an attempt to extract

Name	Type / Size	Default	PK FK	not null	Description
					this batch occurred.
SENT_COUNT	BIGINT	0		X	The number of times this batch was sent. A batch can be sent multiple times if an ACK is not received.
LOAD_COUNT	BIGINT	0		X	The number of times an attempt to load this batch occurred.
DATA_EVENT_COUNT	BIGINT	0		X	The number of data_events that are part of this batch.
RELOAD_EVENT_COUNT	BIGINT	0		X	The number of reload events that are part of this batch.
INSERT_EVENT_COUNT	BIGINT	0		X	The number of insert events that are part of this batch.
UPDATE_EVENT_COUNT	BIGINT	0		X	The number of update events that are part of this batch.
DELETE_EVENT_COUNT	BIGINT	0		X	The number of delete events that are part of this batch.
OTHER_EVENT_COUNT	BIGINT	0		X	The number of other event types that are part of this batch. This includes any events types that are not a reload, insert, update or delete event type.
ROUTER_MILLIS	BIGINT	0		X	The number of milliseconds spent creating this batch.
NETWORK_MILLIS	BIGINT	0		X	The number of milliseconds spent transfering this batch across the network.
FILTER_MILLIS	BIGINT	0		X	The number of milliseconds spent in filters processing data.
LOAD_MILLIS	BIGINT	0		X	The number of milliseconds spent loading the data into the target database.
EXTRACT_MILLIS	BIGINT	0		X	The number of milliseconds spent extracting the data out of the source database.
SQL_STATE	VARCHAR (10)				For a status of error (ER), this is the XOPEN or SQL 99 SQL State.
SQL_CODE	INTEGER	0		X	For a status of error (ER), this is the error code from the database that is specific to the vendor.
SQL_MESSAGE	LONGVARCHAR				For a status of error (ER), this is the error message that describes the error.
FAILED_DATA_ID	BIGINT	0		X	For a status of error (ER), this is the data_id that was being processed when the batch failed.
FAILED_LINE_NUMBER	BIGINT	0		X	The current line number in the CSV for this batch that failed.
LAST_UPDATE_HOSTNAME	VARCHAR (255)				The host name of the process that last did work on this batch.

Name	Type / Size	Default	PK FK	not null	Description
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a process last updated this entry.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.

A.25. INCOMING_BATCH

The incoming_batch is used for tracking the status of loading an outgoing_batch from another node. Data is loaded and commited at the batch level. The status of the incoming_batch is either successful (OK) or error (ER).

Table A.25. INCOMING_BATCH

Name	Type / Size	Default	PK FK	not null	Description
BATCH_ID	BIGINT (50)		PK	X	The id of the outgoing_batch that is being loaded.
NODE_ID	VARCHAR (50)		PK	X	The node_id of the source of the batch being loaded.
CHANNEL_ID	VARCHAR (20)				The channel_id of the batch being loaded.
STATUS	CHAR (2)				The current status of the batch can be loading (LD), successfully loaded (OK), in error (ER) or skipped (SK)
ERROR_FLAG	INTEGER (1)	0			A flag that indicates that this batch was in error during the last synchornization attempt.
NETWORK_MILLIS	BIGINT	0		X	The number of milliseconds spent transfering this batch across the network.
FILTER_MILLIS	BIGINT	0		X	The number of milliseconds spent in filters processing data.
DATABASE_MILLIS	BIGINT	0		X	The number of milliseconds spent loading the data into the target database.
FAILED_ROW_NUMBER	BIGINT	0		X	This numbered data event that failed as read from the CSV.
FAILED_LINE_NUMBER	BIGINT	0		X	The current line number in the CSV for this batch that failed.
BYTE_COUNT	BIGINT	0		X	The number of bytes that were sent as part of this batch.
STATEMENT_COUNT	BIGINT	0		X	The number of statements run to load this batch.
FALLBACK_INSERT_COUNT	BIGINT	0		X	The number of times an update was turned into an insert because the data was not already in the

Name	Type / Size	Default	PK FK	not null	Description
					target database.
FALLBACK_UPDATE_COUNT	BIGINT	0		X	The number of times an insert was turned into an update because a data row already existed in the target database.
IGNORE_COUNT	BIGINT	0		X	The number of times a row was ignored.
MISSING_DELETE_COUNT	BIGINT	0		X	THe number of times a delete did not effect the database because the row was already deleted.
SKIP_COUNT	BIGINT	0		X	The number of times a batch was sent and skipped because it had already been loaded according to incoming_batch.
SQL_STATE	VARCHAR (10)				For a status of error (ER), this is the XOPEN or SQL 99 SQL State.
SQL_CODE	INTEGER	0		X	For a status of error (ER), this is the error code from the database that is specific to the vendor.
SQL_MESSAGE	LONGVARCHAR				For a status of error (ER), this is the error message that describes the error.
LAST_UPDATE_HOSTNAME	VARCHAR (255)				The host name of the process that last did work on this batch.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a process last updated this entry.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.

A.26. LOCK

Contains semaphores that are set when processes run, so that only one server can run a process at a time. Enable this feature by using the cluster.lock.during.xxxx parameters.

Table A.26. LOCK

Name	Type / Size	Default	PK FK	not null	Description
LOCK_ACTION	VARCHAR (50)		PK	X	The process that needs a lock.
LOCKING_SERVER_ID	VARCHAR (255)				The name of the server that currently has a lock. This is typically a host name, but it can be overridden using the -Druntime.symmetric.cluster.server.id=name System property.
LOCK_TIME	TIMESTAMP				The time a lock is aquired. Use the cluster.lock.timeout.ms to specify a lock timeout period.
LAST_LOCK_TIME	TIMESTAMP				Timestamp when a process last updated this

Name	Type / Size	Default	PK FK	not null	Description
					entry.
LAST_LOCKING_SERVER_ID	VARCHAR (255)				The server id of the process that last did work on this batch.

A.27. TRANSFORM_TABLE

Defines a data loader transformation which can be used to map arbitrary tables and columns to other tables and columns.

Table A.27. TRANSFORM_TABLE

Name	Type / Size	Default	PK FK	not null	Description
TRANSFORM_ID	VARCHAR (50)		PK	X	Unique identifier of a specific transform.
source_node_group_id	VARCHAR (50)		PK FK	X	The node group where data changes are captured.
target_node_group_id	VARCHAR (50)		PK FK	X	The node group where data changes will be sent.
TRANSFORM_POINT	VARCHAR (10)			X	The point during the transport of captured data that a transform happens. Support values are EXTRACT or LOAD.
SOURCE_CATALOG_NAME	VARCHAR (255)				Optional name for the catalog the configured table is in.
SOURCE_SCHEMA_NAME	VARCHAR (255)				Optional name for the schema a configured table is in.
SOURCE_TABLE_NAME	VARCHAR (255)			X	The name of the source table that will be transformed.
TARGET_CATALOG_NAME	VARCHAR (255)				Optional name for the catalog a target table is in. Only use this if the target table is not in the default catalog.
TARGET_SCHEMA_NAME	VARCHAR (255)				Optional name of the schema a target table is in. Only use this if the target table is not in the default schema.
TARGET_TABLE_NAME	VARCHAR (255)				Optional name for a target table. Use this if the target table name is different than the source.
UPDATE_FIRST	INTEGER (1)	0			If true, the target actions are attempted as updates first, regardless of whether the source operation was an insert or an update.
DELETE_ACTION	VARCHAR (10)			X	An action to take upon delete of a row. Possible values are: DEL_ROW, UPDATE_COL, or NONE.

Name	Type / Size	Default	PK FK	not null	Description
TRANSFORM_ORDER	INTEGER	1		X	Specifies the order in which to apply transforms if more than one target operation occurs.
COLUMN_POLICY	VARCHAR (10)	SPECIFIED		X	Specifies whether all columns need to be specified or whether they are implied. Possible values are SPECIFIED or IMPLIED.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a user last updated this entry.

A.28. TRANSFORM_COLUMN

Defines the column mappings and optional data transformation for a data loader transformation.

Table A.28. TRANSFORM_COLUMN

Name	Type / Size	Default	PK FK	not null	Description
TRANSFORM_ID	VARCHAR (50)		PK	X	Unique identifier of a specific transform.
INCLUDE_ON	CHAR (1)	*	PK	X	Indicates whether this mapping is included during an insert (I), update (U), delete (D) operation at the target based on the dml type at the source. A value of * represents the fact that you want to map the column for all operations.
TARGET_COLUMN_NAME	VARCHAR (128)		PK	X	Name of the target column.
SOURCE_COLUMN_NAME	VARCHAR (128)				Name of the source column.
PK	INTEGER (1)	0			Indicates whether this mapping defines a primary key to be used to identify the target row. At least one row must be defined as a pk for each transform_id.
TRANSFORM_TYPE	VARCHAR (50)	copy			The name of a specific type of transform. Custom transformers can be configured as extension points.
TRANSFORM_EXPRESSION	LONGVARCHAR				An expression that is specific to the type of transform that is configured in transform_type. See the documentation for each transformer for more details.
TRANSFORM_ORDER	INTEGER	1		X	Specifies the order in which to apply transforms if more than one target operation occurs.

Name	Type / Size	Default	PK FK	not null	Description
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP				Timestamp when a user last updated this entry.

A.29. CONFLICT

Defines how conflicts in row data should be handled during the load process.

Table A.29. CONFLICT

Name	Type / Size	Default	PK FK	not null	Description
CONFLICT_ID	VARCHAR (50)		PK	X	Unique identifier for a specific conflict detection setting.
source_node_group_id	VARCHAR (50)		FK	X	The source node group for which this setting will be applied to. References a node group link.
target_node_group_id	VARCHAR (50)		FK	X	The target node group for which this setting will be applied to. References a node group link.
TARGET_CHANNEL_ID	VARCHAR (20)				Optional channel that this setting will be applied to.
TARGET_CATALOG_NAME	VARCHAR (255)				Optional database catalog that the target table belongs to. Only use this if the target table is not in the default catalog.
TARGET_SCHEMA_NAME	VARCHAR (255)				Optional database schema that the target table belongs to. Only use this if the target table is not in the default schema.
TARGET_TABLE_NAME	VARCHAR (255)				Optional database table that this setting will apply to. If left blank, the setting will be for any table in the channel (if set) and in the specified node group link.
DETECT_TYPE	VARCHAR (128)			X	Indicates the strategy to use for detecting conflicts during a dml action. The possible values are: use_pk_data (manual, fallback, ignore), use_changed_data (manual, fallback, ignore), use_old_data (manual, fallback, ignore), use_timestamp (newer_wins), use_version (newer_wins)
DETECT_EXPRESSION	LONGVARCHAR				An expression that provides additional information about the detection mechanism. If the detection mechanism is use_timestamp or use_version then this expression will be the

Name	Type / Size	Default	PK FK	not null	Description
					name of the timestamp or version column.
RESOLVE_TYPE	VARCHAR (128)			X	Indicates the strategy for resolving update conflicts. The possible values differ based on the detect_type that is specified.
PING_BACK	VARCHAR (128)			X	Indicates the strategy for sending resolved conflicts back to the source system. Possible values are: OFF, SINGLE_ROW, and REMAINING_ROWS.
RESOLVE_CHANGES_ONLY	INTEGER (1)	0			Indicates that when applying changes during an update that only data that has changed should be applied. Otherwise, all the columns will be updated. This really only applies to updates.
RESOLVE_ROW_ONLY	INTEGER (1)	0			Indicates that an action should take place for the entire batch if possible. This applies to a resolve type of 'ignore'. If a row is in conflict and the resolve type is 'ignore', then the entire batch will be ignored.
CREATE_TIME	TIMESTAMP			X	The date and time when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	The date and time when a user last updated this entry.

A.30. INCOMING_ERROR

The captured data change that is in error for a batch. The user can tell the system what to do by updating the resolve columns. Entries in data_error are created when an incoming batch encounters an error.

Table A.30. INCOMING_ERROR

Name	Type / Size	Default	PK FK	not null	Description
BATCH_ID	BIGINT (50)		PK	X	The id of the outgoing_batch that is being loaded.
NODE_ID	VARCHAR (50)		PK	X	The node_id of the source of the batch being loaded.
FAILED_ROW_NUMBER	BIGINT		PK	X	The row number in the batch that encountered an error when loading.
FAILED_LINE_NUMBER	BIGINT	0		X	The current line number in the CSV for this batch that failed.
TARGET_CATALOG_NAME	VARCHAR (255)				The catalog name for the table being loaded.

Name	Type / Size	Default	PK FK	not null	Description
TARGET_SCHEMA_NAME	VARCHAR (255)				The schema name for the table being loaded.
TARGET_TABLE_NAME	VARCHAR (255)			X	The table name for the table being loaded.
EVENT_TYPE	CHAR (1)			X	The type of event captured by this entry. For triggers, this is the change that occurred, which is 'I' for insert, 'U' for update, or 'D' for delete. Other events include: 'R' for reloading the entire table (or subset of the table) to the node; 'S' for running dynamic SQL at the node, which is used for adhoc administration.
BINARY_ENCODING	VARCHAR (10)	HEX		X	The type of encoding the source system used for encoding binary data.
COLUMN_NAMES	LONGVARCHAR			X	The column names defined on the table. The column names are stored in comma-separated values (CSV) format.
PK_COLUMN_NAMES	LONGVARCHAR			X	The primary key column names defined on the table. The column names are stored in comma-separated values (CSV) format.
ROW_DATA	LONGVARCHAR				The row data from the batch as captured from the source. The column values are stored in comma-separated values (CSV) format.
OLD_DATA	LONGVARCHAR				The old row data prior to update from the batch as captured from the source. The column values are stored in CSV format.
CUR_DATA	LONGVARCHAR				The current row data that caused the error to occur. The column values are stored in CSV format.
RESOLVE_DATA	LONGVARCHAR				The capture data change from the user that is used instead of row_data. This is useful when resolving a conflict manually by specifying the data that should load.
RESOLVE_IGNORE	INTEGER (1)	0			Indication from the user that the row_data should be ignored and the batch can continue loading with the next row.
CONFLICT_ID	VARCHAR (50)				Unique identifier for the conflict detection setting that caused the error
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.31. SEQUENCE

A table that supports application level sequence numbering.

Table A.31. SEQUENCE

Name	Type / Size	Default	PK FK	not null	Description
SEQUENCE_NAME	VARCHAR (50)		PK	X	Unique identifier of a specific sequence.
CURRENT_VALUE	BIGINT	0		X	The current value of the sequence.
INCREMENT_BY	INTEGER	1		X	Specify the interval between sequence numbers. This integer value can be any positive or negative integer, but it cannot be 0.
MIN_VALUE	BIGINT	1		X	Specify the minimum value of the sequence.
MAX_VALUE	BIGINT	9999999999		X	Specify the maximum value the sequence can generate.
CYCLE	INTEGER (1)	0			Indicate whether the sequence should automatically cycle once a boundary is hit.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.32. LOAD_FILTER

A table that allows you to dynamically define filters using bsh.

Table A.32. LOAD_FILTER

Name	Type / Size	Default	PK FK	not null	Description
LOAD_FILTER_ID	VARCHAR (50)		PK	X	The id of the load filter.
LOAD_FILTER_TYPE	VARCHAR (10)			X	The type of load filter. Currently 'bsh'. May add 'sql' in the future.
SOURCE_NODE_GROUP_ID	VARCHAR (50)			X	The source node group for the filter.
TARGET_NODE_GROUP_ID	VARCHAR (50)			X	The destination node group for the filter.
TARGET_CATALOG_NAME	VARCHAR (255)				Optional name for the catalog the configured table is in.
TARGET_SCHEMA_NAME	VARCHAR (255)				Optional name for the schema a configured table is in.

Name	Type / Size	Default	PK FK	not null	Description
TARGET_TABLE_NAME	VARCHAR (255)			X	The name of the target table that will trigger the bsh filter.
FILTER_ON_UPDATE	INTEGER (1)	1		X	Whether or not the filter should apply on an update.
FILTER_ON_INSERT	INTEGER (1)	1		X	Whether or not the filter should apply on an insert.
FILTER_ON_DELETE	INTEGER (1)	1		X	Whether or not the filter should apply on a delete.
BEFORE_WRITE_SCRIPT	LONGVARCHAR				The script to apply before the write is completed.
AFTER_WRITE_SCRIPT	LONGVARCHAR				The script to apply after the write is completed.
BATCH_COMPLETE_SCRIPT	LONGVARCHAR				The script to apply on batch complete.
BATCH_COMMIT_SCRIPT	LONGVARCHAR				The script to apply on batch commit.
BATCH_ROLLBACK_SCRIPT	LONGVARCHAR				The script to apply on batch rollback.
HANDLE_ERROR_SCRIPT	LONGVARCHAR				The script to apply when data cannot be processed.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.
LOAD_FILTER_ORDER	INTEGER	1		X	Specifies the order in which to apply load filters if more than one target operation occurs.
FAIL_ON_ERROR	INTEGER (1)	0		X	Whether we should fail the batch if the filter fails.

A.33. TABLE_RELOAD_REQUEST

This table acts as a means to queue up a reload of a specific table. Either the target or the source node may insert into this table to queue up a load. If the target node inserts into the table, then the row will be synchronized to the source node and the reload events will be queued up during routing.

Table A.33. TABLE_RELOAD_REQUEST

Name	Type / Size	Default	PK FK	not null	Description
TARGET_NODE_ID	VARCHAR (50)		PK	X	Unique identifier for the node to receive the table reload.
SOURCE_NODE_ID	VARCHAR (50)		PK	X	Unique identifier for the node that will be the source of the table reload.

Name	Type / Size	Default	PK FK	not null	Description
TRIGGER_ID	VARCHAR (50)		PK	X	Unique identifier for a trigger associated with the table reload. Note the trigger must be linked to the router.
ROUTER_ID	VARCHAR (50)		PK	X	Unique description of the router associated with the table reload. Note the router must be linked to the trigger.
RELOAD_SELECT	LONGVARCHAR				Overrides the initial load select.
RELOAD_DELETE_STMT	LONGVARCHAR				Overrides the initial load delete statement.
RELOAD_ENABLED	INTEGER (1)	0			Indicates that a reload should be queued up.
RELOAD_TIME	TIMESTAMP				The timestamp when the reload was started for this node.
CREATE_TIME	TIMESTAMP				Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.34. GROUPLET

This tables defines named groups to which nodes can belong to based on their external id. Grouplets are used to designate that synchronization should only affect an explicit subset of nodes in a node group.

Table A.34. GROUPLET

Name	Type / Size	Default	PK FK	not null	Description
GROUPLET_ID	VARCHAR (50)		PK	X	Unique identifier for the grouplet.
GROUPLET_LINK_POLICY	CHAR (1)	I		X	Specified whether the external ids in the grouplet_link are included in the group or excluded from the grouplet. In the case of excluded, the grouplet starts with all external ids and removes the excluded ones listed. Use 'I' for inclusive and 'E' for exclusive.
DESCRIPTION	VARCHAR (255)				A description of this grouplet.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.35. GROUPLET_LINK

This tables defines nodes belong to a grouplet based on their external.id

Table A.35. GROUPLET_LINK

Name	Type / Size	Default	PK FK	not null	Description
grouplet_id	VARCHAR (50)		PK FK	X	Unique identifier for the grouplet.
EXTERNAL_ID	VARCHAR (50)		PK	X	Provides a means to select the nodes that belong to a grouplet.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

A.36. TRIGGER_ROUTER_GROUPLET

This tables defines what grouplets are associated with what trigger routers and if they apply to source node groups or target node groups

Table A.36. TRIGGER_ROUTER_GROUPLET

Name	Type / Size	Default	PK FK	not null	Description
grouplet_id	VARCHAR (50)		PK FK	X	Unique identifier for the grouplet.
trigger_id	VARCHAR (50)		PK FK	X	The id of a trigger.
router_id	VARCHAR (50)		PK FK	X	The id of a router.
APPLIES_WHEN	CHAR (1)		PK	X	Indicates the side that a grouplet should be applied to. Use 'T' for target and 'S' for source and 'B' for both source and target.
CREATE_TIME	TIMESTAMP			X	Timestamp when this entry was created.
LAST_UPDATE_BY	VARCHAR (50)				The user who last updated this entry.
LAST_UPDATE_TIME	TIMESTAMP			X	Timestamp when a user last updated this entry.

Appendix B. Parameters

There are two kinds of parameters that can be used to configure the behavior of SymmetricDS: *Startup Parameters* and *Runtime Parameters* . Startup Parameters are required to be in a system property or a property file, while Runtime Parameters can also be found in the Parameter table from the database. Parameters are re-queried from their source at a configured interval and can also be refreshed on demand by using the JMX API. The following table shows the source of parameters and the hierarchy of precedence.

Table B.1. Parameter Locations

Location	Required	Description
symmetric-default.properties	Y	Packaged inside symmetric-core jar file. This file has all the default settings along with descriptions.
conf/symmetric.properties	N	Changes to this file in the conf directory of a standalone install apply to all engines in the JVM.
symmetric-override.properties	N	Changes to this file, provided by the end user in the JVM's classpath, apply to all engines in the JVM.
engines/*.properties	N	Properties for a specific engine or node that is hosted in a standalone install.
Java System Properties	N	Any SymmetricDS property can be passed in as a -D property to the runtime. It will take precedence over any properties file property.
Parameter table	N	A table which contains SymmetricDS parameters. Parameters can be targeted at a specific node group and even at a specific external id. These settings will take precedence over all of the above.
IParameterFilter	N	An extension point which allows parameters to be sourced from another location or customized. These settings will take precedence over all of the above.

B.1. Startup Parameters

Startup parameters are read once from properties files and apply only during start up. The following properties are used:

engine.name
This is the engine name. This should be set if you have more than one engine running in the same JVM. It is used to name the JMX management bean. [Default: Default]

registration.url
The URL where this node can connect for registration to receive its configuration. This property is only valid if you use the default IRuntimeConfiguration implementation. [Default:]

sync.url
The URL where this node can be contacting for synchronization.

[Default: http://localhost:8080/sync]

group.id
The node group id for this node. [Default: default]

external.id
The secondary identifier for this node that has meaning to the system where it is deployed. While the node id is a generated sequence number, the external ID could have meaning in the user's domain, such as a retail store number. [Default:]

db.driver
The class name of the JDBC driver. If `db.jndi.name` is set, this property is ignored.
[Default: com.mysql.jdbc.Driver]

db.url
The JDBC URL used to connect to the database. If `db.jndi.name` is set, this property is ignored.
[Default: jdbc:mysql://localhost/symmetric]

db.user
The database username, which is used to login, create, and update SymmetricDS tables. To use an encrypted username, see Section 5.7, Encrypted Passwords (p. 73) . If `db.jndi.name` is set, this property is ignored. [Default: symmetric]

db.password
The password for the database user. To use an encrypted password, see Section 5.7, Encrypted Passwords (p. 73) . If `db.jndi.name` is set, this property is ignored. [Default:]

db.pool.initial.size
The initial size of the connection pool. If `db.jndi.name` is set, this property is ignored. [Default: 5]

db.pool.max.active
The maximum number of connections that will be allocated in the pool. If `db.jndi.name` is set, this property is ignored. [Default: 10]

db.pool.max.wait.millis
This is how long a request for a connection from the datasource will wait before giving up. If `db.jndi.name` is set, this property is ignored. [Default: 30000]

db.pool.min.evictable.idle.millis
This is how long a connection can be idle before it will be evicted. If `db.jndi.name` is set, this property is ignored. [Default: 120000]

db.sql.query.timeout.seconds
The timeout in seconds for queries running on the database. [Default: 300]

db.jdbc.streaming.results.fetch.size
This is the default fetch size for streaming result sets into memory from the database.
[Default: 1000]

auto.config.database
If this is true, the configuration and runtime tables used by SymmetricDS are automatically created

during startup. [Default: true]

auto.sync.configuration
If this is true, create triggers for the SymmetricDS configuration table that will synchronize changes to node groups that pull from the node where this property is set. [Default: true]

https.allow.self.signed.certs
If this is true, a Symmetric client node to accept self signed certificates. [Default: true]

http.basic.auth.username
If specified, a Symmetric client node will use basic authentication when communicating with its server node using the given user name. [Default:]

http.basic.auth.password
If specified, the password used for basic authentication. [Default:]

https.verified.server.names
A list of comma separated server names that will always verify when using https. This is useful if the URL's hostname and the server's identification hostname don't match exactly using the default rules for the JRE. A special value of "all" may be specified to allow all hostnames to verify. [Default:]

sync.table.prefix
When symmetric tables are created and accessed, this is the prefix to use for the table name.
[Default: sym]

start.route.job
Whether the route job is enabled for this node. [Default: true]

job.route.period.time.ms
This is how often the route job will be run if enabled. [Default: 10000]

start.push.job
Whether the push job is enabled for this node. [Default: true]

job.push.period.time.ms
This is how often the push job will be run if enabled. [Default: 60000]

start.pull.job
Whether the pull job is enabled for this node. [Default: true]

job.pull.period.time.ms
This is how often the pull job will be run if enabled. [Default: 60000]

start.purge.job
Whether the purge jobs are enabled for this node. [Default: true]

job.purge.incoming.cron
This is how often the incoming batch purge job will be run if enabled. [Default: 0 0 0 * * *]

job.purge.outgoing.cron
This is how often the outgoing batch and data purge job will be run if enabled. [Default: 0 0 0 * * *]

job.purge.datagaps.cron
This is how often the data gaps purge job will be run if enabled. [Default: 0 0 0 * * *]

start.synctriggers.job
Whether the sync triggers job is enabled for this node. [Default: true]

job.synctriggers.cron
This is how often the sync triggers job will be run if enabled. [Default: 0 0 0 * * *]

start.stage.management.job
Whether the stage directory management/purge job is enabled for this node. [Default: true]

job.stage.management.period.time.ms
This is how often the stage management job will be run if enabled. [Default: 15000]

start.stat.flush.job
Whether the statistics flush (to database) job is enabled for this node. [Default: true]

job.stat.flush.cron
This is how often accumulated statistics will be flushed out to the database from memory. [Default: 0 0/5 * * * *]

start.heartbeat.job
Whether the heartbeat job is enabled for this node. The heartbeat job simply inserts an event to update the heartbeat_time column on the node table for the current node. [Default: true]

job.heartbeat.period.time.ms
This is how often the heartbeat job runs. Note that this doesn't mean that a heartbeat is performed this often. See heartbeat.sync.on.push.period.sec to change how often the heartbeat is sync'd.
[Default: 1000]

heartbeat.sync.on.push.period.sec
This is the number of seconds between when the sym_node table's heartbeat_time column is updated by the heartbeat job. [Default: 900]

start.watchdog.job
Whether the watchdog job is enabled for this node. The watchdog job monitors child nodes to detect if they are offline. Refer to for more information.
[Default: true]

job.watchdog.period.time.ms
This is how often the watchdog job will be run if enabled. [Default: 3600000]

schema.version
This is hook to give the user a mechanism to indicate the schema version that is being synchronized. This property is only valid if you use the default IRuntimeConfiguration implementation.
[Default: ?]

B.2. Runtime Parameters

Runtime parameters are read periodically from properties files or the database. The following properties are used:

auto.registration
If this is true, registration is opened automatically for nodes requesting it. [Default: false]

auto.reload
If this is true, a reload is automatically sent to nodes when they register. [Default: false]

auto.reload.reverse
If this is true, a reload is automatically sent from a node after they register. [Default: false]

auto.update.node.values.from.properties
Update the node row in the database from the local properties during a heartbeat operation.
[Default: true]

http.concurrent.workers.max
This is the number of HTTP concurrent push/pull requests symmetric will accept. This is controlled by the NodeConcurrencyFilter. The maximum number of database connections in the database pool should be set to twice this number.[Default: 20]

offline.node.detection.period.minutes
This is the minimum number of minutes that a child node has been offline before taking action. Refer to Section 5.10.15, IOfflineServerListener (p. 80) for more information. [Default: 120]

outgoing.batches.peek.ahead.window.after.max.size
This is the maximum number of events that will be peeked at to look for additional transaction rows after the max batch size is reached. The more concurrency in your db and the longer the transaction takes the bigger this value might have to be. [Default: 100]

incoming.batches.skip.duplicates
Whether or not to skip duplicate batches that are received. A duplicate batch is identified by the batch ID already existing in the incoming batch table. If this happens, it means an acknowledgement was lost due to failure or there is a bug. Accepting a duplicate batch in this case can mean overwriting data with old data. Another cause of duplicates is when the batch sequence number is reset, which might happen in a lab environement. Skipping a duplicate batch in this case would prevent data changes from loading. Generally, in a production envionment, this setting should be true. [Default: true]

num.of.ack.retries
This is the number of times we will attempt to send an ACK back to the remote node when pulling and loading data. [Default: 5]

time.between.ack.retries.ms
This is the amount of time to wait between trying to send an ACK back to the remote node when pulling and loading data. [Default: 5000]

dataextractor.enabled

Enable or disable all data extraction at a node for all channels other than the config channel.
[Default: true]

dataloader.enabled

Enable or disable all data loading at a node for all channels other than the config channel.
[Default: true]

cluster.server.id

Set this if you want to give your server a unique name to be used to identify which server did what action. Typically useful when running in a clustered environment. This is currently used by the ClusterService when locking for a node. [Default:]

cluster.lock.timeout.ms

Time limit of lock before it is considered abandoned and can be broken. [Default: 1800000]

cluster.lock.enabled

[Default: false]

initial.load.delete.first

Set this if tables should be purged prior to an initial load. [Default: false]

initial.load.delete.first.sql

This is the SQL statement that will be used for purging a table during an initial load, provided initial.load.delete.first is set to true. [Default: delete from %s]

initial.load.create.first

Set this if tables (and their indexes) should be created prior to an initial load. [Default: false]

http.timeout.ms

Sets both the connection and read timeout on the internal HttpUrlConnection. [Default: 600000s]

http.compression

Whether or not to use compression over HTTP connections. Currently, this setting only affects the push connection of the source node. Compression on a pull is enabled using a filter in the web.xml for the PullServlet. [Default: true]

web.compression.disabled

Disable compression from occurring on Servlet communication. This property only affects the outbound HTTP traffic streamed by the PullServlet and PushServlet. [Default: false]

compression.level

Set the compression level this node will use when compressing synchronization payloads. Valid values include: NO_COMPRESSION = 0, BEST_SPEED = 1, BEST_COMPRESSION = 9, DEFAULT_COMPRESSION = -1 [Default: -1]

compression.strategy

Set the compression strategy this node will use when compressing synchronization payloads. Valid values include: FILTERED = 1, HUFFMAN_ONLY = 2, DEFAULT_STRATEGY = 0 [Default: 0]

stream.to.file.enabled

Save data to the file system before transporting it to the client or loading it to the database if the

number of bytes is past a certain threshold. This allows for better compression and better use of database and network resources. Statistics in the batch tables will be more accurate if this is set to true because each timed operation is independent of the others. [Default: true]

stream.to.file.threshold.bytes
If stream.to.file.enabled is true, then the threshold number of bytes at which a file will be written is controlled by this property. Note that for a synchronization the entire payload of the synchronization will be buffered in memory up to this number (at which point it will be written and continue to stream to disk) [Default: 32767]

job.random.max.start.time.ms
When starting jobs, symmetric attempts to randomize the start time to spread out load. This is the maximum wait period before starting a job. [Default: 10000]

purge.retention.minutes
This is the retention for how long synchronization data will be kept in the SymmetricDS synchronization tables. Note that data will be purged only if the purge job is enabled. [Default: 7200]

Appendix C. Database Notes

Each database management system has its own characteristics that results in feature coverage in SymmetricDS. The following table shows which features are available by database.

Table C.1. Support by Database

Database	Versions supported	Transaction Identifier	Data Capture	Conditional Sync	Update Loop Prevention	BLOB Sync
Oracle	10g and above	Y	Y	Y	Y	Y
MySQL	5.0.2 and above	Y	Y	Y	Y	Y
PostgreSQL	8.2.5 and above	Y (8.3 and above only)	Y	Y	Y	Y
Greenplum	8.2.15 and above	N	N	N	Y	N
SQL Server	2005 and above	Y	Y	Y	Y	Y
SQL Server Azure	Tested on 11.00.2065	Y	Y	Y	Y	Y
HSQLDB	1.8	Y	Y	Y	Y	Y
HSQLDB	2.0	N	Y	Y	Y	Y
H2	1.x	Y	Y	Y	Y	Y
Apache Derby	10.3.2.1	Y	Y	Y	Y	Y
IBM DB2	9.5	N	Y	Y	Y	Y
Firebird	2.0	Y	Y	Y	Y	Y
Informix	11	N	Y	Y	Y	N
Interbase	9.0	N	Y	Y	Y	Y
SQLite	3.x	N	Y	Y	Y	Y

C.1. Oracle

While BLOBs are supported on Oracle, the LONG data type is not. LONG columns cannot be accessed from triggers.

Note that while Oracle supports multiple triggers of the same type to be defined, the order in which the triggers occur appears to be arbitrary.

The SymmetricDS user generally needs privileges for connecting and creating tables (including indexes), triggers, sequences, and procedures (including packages and functions). The following is an example of the needed grant statements:

```
GRANT CONNECT TO SYMMETRIC;
GRANT RESOURCE TO SYMMETRIC;
GRANT CREATE ANY TRIGGER TO SYMMETRIC;
GRANT EXECUTE ON UTL_RAW TO SYMMETRIC;
```

Partitioning the DATA table by channel can help insert, routing and extraction performance on concurrent, high throughput systems. TRIGGERs should be organized to put data that is expected to be inserted concurrently on separate CHANNELs. The following is an example of partitioning. Note that both the table and the index should be partitioned. The default value allows for more channels to be added without having to modify the partitions.

```
CREATE TABLE SYM_DATA
(
    data_id INTEGER NOT NULL ,
    table_name VARCHAR2(50) NOT NULL,
    event_type CHAR(1) NOT NULL,
    row_data CLOB,
    pk_data CLOB,
    old_data CLOB,
    trigger_hist_id INTEGER NOT NULL,
    channel_id VARCHAR2(20),
    transaction_id VARCHAR2(1000),
    source_node_id VARCHAR2(50),
    external_data VARCHAR2(50),
    create_time TIMESTAMP
) PARTITION BY LIST (channel_id) (
PARTITION P_CONFIG VALUES ('config'),
PARTITION P_CHANNEL_ONE VALUES ('channel_one'),
PARTITION P_CHANNEL_TWO VALUES ('channel_two'),
...
PARTITION P_CHANNEL_N VALUES ('channel_n'),
PARTITION P_DEFAULT VALUES (DEFAULT));
```

```
CREATE UNIQUE INDEX IDX_D_CHANNEL_ID ON SYM_DATA (DATA_ID, CHANNEL_ID)  LOCAL
(
 PARTITION I_CONFIG,
 PARTITION I_CHANNEL_ONE,
 PARTITION I_CHANNEL_TWO,
 ...
 PARTITION I_CHANNEL_N,
 PARTITION I_DEFAULT
);
```

Note also that, for Oracle, you can control the amount of precision used by the Oracle triggers with the parameter `oracle.template.precision`, which defaults to a precision of 30,10.

If the following Oracle error 'ORA-01489: result of string concatenation is too long' is encountered you might need to set `use_capture_lobs` to 1 on in the TRIGGER table and resync the triggers. The error can happen when the captured data in a row exceeds 4k and lob columns do not exist in the table. By enabling

`use_capture_lobs` the concatanated varchar string is cast to a clob which allows a length of more than 4k.

C.2. MySQL

MySQL supports several storage engines for different table types. SymmetricDS requires a storage engine that handles transaction-safe tables. The recommended storage engine is InnoDB, which is included by default in MySQL 5.0 distributions. Either select the InnoDB engine during installation or modify your server configuration. To make InnoDB the default storage engine, modify your MySQL server configuration file (`my.ini` on Windows, `my.cnf` on Unix):

```
default-storage_engine = innodb
```

Alternatively, you can convert tables to the InnoDB storage engine with the following command:

```
alter table t engine = innodb;
```

On MySQL 5.0, the SymmetricDS user needs the SUPER privilege in order to create triggers.

```
grant super on *.* to symmetric;
```

On MySQL 5.1, the SymmetricDS user needs the TRIGGER and CREATE ROUTINE privileges in order to create triggers and functions.

```
grant trigger on *.* to symmetric;
```

```
grant create routine on *.* to symmetric;
```

MySQL allows '0000-00-00 00:00:00' to be entered as a value for datetime and timestamp columns. JDBC can not deal with a date value with a year of 0. In order to work around this SymmetricDS can be configured to treat date and time columns as varchar columns for data capture and data load. To enable this feature set the `db.treat.date.time.as.varchar.enabled` property to `true`.

C.3. PostgreSQL

Starting with PostgreSQL 8.3, SymmetricDS supports the transaction identifier. Binary Large Object (BLOB) replication is supported for both byte array (BYTEA) and object ID (OID) data types.

In order to function properly, SymmetricDS needs to use session variables. On PostgreSQL, session variables are enabled using a custom variable class. Add the following line to the `postgresql.conf` file of PostgreSQL server:

```
custom_variable_classes = 'symmetric'
```

This setting is required, and SymmetricDS will log an error and exit if it is not present.

Before database triggers can be created by in PostgreSQL, the plpgsql language handler must be installed on the database. The following statements should be run by the administrator on the database:

```
CREATE FUNCTION plpgsql_call_handler() RETURNS language_handler AS
    '$libdir/plpgsql' LANGUAGE C;

CREATE FUNCTION plpgsql_validator(oid) RETURNS void AS
    '$libdir/plpgsql' LANGUAGE C;

CREATE TRUSTED PROCEDURAL LANGUAGE plpgsql
    HANDLER plpgsql_call_handler
    VALIDATOR plpgsql_validator;
```

If you want SymmetricDS to install into a schema other than public you should alter the database user to set the default schema.

```
        alter user {user name} set search_path to {schema name};
```

In addition, you will likely need the follow privelegdes as well:

```
        GRANT USAGE ON SCHEMA {schema name} TO {user name};
        GRANT CREATE ON SCHEMA {schema name} TO {user name};
```

C.4. Greenplum

Greenplum is a data warehouse based on PostgreSQL. It is supported as a target platform in SymmetricDS. For the best performance, the SymmetricDS Pro PostgreSQL bulk loader should be used.

C.5. MS SQL Server

SQL Server was tested using the jTDS JDBC driver.

C.6. HSQLDB

HSQLDB was implemented with the intention that the database be run embedded in the same JVM process as SymmetricDS. Instead of dynamically generating static SQL-based triggers like the other databases, HSQLDB triggers are Java classes that re-use existing SymmetricDS services to read the configuration and insert data events accordingly.

The transaction identifier support is based on SQL events that happen in a 'window' of time. The trigger(s) track when the last trigger fired. If a trigger fired within X milliseconds of the previous firing, then the current event gets the same transaction identifier as the last. If the time window has passed, then

a new transaction identifier is generated.

C.7. H2

The H2 database allows only Java-based triggers. Therefore the H2 dialect requires that the SymmetricDS jar file be in the database's classpath.

C.8. Apache Derby

The Derby database can be run as an embedded database that is accessed by an application or a standalone server that can be accessed from the network. This dialect implementation creates database triggers that make method calls into Java classes. This means that the supporting JAR files need to be in the classpath when running Derby as a standalone database, which includes symmetric-ds.jar and commons-lang.jar.

C.9. IBM DB2

The DB2 Dialect uses global variables to enable and disable node and trigger synchronization. These variables are created automatically during the first startup. The DB2 JDBC driver should be placed in the "lib" folder.

Currently, the DB2 Dialect for SymmetricDS does not provide support for transactional synchronization. Large objects (LOB) are supported, but are limited to 16,336 bytes in size. The current features in the DB2 Dialect have been tested using DB2 9.5 on Linux and Windows operating systems.

There is currently a bug with the retrieval of auto increment columns with the DB2 9.5 JDBC drivers that causes some of the SymmetricDS configuration tables to be rebuilt when auto.config.database=true. The DB2 9.7 JDBC drivers seem to have fixed the issue. They may be used with the 9.5 database.

A system temporary tablespace with too small of a page size may cause the following trigger build errors:

```
SQL1424N Too many references to transition variables and transition table
columns or the row length for these references is too long. Reason
code="2". LINE NUMBER=1. SQLSTATE=54040
```

Simply create a system temporary tablespace that has a bigger page size. A page size of 8k will probably suffice.

C.10. Firebird

The Firebird Dialect requires the installation of a User Defined Function (UDF) library in order to provide functionality needed by the database triggers. SymmetricDS includes the required UDF library,

called SYM_UDF, in both source form (as a C program) and as pre-compiled libraries for both Windows and Linux. The SYM_UDF library is copied into the UDF folder within the Firebird installation directory.

For Linux users:

cp databases/firebird/sym_udf.so /opt/firebird/UDF

For Windows users:

copy databases\firebird\sym_udf.dll C:\Program Files\Firebird\Firebird_2_0\UDF

The following limitations currently exist for this dialect:

- The outgoing batch does not honor the channel size, and all outstanding data events are included in a batch.
- Syncing of Binary Large Object (BLOB) is limited to 16K bytes per column.
- Syncing of character data is limited to 32K bytes per column.

C.11. Informix

The Informix Dialect was tested against Informix Dynamic Server 11.50, but older versions may also work. You need to download the Informix JDBC Driver (from the [IBM Download Site](#)) and put the `ifxjdbc.jar` and `ifxlang.jar` files in the SymmetricDS `lib` folder.

Make sure your database has logging enabled, which enables transaction support. Enable logging when creating the database, like this:

```
CREATE DATABASE MYDB WITH LOG;
```

Or enable logging on an existing database, like this:

```
ondblog mydb unbuf log
ontape -s -L 0
```

The following features are not yet implemented:

- Syncing of Binary and Character Large Objects (LOB) is disabled.
- There is no transaction ID recorded on data captured, so it is possible for data to be committed within different transactions on the target database. If transaction synchronization is required, either specify a custom transaction ID or configure the synchronization so data is always sent in a single batch. A custom transaction ID can be specified with the tx_id_expression on TRIGGER. The batch size is controlled with the max_batch_size on CHANNEL. The pull and push jobs have runtime properties to control their interval.

C.12. Interbase

The Interbase Dialect requires the installation of a User Defined Function (UDF) library in order to provide functionality needed by the database triggers. SymmetricDS includes the required UDF library, called SYM_UDF, in both source form (as a C program) and as pre-compiled libraries for both Windows and Linux. The SYM_UDF library is copied into the UDF folder within the Interbase installation directory.

For Linux users:

cp databases/interbase/sym_udf.so /opt/interbase/UDF

For Windows users:

copy databases\interbase\sym_udf.dll C:\CodeGear\InterBase\UDF

The Interbase dialect currently has the following limitations:
- Data capture is limited to 4 KB per row, including large objects (LOB).
- There is no transaction ID recorded on data captured. Either specify a tx_id_expression on the TRIGGER table, or set a max_batch_size on the CHANNEL table that will accommodate your transactional data.

C.13. SQLite

For SQLite, the implementation of sync-on-incoming back and the population of a source node if in the sym data rows relies on use of a context table (by default, called sym_context) to hold a boolean and node id in place of the more common methods of using temp tables (which are unaccessible from triggers) or functions (which are not available). The context table assumes there's a single thread updating the database at any one time. If that is not the case in the future, the current implementation of sync on incoming batch will be unreliable.

Nodes using SQLite should have the `jobs.synchronized.enable` parameter set to `true`. This parameter causes the jobs and push/pull threads to all run in a synchronized fashion, which is needed in the case of SQLite.

The SQLite dialect has the following limitations:
- There is no transaction ID recorded on data captured. Either specify a tx_id_expression on the TRIGGER table, or set a max_batch_size on the CHANNEL table that will accommodate your transactional data.
- Due to the single threaded access to SQLite, the following parameter should be set to true: `jobs.synchronized.enable`.

Appendix D. Data Format

The SymmetricDS Data Format is used to stream data from one node to another. The data format reader and writer are pluggable with an initial implementation using a format based on Comma Separated Values (CSV). Each line in the stream is a record with fields separated by commas. String fields are surrounded with double quotes. Double quotes and backslashes used in a string field are escaped with a backslash. Binary values are represented as a string with hex values in "\0xab" format. The absence of any value in the field indicates a null value. Extra spacing is ignored and lines starting with a hash are ignored.

The first field of each line gives the directive for the line. The following directives are used:

nodeid, {node_id}
Identifies which node the data is coming from. Occurs once in CSV file.

binary, {BASE64|NONE|HEX}
Identifies the type of decoding the loader needs to use to decode binary data in the pay load. This varies depending on what database is the source of the data.

channel, {channel_id}
Identifies which channel a batch belongs to. The SymmetricDS data loader expects the channel to be specified before the batch.

batch, {batch_id}
Uniquely identifies a batch. Used to track whether a batch has been loaded before. A batch of -9999 is considered a virtual batch and will be loaded, but will not be recorded in incoming_batch.

schema, {schema name}
The name of the schema that is being targeted.

catalog, {catalog name}
The name of the catalog that is being targeted.

table, {table name}
The name of the table that is being targeted.

keys, {column name...}
Lists the column names that are used as the primary key for the table. Only needs to occur after the first occurrence of the table.

columns, {column name...}
Lists all the column names (including key columns) of the table. Only needs to occur after the first occurrence of the table.

insert, {column value...}
Insert into the table with the values that correspond with the columns.

update, {new column value...},{old key value...}
Update the table using the old key values to set the new column values.

old, {old column value...}
Represent all the old values of the data. This data can be used for conflict resolution.

delete, {old key value...}
Delete from the table using the old key values.

sql, {sql statement}
Optional notation that instructs the data loader to run the accompanying SQL statement.

bsh, {bsh script}
Optional notation that instructs the data loader to run the accompanying [BeanShell](#) snippet.

create, {xml}
Optional notation that instructs the data loader to run the accompanying [DdlUtils](#) XML table definition in order to create a database table.

commit, {batch_id}
An indicator that the batch has been transmitted and the data can be committed to the database.

Example D.1. Data Format Stream

```
nodeid, 1001
channel, pricing
binary, BASE64
batch, 100
schema,
catalog,
table, item_selling_price
keys, price_id
columns, price_id, price, cost
insert, 55, 0.65, 0.55
schema,
catalog,
table, item
keys, item_id
columns, item_id, price_id, name
insert, 110000055, 55, "Soft Drink"
delete, 110000001
schema,
catalog,
table, item_selling_price
update, 55, 0.75, 0.65, 55
commit, 100
```

Appendix E. Upgrading from 2.x

Please test carefully when upgrading SymmetricDS 2 to SymmetricDS 3. Note that
OUTGOING_BATCH table's primary key changed. The automatic upgrade backs up and copies the
table. This might take some time if the table is large.

The following parameters are no longer supported:

- `db.spring.bean.name` - The connection pool is no longer wired in via the Spring Framework

- `db.tx.timeout.seconds` - Transactions are no longer managed by the Spring Framework

- `db.default.schema` - The default schema is always the schema associated with the database user

- `db.jndi.name` - JNDI data sources are no longer supported

- `auto.upgrade` - Database upgrade is controlled by `auto.config.database`

- `routing.data.reader.type` - As of this release, there is only one data reader type.

- `job.purge.max.num.data.events.to.delete.in.tx` - The name of this property changed to
  `job.purge.max.num.data.event.batches.to.delete.in.tx`

- `web.base.servlet.path` - No longer needed

- `dataloader.allow.missing.delete` - Controlled by conflict detection and resolution

- `dataloader.enable.fallback.insert` - Controlled by conflict detection and resolution

- `dataloader.enable.fallback.update` - Controlled by conflict detection and resolution

- `dataloader.enable.fallback.savepoint` - No longer needed

- `db.force.delimited.identifier.mode.on` - No longer needed

- `db.force.delimited.identifier.mode.off` - No longer needed

The way extension points work has changed. SymmetricDS services are no longer Spring injectable into
extension points. Please use the `ISymmetricEngineAware` interface to get a handle to the engine which gives
access to services.

The following extension points are no longer supported:

- `IDataLoaderFilter` - Replaced by IDatabaseWriterFilter

- `IBatchListener` - Replaced by IDatabaseWriterFilter

- `IExtractorFilter` - No longer supported. Rarely used.

- `IColumnFilter` - No longer needed. Please use the transformation feature.

Appendix F. Version Numbering

The software is released with a version number based on the [Apache Portable Runtime Project](#) version guidelines. In summary, the version is denoted as three integers in the format of MAJOR.MINOR.PATCH. Major versions are incompatible at the API level, and they can include any kind of change. Minor versions are compatible with older versions at the API and binary level, and they can introduce new functions or remove old ones. Patch versions are perfectly compatible, and they are released to fix defects.