

# MySQL Fabric 1.5

---

## Abstract

MySQL Fabric is a system for managing a farm of MySQL servers. MySQL Fabric provides an extensible and easy to use system for managing a MySQL deployment for sharding and high-availability.

This document describes MySQL Fabric, beginning with a short introduction, providing instructions on how to download and install MySQL Fabric, and a quick-start guide to help you begin using and experimenting with MySQL Fabric. Later sections provide details for MySQL Fabric-aware connectors.

For legal information, see the [Legal Notices](#).

For help with using MySQL, please visit either the [MySQL Forums](#) or [MySQL Mailing Lists](#), where you can discuss your issues with other MySQL users.

For additional documentation on MySQL products, including translations of the documentation into other languages, and downloadable versions in variety of formats, including HTML and PDF formats, see the [MySQL Documentation Library](#).

**Licensing information.** This product may include third-party software, used under license. If you are using a *Commercial* release of MySQL Fabric, see [this document](#) for licensing information, including licensing information relating to third-party software that may be included in this Commercial release. If you are using a *Community* release of MySQL Fabric, see [this document](#) for licensing information, including licensing information relating to third-party software that may be included in this Community release.

Document generated on: 2017-07-17 (revision: 52970)

---

---

# Table of Contents

Preface and Legal Notices .....	v
1 Introduction to Fabric .....	1
1.1 Fabric Prerequisites .....	1
1.2 Fabric Concepts .....	2
2 Installing and Configuring MySQL Fabric .....	3
2.1 Downloading MySQL Fabric .....	3
2.2 Installing MySQL Fabric .....	3
2.3 Configuring MySQL Fabric .....	3
2.3.1 Create the Associated MySQL Users .....	4
2.3.2 Configuration File .....	7
2.3.3 Configuration File Sections .....	9
2.3.4 The Configuration Parameter (--param) .....	14
2.4 Starting and Stopping MySQL Fabric Nodes .....	14
2.5 Old Configuration System .....	15
3 Quick Start .....	19
3.1 Example: Fabric and Replication .....	19
3.2 Example: Fabric and Sharding .....	23
3.2.1 Introduction to Sharding .....	23
3.2.2 Sharding Scenario .....	24
4 The <code>mysqfabric</code> Utility .....	29
4.1 Getting Help .....	29
4.2 Dump Commands .....	30
4.3 Event Commands .....	31
4.4 Group Commands .....	32
4.5 Manage Commands .....	36
4.6 Provider Commands .....	37
4.7 Role Commands .....	38
4.8 Server Commands .....	38
4.9 Sharding Commands .....	42
4.10 Snapshot Commands .....	46
4.11 Statistics Commands .....	47
4.12 Threat Commands .....	47
4.13 User Commands .....	48
5 Fabric Utility Command Matrix .....	51
6 Backing Store .....	57
6.1 Backing Store Tables .....	57
6.2 Protecting the Backing Store .....	60
7 Using MySQL Fabric with Pacemaker and Corosync .....	61
7.1 Introduction .....	61
7.2 Pre-requisites .....	62
7.3 Target Configuration .....	62
7.4 Setting up and testing your system .....	63
7.4.1 Configure Network .....	63
7.4.2 Install all packages .....	63
7.4.3 Configure DRBD .....	64
7.4.4 Configure MySQL Server .....	66
7.4.5 Configure MySQL Fabric .....	67
7.4.6 Configure Corosync and Pacemaker .....	68
7.5 Key administrative tasks .....	71
8 Using Connector/Python with MySQL Fabric .....	73
8.1 Installing Connector/Python with MySQL Fabric Support .....	74

8.2	Requesting a Fabric Connection .....	74
8.3	Providing Information to Choose a MySQL Server .....	76
9	Using Connector/J with MySQL Fabric .....	79
9.1	Installing Connector/J with MySQL Fabric Support .....	79
9.2	Loading the Driver and Requesting a Fabric Connection .....	80
9.3	Providing Information to Choose a MySQL Server .....	80
9.4	MySQL Fabric Configuration for Running Samples .....	82
9.5	Running Tests .....	83
9.6	Running Demonstration Programs .....	84
9.7	A Complete Example: Working with Employee Data .....	84
9.8	How Connector/J Chooses a MySQL Server .....	88
9.9	Using Hibernate with MySQL Fabric .....	89
9.10	Connector/J Fabric Support Reference .....	92
9.10.1	Connection Properties .....	92
9.10.2	FabricMySQLConnection API .....	93
10	Using Connector/Net with MySQL Fabric .....	95
10.1	System Requirements .....	95
10.2	Set up the MySQL Fabric Plugin .....	95
10.3	Using MySQL Fabric Groups .....	97
10.4	Using Ranged Sharding .....	99
11	MySQL Workbench and MySQL Fabric Integration .....	103
12	MySQL Fabric Frequently Asked Questions .....	105

---

# Preface and Legal Notices

This document describes MySQL Fabric, a system for managing a farm of MySQL servers. MySQL Fabric provides an extensible and easy to use system for managing a MySQL deployment for sharding and high-availability.

## Legal Notices

Copyright © 1997, 2017, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

**U.S. GOVERNMENT END USERS:** Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

## Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

---

# Chapter 1 Introduction to Fabric

## Table of Contents

1.1 Fabric Prerequisites .....	1
1.2 Fabric Concepts .....	2

To take advantage of Fabric, an application requires an augmented version of a MySQL connector which accesses Fabric using the XML-RPC protocol. Currently, Connector/Python and Connector/J are fabric-aware.

Fabric manages sets of MySQL Servers that have Global Transaction Identifiers (GTIDs) enabled to check and maintain consistency among servers. Sets of servers are called *high-availability groups*. Information about all of the servers and groups is managed by a separate MySQL instance, which cannot be a member of the Fabric high-availability groups. This server instance is called the *backing store*.

Fabric organizes servers in groups (called *high-availability groups*) for managing different shards or simply for providing high availability. For example, if standard asynchronous replication is in use, Fabric may be configured to automatically monitor the status of servers in a group. If the current master in a group fails, it elects a new one if a server in the group can become a master.

Besides the high-availability operations such as failover and switchover, Fabric also permits shard operations such as shard creation and removal.

Fabric is written in Python and includes a special library that implements all of the functionality provided. To interact with Fabric, a special utility named `mysqlfabric` provides a set of commands you can use to create and manage groups, define and manipulate sharding, and much more.

## 1.1 Fabric Prerequisites

Fabric is designed to work with MySQL servers version 5.6.10 and later. The `mysqlfabric` utility requires Python 2 (2.6 or later) and Connector/Python 1.2.1 or later.

You must have a MySQL instance available to install the backing store. This server must not be a member of a Fabric group.



### Note

The backing store must be a MySQL server 5.6 or later.

To utilize Fabric in your applications, you must have a Fabric-aware connector installed on the system where the application is running. For more information about using a connector with Fabric, see the appropriate connector-specific section ([Chapter 8, Using Connector/Python with MySQL Fabric](#), [Chapter 9, Using Connector/J with MySQL Fabric](#)).

In summary, the following items indicate the prerequisites for using MySQL Fabric:

- MySQL Server 5.6.10 or later for MySQL servers managed by Fabric. MySQL Server requirements: `gtid_mode` (GTID), `bin_log` (binary logging), and `log_slave_updates` enabled, with `server_id` properly configured.
- MySQL server 5.6.x or later for the backing store.

- Python 2 (2.6 or later) for the `mysqlfabric` utility.
- A Fabric-aware connector to use Fabric in applications. Permitted connectors and versions are:
  - Connector/Python 1.2.1 or later
  - Connector/J 5.1.27 or later

## 1.2 Fabric Concepts

This section describes some of the concepts used in Fabric.

A *high-availability group*, or simply *group*, is a collection of servers. It is used to associate the servers in a set. This association may be a set of replication-enabled servers, the servers participating in a sharding solution, and so forth.

A *group identifier* is the name we give a group or members of the group. A group identifier is a name that matches the regular expression `[a-zA-Z0-9_-]+`. Examples of legal identifiers are `my_group`, `employees`, and `shard1`.

A *global group* stores all updates that must be propagated to all shards that are part of a sharding scheme.

A *node* or *fabric node* is an instance of the Fabric system running on a server. To use the features of Fabric, at least one Fabric node must be running.

*Sharding* refers to the Fabric feature that permits the distribution of data across several servers. There are many uses of sharding but the most effective use of sharding enables distributing the work of writing across several servers for improved write speeds.

A *shard* is a horizontal partition or segment of data in a table.

*Primary* refers to a member of a group that is designated as master in the sense that it can accept read-write transactions.

*Secondary* refers to a member of a group that can be a candidate to replace the master during switchover or failover and that can accept read-only transactions.



---

# Chapter 2 Installing and Configuring MySQL Fabric

## Table of Contents

2.1 Downloading MySQL Fabric .....	3
2.2 Installing MySQL Fabric .....	3
2.3 Configuring MySQL Fabric .....	3
2.3.1 Create the Associated MySQL Users .....	4
2.3.2 Configuration File .....	7
2.3.3 Configuration File Sections .....	9
2.3.4 The Configuration Parameter (--param) .....	14
2.4 Starting and Stopping MySQL Fabric Nodes .....	14
2.5 Old Configuration System .....	15

To use MySQL Fabric, you must have a set of MySQL server instances running MySQL 5.6.10 or higher. One server is required for the backing store and at least one server must be added to a group. To use the replication features of MySQL Fabric, a replication topology of a master and at least one slave is required. To use the sharding features of MySQL Fabric, you should have the number of servers corresponding to the depth of the shards (the number of segments).



### Note

The configuration system changed in Fabric 1.5.5. The current system is documented here, and the old system is documented under [Section 2.5, “Old Configuration System”](#).

For instructions on how to download and install MySQL server, see the online MySQL reference manual ([Installing and Upgrading MySQL](#)).

## 2.1 Downloading MySQL Fabric

Download a version of MySQL Fabric from the MySQL Developer Zone website (<http://dev.mysql.com/downloads/utilities/>). Packaged downloads are available for a variety of servers. Download the package that matches your platform and extract the files.

The above website also provides the MySQL Fabric-aware connectors. Download the connector you want to use with a MySQL Fabric application. For more information about how to install and get started using MySQL Fabric in your applications, see the appropriate connector-specific section ([Chapter 8, Using Connector/Python with MySQL Fabric](#), [Chapter 9, Using Connector/J with MySQL Fabric](#)).

## 2.2 Installing MySQL Fabric

To install MySQL Fabric, install MySQL Utilities 1.5.6. For more information, see [How to Install MySQL Utilities](#).

## 2.3 Configuring MySQL Fabric

Configuring MySQL Fabric requires creating separate MySQL users to access the backing store and managed MySQL servers, and editing the configuration file with the MySQL user details for all of these users.

Each managed MySQL Server has the following requirements: `gtid_mode` (GTID), `bin_log` (binary logging), and `log_slave_updates` enabled, with `server_id` properly configured.



**Note**

The configuration system changed in MySQL Fabric 1.5.5. For information about the previous configuration system, see [Section 2.5, “Old Configuration System”](#).

## 2.3.1 Create the Associated MySQL Users

Fabric uses four different types of users, each with a different set of required privileges.



**Note**

The backup and restore users were added in Fabric 1.5.5.

- **Backing store user:** stores Fabric specific information, and is only created on the Fabric backing store MySQL server. For additional information, see [Chapter 6, Backing Store](#)
- **Server user:** accesses the managed MySQL servers, and is created on each managed MySQL server.
- **Backup user:** executes backup operations, such as `mysqldump`, and is created on each managed MySQL server.
- **Restore user:** executes restore operations that typically use the `mysql` client, and is created on each managed MySQL server.

### Privileges

It is possible to use the same MySQL account for the server user, backup user, and restore user, but in this case the user would have the sum of privileges of the three users. This would result in a very powerful user and is therefore not recommended for production use.

However, for a quick and simple temporary trial, it may be easiest to set the users for all accounts using the same user name and password, such as root.

## The Backing Store (Fabric) User

The first thing you must have is a user account on the MySQL server that you plan to use for your backing store. The user account information is stored in the Fabric configuration file.

The backing store database and its associated user are defined under the **[storage]** using *user* for the user name and *password* as the password.

The Fabric user account on the backing store requires the following privileges on the backing store database:

```
ALTER          - alter some database objects
CREATE         - create most database objects
CREATE VIEW    - create views
DELETE        - delete rows
DROP          - drop most database objects
EVENT         - manage events
REFERENCES    - foreign keys
INDEX         - create indexes
INSERT       - insert rows
```

```
SELECT      - select rows
UPDATE     - update rows
```

Example statements to create this user, to be executed on the backing store MySQL server:



**Note**

MySQL Fabric creates this database based on `fabric.cfg`, which in our example is named `mysql_fabric`. In other words, do not execute `CREATE database mysql_fabric;` here.

```
CREATE USER 'fabric_store'@'localhost'
  IDENTIFIED BY 'secret';

GRANT ALTER, CREATE, CREATE VIEW, DELETE, DROP, EVENT,
  INDEX, INSERT, REFERENCES, SELECT, UPDATE ON mysql_fabric.*
  TO 'fabric_store'@'localhost';
```



**Note**

The "REFERENCES" privilege is only required when working with MySQL 5.7 and above. MySQL Fabric does not check for this privilege on earlier versions.

For additional information about using and setting up the backing store, see [Chapter 6, Backing Store](#).

## The Server User

MySQL Fabric uses the server user account to access all MySQL servers that it manages. In other words, this user must be created on all managed MySQL servers.

The server account is defined under the `[servers]` section using `user` for the user name and `password` as the password.

The Fabric server user account on the managed MySQL servers requires the following privileges in global scope:

```
DELETE      - prune_shard
PROCESS    - list sessions to kill
RELOAD     - RESET SLAVE
REPLICATION CLIENT - SHOW SLAVE STATUS
REPLICATION SLAVE - SHOW SLAVE HOSTS
```

The Fabric server user account on the managed MySQL servers requires the following privileges on `mysql_fabric.*`:

```
ALTER      - alter some database objects
CREATE     - create most database objects
DELETE     - delete rows
DROP      - drop most database objects
INSERT    - insert rows
SELECT    - select rows
UPDATE    - update rows
```

Example statements to create the server user, to be executed on each managed MySQL server:

```
CREATE USER 'fabric_server'@'localhost'
  IDENTIFIED BY 'secret';
```

```
GRANT DELETE, PROCESS, RELOAD, REPLICATION CLIENT,  
REPLICATION SLAVE, SELECT, SUPER, TRIGGER ON *.*  
TO 'fabric_server'@'localhost';  
  
GRANT ALTER, CREATE, DELETE, DROP, INSERT, SELECT, UPDATE  
ON mysql_fabric.* TO 'fabric_server'@'localhost';
```

### The Backup User

If you want to use sharding, or clone a MySQL server with the intention to add it to a High-Availability (HA) group, then you must define restore and backup users. Like the server user, these users must be created on all managed servers.

The backup account is defined under the **[servers]** section using *backup\_user* for the user name and *backup\_password* as the password.

The backup account on the managed MySQL servers requires the following privileges in global scope if *mysqldump* is used:

```
EVENT          - show event information  
EXECUTE        - show routine information inside views  
REFERENCES     - foreign keys  
SELECT         - read data  
SHOW VIEW     - SHOW CREATE VIEW  
TRIGGER        - show trigger information
```

Example statements to create the backup user, to be executed on each managed MySQL server:

```
CREATE USER 'fabric_backup'@'localhost'  
IDENTIFIED BY 'secret';  
  
GRANT EVENT, EXECUTE, REFERENCES, SELECT, SHOW VIEW, TRIGGER ON *.*  
TO 'fabric_backup'@'localhost';
```



#### Note

The "REFERENCES" privilege is only required when working with MySQL 5.7 and above. MySQL Fabric does not check for this privilege on earlier versions.

### The Restore User

If you want to use sharding, or clone a server with the intention to add it to a High-Availability (HA) group, then you must define restore and backup users. Like the server user, these users must be created on all managed servers.

The restore account is defined under the **[servers]** section using *restore\_user* for the user name and *restore\_password* as the password.

The restore account on the managed MySQL servers requires the following privileges in global scope if *mysqldump* is used:

```
ALTER          - ALTER DATABASE  
ALTER ROUTINE  - ALTER {PROCEDURE|FUNCTION}  
CREATE         - CREATE TABLE  
CREATE ROUTINE - CREATE {PROCEDURE|FUNCTION}  
CREATE TABLESPACE - CREATE TABLESPACE
```

```
CREATE VIEW      - CREATE VIEW
DROP            - DROP TABLE (used before CREATE TABLE)
EVENT          - DROP/CREATE EVENT
INSERT         - write data
LOCK TABLES   - LOCK TABLES (--single-transaction)
REFERENCES     - Create tables with foreign keys
SELECT        - LOCK TABLES (--single-transaction)
SUPER         - SET @@SESSION.SQL_LOG_BIN = 0
TRIGGER        - CREATE TRIGGER
```



**Note**

Although the "CREATE TABLESPACE" and "REFERENCES" privileges are only required when working with MySQL 5.7 and above, MySQL Fabric still checks for them to help simplify the upgrade process to MySQL 5.7.

Example statements to create the restore user, to be executed on each managed MySQL server:

```
CREATE USER 'fabric_restore'@'localhost'
  IDENTIFIED BY 'secret';

GRANT ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TABLESPACE, CREATE VIEW,
  DROP, EVENT, INSERT, LOCK TABLES, REFERENCES, SELECT, SUPER,
  TRIGGER ON *.* TO 'fabric_restore'@'localhost';
```

## 2.3.2 Configuration File

The location of the MySQL Fabric configuration file varies depending on the operating system it is installed on and how you installed it. The table below lists the default configuration file locations for pre-built packages from <http://dev.mysql.com/downloads/utilities/>. Alternatively, the optional `--config` option accepts to use a path to a Fabric configuration file, and if defined, the file is loaded and used instead of the default configuration file location.

**Table 2.1 Default MySQL Fabric configuration file location**

Platform	Package	Location
Microsoft Windows	mysql-utilities-1.5.6-win32.msi	<code>UTILITIES_INSTALLDIR/etc/mysql/fabric.cfg</code>
Ubuntu Linux 14.04	mysql-utilities_1.5.6-1ubuntu14.04_all.deb	<code>/etc/mysql/fabric.cfg</code>
Debian Linux 6.0	mysql-utilities_1.5.6-1debian6.0_all.deb	<code>/etc/mysql/fabric.cfg</code>
Red Hat Enterprise Linux 6 / Oracle Linux 6	mysql-utilities-1.5.6-1.el6.noarch.rpm	<code>/etc/mysql/fabric.cfg</code>
OS X	mysql-utilities-1.5.6-osx10.9.dmg	<code>/etc/mysql/fabric.cfg</code>

Modify the configuration file and include the users and passwords defined in the previous step (Section 2.3.1, "Create the Associated MySQL Users"), here is an example Fabric configuration file:

```
[DEFAULT]
prefix = /usr/local
sysconfdir = /usr/local/etc
logdir = /var/log

[storage]
```

## Configuration File

---

```
address = localhost:3306
user = fabric_store
password = secret
database = mysql_fabric
auth_plugin = mysql_native_password
connection_timeout = 6
connection_attempts = 6
connection_delay = 1

[servers]
user = fabric_server
password = secret
backup_user = fabric_backup
backup_password = secret
restore_user = fabric_restore
restore_password = secret
unreachable_timeout = 5

[protocol.xmlrpc]
address = localhost:32274
threads = 5
user = admin
password = secret
disable_authentication = no
realm = MySQL Fabric
ssl_ca =
ssl_cert =
ssl_key =

[protocol.mysql]
address = localhost:32275
user = admin
password = secret
disable_authentication = no
ssl_ca =
ssl_cert =
ssl_key =

[executor]
executors = 5

[logging]
level = INFO
url = file:///var/log/fabric.log

[sharding]
mysqldump_program = /usr/bin/mysqldump
mysqlclient_program = /usr/bin/mysql

[statistics]
prune_time = 3600

[failure_tracking]
notifications = 300
notification_clients = 50
notification_interval = 60
failover_interval = 0
detections = 3
detection_interval = 6
detection_timeout = 1
prune_time = 3600

[connector]
ttl = 1
```

Each section has one or more variables defined that provide key information to the MySQL Fabric system libraries. You might not have to change any of these variables other than the users and passwords. For

more information on the sections and variables in the configuration file, see [Section 2.3.3, “Configuration File Sections”](#).

## 2.3.3 Configuration File Sections

The MySQL Fabric configuration file contains all the information necessary to run the MySQL Fabric utility. In addition, it serves as a configuration file for the utilities from within MySQL Fabric.

Each section has one or more variables defined that provide key information to the MySQL Fabric system libraries.



### Note

The **[client]** section was removed in MySQL Fabric 1.5.5. Instead, use the *restore\_user*, *restore\_password*, *backup\_user*, and *backup\_password* under the **[servers]** section to configure users for the backup and restore utilities, such as `mysqldump` and the `mysql` client.

### 2.3.3.1 Section DEFAULT

The `DEFAULT` section contains information on the installation paths for MySQL Fabric. This section is generated as part of the installation and should normally not be modified.

<code>prefix</code>	The installation prefix used when installing the <code>mysql.fabric</code> package and the binaries.
<code>sysconfdir</code>	The location of the system configuration files. Normally located in the <code>etc</code> directory under the directory given in <code>prefix</code> , but in some situations this might be different.
<code>logdir</code>	Configures the directory where log files are located by default. Normally, the <code>logging URL</code> contains the absolute path, but in the event that the path is relative, it is relative to this directory.

### 2.3.3.2 Section storage

This section contains information that the MySQL Fabric node uses for the connection to the backing store. For more information on the backing store, see [Chapter 6, Backing Store](#).

<code>address</code>	This is the address of the backing store in the form <code>host:port</code> . The port is optional and if not provided, defaults to 3306.
<code>user</code>	User name to use when connecting to the backing store.
<code>password</code>	The password to use when connecting to the backing store. If no password option is in the configuration file, a password is required on the terminal when the MySQL Fabric node is started. Although it is possible to set an empty password by not providing a value to the option, it is not recommended.
<code>database</code>	The database where the tables containing information about the MySQL Fabric farm is stored, typically <code>fabric</code> .
<code>auth_plugin</code>	The authentication plugin to use when connecting to the backing store. This option is passed to the connector when connecting to the backing store. For more information on authentication plugins, see <a href="#">Connector/Python Connection Arguments</a> .

<code>connection_timeout</code>	Timeout for the connection to the backing store, in seconds. This option is passed to the connector when connecting to the backing store. This is the maximum amount of time that MySQL Fabric waits for access to the backing store to complete before aborting. For more information on authentication plugins, see <a href="#">Connector/Python Connection Arguments</a> .
<code>connection_attempts</code>	The number of attempts to reconnect to the backing store before giving up. This is the maximum number of times MySQL Fabric attempts to create a connection to the backing store before aborting. The default is 0 retries.
<code>connection_delay</code>	The delay between attempts to connect to the backing store in seconds. The default is 0 seconds.

### 2.3.3.3 Section servers

This section contains information that MySQL Fabric uses to connect to the servers being managed.



#### Note

The `backup_user`, `backup_password`, `restore_user`, and `restore_password` options were added in MySQL Fabric 1.5.5.

<code>user</code>	User name to use when connecting to the managed server.
<code>password</code>	Password to use when connecting to the managed servers.
<code>backup_user</code>	User name to use when backing up the MySQL server.
<code>backup_password</code>	Password to use when backing up the MySQL servers with the <b>backup_user</b> user.
<code>restore_user</code>	User name to use when restoring the MySQL server.
<code>restore_password</code>	Password to use when restoring the MySQL servers with the <b>restore_user</b> user.
<code>unreachable_timeout</code>	Used for the connection timeout when checking faulty servers, or servers that are new to the farm. Hence, for servers that can potentially be unreachable. Defaults to 5, can be a value between 1-60.

### 2.3.3.4 Section protocol.xmlrpc

This section contains information about how the client connects to a MySQL Fabric node and configuration parameters for the XML-RPC protocol on the server.

<code>address</code>	Host and port of XML-RPC server. The host is only used by the client when connecting to the MySQL Fabric node, but the port is used by the server when starting the protocol server and by the client when reading how to connect to the XML-RPC server. The port number is typically 32274, and the host is typically <code>localhost</code> .
<code>threads</code>	Number of threads that the XML-RPC server uses for processing requests. This determines the number of concurrent requests that MySQL Fabric accepts.
<code>user</code>	User that the client uses to connect to the XML-RPC server.



password	Password used when the client connects to the server. If no password is provided, the client requests a password on the command-line.
<code>disable_authentication</code>	Whether to disable authentication or not. Disabling authentication can be useful when experimenting in a closed environment, it is <i>not</i> recommended for normal usage. Alternatives are <code>yes</code> or <code>no</code> and are case-insensitive.
realm	The realm (as defined in RFC 2617) the XML-RPC server identifies as when authenticating.
ssl_ca	Path to a file containing a list of trusted SSL certification authorities (CAs).
ssl_cert	The name of the SSL certificate file to use for establishing a secure connection.
ssl_key	The name of the SSL key file to use for establishing a secure connection.

### 2.3.3.5 Section `protocol.mysql`

This section contains information about how the client connects to a MySQL Fabric node using the MySQL Client/Server protocol.

address	Host and port of a MySQL Fabric node. The port number is typically 32275, and the host is typically <code>localhost</code> .
user	User that the client uses to connect to the MySQL Fabric node.
password	Password used when the client connects to the MySQL Fabric node. If no password is provided, the client requests a password on the command-line.
<code>disable_authentication</code>	Whether to disable authentication or not. Disabling authentication can be useful when experimenting in a closed environment, it is <i>not</i> recommended for normal usage. Alternatives are <code>yes</code> or <code>no</code> and are case-insensitive.
ssl_ca	Path to a file containing a list of trusted SSL certification authorities (CAs).
ssl_cert	The name of the SSL certificate file to use for establishing a secure connection.
ssl_key	The name of the SSL key file to use for establishing a secure connection.

### 2.3.3.6 Section `executor`

This section contains parameters to configure the executor. The executor executes procedures in a serial order, which guarantees that requests do not conflict. The requests received are mapped to procedures which can be executed directly or scheduled through the executor. Procedures scheduled through the executor are processed within the context of threads spawned by the executor. Usually, read operations are immediately executed by the XML-RPC session thread and write operations are scheduled and executed through the executor.

executors	The number of executor threads that the executor uses when processing requests.
working_directory	The directory Fabric uses by default to store files. If the option is not found, the working directory is the directory from where the process was launched.



**Note**

This option was added in Fabric 1.5.7.

### 2.3.3.7 Section logging

MySQL Fabric logs information about its activities to the standard output when started as a regular process. However, when started as a daemon, it writes information to a file configured by the the option [Fabric URL used for logging](#).

level	The log level to use when generating the log. Acceptable values are <a href="#">CRITICAL</a> , <a href="#">ERROR</a> , <a href="#">WARNING</a> , <a href="#">INFO</a> , and <a href="#">DEBUG</a> . The default is <a href="#">INFO</a> .
url	The URL to use for logging. Supported protocols are currently <a href="#">file</a> and <a href="#">syslog</a> . The <a href="#">file</a> protocol creates a rotating file handler, while the <a href="#">syslog</a> protocol logs messages using the system logger <a href="#">syslogd</a> .

The [file](#) handler accepts either a relative path or an absolute path. If a relative path is provided, it is relative to [Configure default log directory](#).

The [syslog](#) handler accepts either a path (for example [syslog:///dev/log](#)) or a hostname and an optional port (for example, [syslog://localhost:555](#), and [syslog://my.example.com](#)). If no port is provided, it defaults to 541, which is the default port for the syslog daemon.

### 2.3.3.8 Section sharding

To perform operations such as moving and splitting shards, MySQL Fabric relies on the [mysqldump](#) and [mysql](#) client programs. These programs can be installed in different locations and if they are not in the path for the MySQL Fabric node, this section configures where they can be found.

mysqldump_program	Path to the <a href="#">mysqldump</a> program.
mysqlclient_program	Path to the <a href="#">mysql</a> program.

### 2.3.3.9 Section statistics

Connectors and other external entities log any errors while accessing servers so that MySQL Fabric can monitor server health and act accordingly. For example, MySQL Fabric promotes a new master after receiving [notifications](#) from the number of clients configured in [notification\\_clients](#)) within the time interval configured in [notification\\_interval](#). If a server is considered unstable but it is not a master, it is simply marked as faulty. To avoid making the system unstable, a new master can only be automatically promoted after the [failover\\_interval](#) has been elapsed since the last promotion. In order to ease the adoption of MySQL Fabric, a built-in failure detector is provided. If the failure detector is enabled to monitor a group, a new master is promoted after 3 failed successive attempts to access the current master within the time interval configured in [failover\\_interval](#). The failure detection routine tries to connect to servers in a group and uses the value configured in [detection\\_timeout](#) as timeout.

prune_time	How often the internal event log is pruned, in seconds and also the age of events in the event log that is used to present statistics.
notifications	Number of issues before the server is considered unstable.
notification_clients	Number of different sources that should report issues on a server before it is considered unstable.
notification_interval	Amount of time in seconds that is used when deciding if a server is unstable. Issues older than this are not considered when deciding.
failover_interval	Minimum time in seconds between subsequent automatic promotions. This parameter is used to prevent the system entering a sequence of promotions that could disable the system.
detections	Number of successive failed attempts to contact the server after which the built-in failure detector considers the server unstable.
detection_timeout	Timeout in seconds used when contacting the server. If the server does not respond within this time period, it is recorded as a failed attempt to contact the server.
prune_time	Maximum age in seconds for reported issues in the error log. Issues older than this are removed from the error log.

### 2.3.3.10 Section failure\_tracking

This section contains parameters for the failure management system.

notifications	The notification threshold. If more than this number of notifications arrive in the notification interval and the number of distinct notification clients are over the notification client threshold, the server is considered dead and failover is triggered.
notification_clients	The number of distinct notification clients that need to report. If more than this number of distinct notification clients are over the notification client threshold and the number of notifications above the notification threshold arrive in the notification interval, the server is considered dead and failover is triggered.
notification_interval	The notification interval in seconds. Only notifications arriving within this time frame can trigger a failover.
failover_interval	The minimum interval between failover operations in seconds. In order to avoid making the system unstable, failover operations are not triggered unless at least this much time has expired since the last failover.
detections	This parameter is for the built-in failure detector. If more than this number of failures to contact the server occurs during the detection interval, the server is considered unstable and a failover is triggered.
detection_interval	This parameter configures the detection interval for the built-in failure detector, in seconds.
detection_timeout	This parameter configures the detection timeout used when attempting to contact the servers in the group.

`prune_time` This is the maximum age of events in the failure detector's error log and is also the interval for how often the error log is pruned.

### 2.3.3.11 Section connector

Connectors that are MySQL Fabric-aware contact MySQL Fabric to fetch information on groups, shards, and servers, and then cache the results locally for a time period to improve performance. This section contains configuration parameters passed to the connectors.

`ttl` The Time To Live (TTL), measured in seconds, is passed together with other information to the connector. This is used by the connector to invalidate the caches, and reload them from a MySQL Fabric node, after the TTL has expired.

## 2.3.4 The Configuration Parameter (--param)

The `--param` option allows you to override configuration options at runtime. The syntax is `--param=section.option=value`. For example:

```
shell> mysqlfabric manage setup --param=storage.user=fabric_store --param=storage.password=secret
shell> mysqlfabric --param=storage.user=fabric_store --param=storage.password=secret manage setup

shell> mysqlfabric manage setup --param=storage.address=localhost:13000 \
--param=storage.user=root --param=protocol.xmlrpc.password=secret
```

For additional information about the available options, see [Section 2.3.2, “Configuration File”](#).

## 2.4 Starting and Stopping MySQL Fabric Nodes

To start or stop MySQL Fabric nodes, use the `mysqlfabric` command (see [Chapter 4, \*The mysqlfabric Utility\*](#)). This command requires that MySQL Fabric and Connector/Python be installed, and assumes that you have set up the backing store.

The following command starts a MySQL Fabric node and should be run on one of the servers listed in the `[protocol.xmlrpc]` section in the configuration file.

```
shell> mysqlfabric manage start
```

This command starts a MySQL Fabric node on the machine where it is executed and prints the log to standard output. Thus, this is the machine where you installed the MySQL Fabric and Connector/Python software and is also the machine listed in the configuration file `[protocol.xmlrpc]` section. To follow the examples in the quick-start section, you must use `localhost` for the host name.

To put the MySQL Fabric node in the background, add the `--daemonize` option. However, this diverts the log to the syslog file. While experimenting with MySQL Fabric, you may find it more convenient not to use `--daemonize` so that the log is written to your terminal.

Use this command to stop a MySQL Fabric node:

```
shell> mysqlfabric manage stop
```

This command contacts the MySQL Fabric server running at the address mentioned in the `[protocol.xmlrpc]` section and instructs it to stop.

## 2.5 Old Configuration System



### Important

This documentation describes the MySQL Fabric configuration before version 1.5.5. The previous configuration documentation is archived here for informational and upgrade purposes.

Configuring MySQL Fabric requires creating a MySQL user to access the backing store, and editing the configuration file with the MySQL user details. This section assumes you have already set up the backing store. See [Chapter 6, Backing Store](#) for more information.

### Create a MySQL User

The first thing you must have is a user account on the MySQL server that you plan to use for your backing store. The user account information is stored in the configuration file.

The user account must have full privileges for the database named `fabric`. To create the user and grant the privileges needed, use the following statements:

```
CREATE USER 'fabric'@'localhost' IDENTIFIED BY 'secret';
GRANT ALL ON fabric.* TO 'fabric'@'localhost';
```

In the preceding example, substitute a password of your choice (replace `'secret'`). Also, if you are going to run MySQL Fabric on a host other than where the backing store resides, substitute the `'localhost'` for the host name.

MySQL Fabric uses the same user account, who must have all privileges on all databases, to access all MySQL servers that it manages. The user and password are defined in the configuration file as shown below. To create this user and grant all the necessary privileges, execute the following command on all MySQL servers:

```
CREATE USER 'fabric'@'localhost' IDENTIFIED BY 'secret';
GRANT ALL ON *.* TO 'fabric'@'localhost';
```

In the preceding example, substitute a password of your choice (replace `'secret'`). Also, if you are going to run MySQL Fabric on a host other than where the managed MySQL servers reside, substitute the `'localhost'` for the Fabric's host name.

### Configuration File

The next step is to modify the configuration file with the user and password we defined in the previous step. Open the configuration file:

### MySQL Fabric configuration file location

Table 2.2 MySQL Fabric configuration file location

Platform	Package	Location
Microsoft Windows	mysql-utilities-1.5.6-win32.msi	<code>UTILITIES_INSTALLDIR/etc/mysql/fabric.cfg</code>
Ubuntu Linux 14.04	mysql-utilities_1.5.6-1ubuntu14.04_all.deb	<code>/etc/mysql/fabric.cfg</code>
Debian Linux 6.0	mysql-utilities_1.5.6-1debian6.0_all.deb	<code>/etc/mysql/fabric.cfg</code>

Platform	Package	Location
Red Hat Enterprise Linux 6 / Oracle Linux 6	mysql- utilities-1.5.6-1.el6.noarch.rpm	/etc/mysql/fabric.cfg

The following shows the content of the configuration file and the modifications necessary. In the [\[storage\]](#) section, store the user and password of the user created in the previous step.

```
[DEFAULT]
prefix = /usr/local
sysconfdir = /usr/local/etc
logdir = /var/log

[storage]
address = localhost:3306
user = fabric
password = secret
database = fabric
auth_plugin = mysql_native_password
connection_timeout = 6
connection_attempts = 6
connection_delay = 1

[servers]
user = fabric
password =
backup_user = fabric
backup_password =
restore_user = fabric
restore_password =
unreachable_timeout = 5

[protocol.xmlrpc]
address = localhost:32274
threads = 5
user = admin
password =
disable_authentication = no
realm = MySQL Fabric
ssl_ca =
ssl_cert =
ssl_key =

[protocol.mysql]
address = localhost:32275
user = admin
password =
disable_authentication = no
ssl_ca =
ssl_cert =
ssl_key =

[executor]
executors = 5

[logging]
level = INFO
url = file:///var/log/fabric.log

[sharding]
mysqldump_program = /usr/bin/mysqldump
mysqlclient_program = /usr/bin/mysql

[statistics]
prune_time = 3600
```

```
[failure_tracking]
notifications = 300
notification_clients = 50
notification_interval = 60
failover_interval = 0
detections = 3
detection_interval = 6
detection_timeout = 1
prune_time = 3600

[connector]
ttl = 1
```

Each section has one or more variables defined that provide key information to the MySQL Fabric system libraries. You should not have to change any of these variables other than the user and password for the backing store (in the [storage](#) section).





---

## Chapter 3 Quick Start

### Table of Contents

3.1 Example: Fabric and Replication .....	19
3.2 Example: Fabric and Sharding .....	23
3.2.1 Introduction to Sharding .....	23
3.2.2 Sharding Scenario .....	24

This section demonstrates how to get started using MySQL Fabric. Two examples are included in this section: one for using MySQL Fabric with replication to demonstrate how Fabric reduces the overhead of directing reads and writes from applications, and another showing how Fabric makes using sharding much easier.

If you have not installed and configured Fabric, please refer to the previous sections before proceeding with the examples.

### 3.1 Example: Fabric and Replication

This section presents a quick start for using MySQL replication features in Fabric. To run this example, you should have four server instances (running MySQL version 5.6.10 or later). The commands in this example are executed on the same server host as the backing store (which happens to be the same host where Fabric was installed). You must also have a Fabric node started on that host.

The replication features in Fabric focus on providing high availability. While these features continue to evolve, the most unique feature of Fabric replication is the ability to use a Fabric-aware connector to seamlessly direct reads and writes to the appropriate servers.

This redirection is achieved through the use of one of the central concepts in Fabric: a high-availability group that uses a high-availability solution for providing resilience to failures. Currently, only asynchronous primary backup replication is supported. As long as the primary is alive and running, it handles all write operations. Secondaries replicate everything from the primary to stay up to date and might be used to scale out read operations.

### Creating a High-Availability Group

The first step consists of creating a group, here designated `my_group`. After doing so, you can add servers to the group. In this case, we add three servers, `localhost:3307`, `localhost:3308`, and `localhost:3309`.

Fabric accesses each added server using the user and password provided in the configuration file to guarantee that they are alive and accessible. If these requirements are not fulfilled, the operation fails and the server is not added to the group.



#### Note

Each managed MySQL Server has the following requirements: `gtid_mode` (GTID), `bin_log` (binary logging), and `log_slave_updates` enabled, with `server_id` properly configured.

The following demonstrates the commands to execute these steps.

```
shell> mysqlfabric group create my_group
Procedure :
{ uuid      = d4e60ed4-fd36-4df6-8004-d034202c3698,
  finished  = True,
  success   = True,
  return    = True,
  activities =
}
shell> mysqlfabric group add my_group localhost:3307
Procedure :
{ uuid      = 6a33ed29-ccf8-437f-b436-daf07db7a1fc,
  finished  = True,
  success   = True,
  return    = True,
  activities =
}
shell> mysqlfabric group add my_group localhost:3308
Procedure :
{ uuid      = 6892bc49-3ab7-4bc2-891d-57c4a1577081,
  finished  = True,
  success   = True,
  return    = True,
  activities =
}
shell> mysqlfabric group add my_group localhost:3309
Procedure :
{ uuid      = 7943b27f-2da5-4dcf-ala4-24aed8066bb4,
  finished  = True,
  success   = True,
  return    = True,
  activities =
}
```

To show information about the set of servers in a group, use this command:

```
shell> mysqlfabric group lookup_servers my_group
```

To get detailed information about the group health, use this command:

```
shell> mysqlfabric group health my_group
```

## Promoting and Demoting Servers

After executing the steps in setting up a high-availability group, Fabric does not become aware of any replication topology that was already in place. It is necessary to promote one of them to primary (that is, master) and make the remaining servers secondaries (that is, slaves). To do so, execute this command:

```
shell> mysqlfabric group promote my_group
```

If there is a primary in a group, any server added subsequently is automatically set as secondary. Setting a different server as primary can be done through the same command, which demotes the current primary and elects a new one. If the current primary has failed, the same command (which can be triggered either manually or automatically) can be used to elect a new one.



### Note

A server marked as "faulty" cannot be promoted to a secondary or primary status in one step. The faulty server status must first be changed to the "spare" status. For example, use `mysqlfabric server set_status server-address spare`.

## Activating or Deactivating a Failure Detector

If the primary fails, you may want to automatically promote one of the secondaries as primary and redirect the remaining secondaries to the new primary. To do this, execute the following command:

```
shell> mysqlfabric group activate my_group
```

If the failure detector discovers that a primary has crashed, it marks it as faulty and triggers a failover routine. This is not done automatically because there may be users who prefer to use an external failure detector or want to do things manually. To deactivate the failure detector, execute the following command:

```
shell> mysqlfabric group deactivate my_group
```

## Executing Updates and Queries

Executing queries with a Fabric-aware connector is easy. The following example shows a section of code designed to add employees and search for them. Notice that we simply import the fabric package from the Connector/Python library and provide the Fabric connection parameters such as the location of the Fabric node (as specified in the `[protocol.xmlrpc]` configuration file section) and user credentials.

```
import mysql.connector
from mysql.connector import fabric

def add_employee(conn, emp_no, first_name, last_name):
    conn.set_property(group="my_group", mode=fabric.MODE_READWRITE)
    cur = conn.cursor()
    cur.execute("USE employees")
    cur.execute(
        "INSERT INTO employees VALUES (%s, %s, %s)",
        (emp_no, first_name, last_name)
    )
    # We need to keep track of what we have executed in order to,
    # at least, read our own updates from a slave.
    cur.execute("SELECT @@global.gtid_executed")
    for row in cur:
        print "Transactions executed on the master", row[0]
    return row[0]

def find_employee(conn, emp_no, gtid_executed):
    conn.set_property(group="my_group", mode=fabric.MODE_READONLY)
    cur = conn.cursor()
    # Guarantee that a slave has applied our own updates before
    # reading anything.
    cur.execute(
        "SELECT WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS('%s', 0)" %
        (gtid_executed, )
    )
    for row in cur:
        print "Had to synchronize", row, "transactions."
    cur.execute("USE employees")
    cur.execute(
        "SELECT first_name, last_name FROM employees "
        "WHERE emp_no = %s", (emp_no, )
    )
    for row in cur:
        print "Retrieved", row

# Address of the Fabric, not the host we are going to connect to.
conn = mysql.connector.connect(
    fabric={"host" : "localhost", "port" : 32274,
```

```
        "username": "admin", "password" : "adminpass"
    },
    user="webuser", password="webpass", autocommit=True
)

conn.set_property(group="my_group", mode=fabric.MODE_READWRITE)
cur = conn.cursor()
cur.execute("CREATE DATABASE IF NOT EXISTS employees")
cur.execute("USE employees")
cur.execute("DROP TABLE IF EXISTS employees")
cur.execute(
    "CREATE TABLE employees ( "
    "  emp_no INT, "
    "  first_name CHAR(40), "
    "  last_name CHAR(40)"
    ")"
)

gtid_executed = add_employee(conn, 12, "John", "Doe")
find_employee(conn, 12, gtid_executed)
```

You can copy this code into a Python module named `test_fabric_query.py` and execute it with the following command:

```
shell> python ./test_fabric_query.py
(u'John', u'Doe')
```

## Group Maintenance

To find out which servers are in a group, use this command:

```
shell> mysqlfabric group lookup_servers my_group
Command :
{ success      = True
  return      = [
    {'status': 'PRIMARY', 'server_uuid': 'bbe6f7c1-b6c3-11e3-aaa2-58946b051f64',
      'mode': 'READ_WRITE', 'weight': 1.0, 'address': 'localhost:3307'
    },
    {'status': 'SECONDARY', 'server_uuid': '0c9e67d0-8194-11e2-8a7c-f0def124dcc5',
      'mode': 'READ_ONLY', 'weight': 1.0, 'address': 'localhost:3308'
    },
    {'status': 'SECONDARY', 'server_uuid': '0c67e5b1-8194-11e2-8a7c-f0def124dcc5',
      'mode': 'READ_ONLY', 'weight': 1.0, 'address': 'localhost:3309'
    },
  ],
  activities =
}
```

In this example, there are three servers identified by their UUID values. The server running at `localhost:3307` is a primary, whereas the other servers are secondaries.

It is sometimes necessary to take secondaries offline for maintenance. However, before stopping a server, it is important to remove it from the group to avoid having the Fabric failure detector trigger any action. This can be done through the following commands. `server_uuid` should be replaced with a server UUID value (a value of the form `d2369bc9-2009-11e3-93c6-f0def14a00f4`).

```
shell> mysqlfabric group remove my_group server_uuid
```

A primary cannot be removed from a group. To disable any primary in a group, execute this command:

```
shell> mysqlfabric group demote my_group
```

If a group contains no servers, it is empty and can be destroyed (removed) with this command:

```
shell> mysqlfabric group destroy my_group
```

It is also possible to force removal of a nonempty group by specifying the parameter `--force`. This command removes all servers from `my_group` and removes the group.

```
shell> mysqlfabric group destroy my_group --force
```

## 3.2 Example: Fabric and Sharding

This example explores sharding. The essence of a sharding solution that uses MySQL involves partitioning the data into independent sets (independent MySQL Servers) and directing each client to the partition (MySQL Server) that has the data the client wants to modify.

This architecture scales the write operations for a given dataset since resource demands are distributed across the partitions (MySQL Servers) of the data set. Each partition is referred to as a shard.

### 3.2.1 Introduction to Sharding

The Fabric sharding implementation requires you to provide the sharding key explicitly while executing a query. To define sharding over a set of tables using the sharding mechanism built into Fabric, it is important to understand two concepts and how they relate.

#### Shard Mapping

A shard mapping serves to bring a database object (a database table) into the Fabric sharding system. The mapping is a way of informing Fabric that we want a particular sharding scheme (range, hash, list, and so forth.) to be used on a database table, using a value in a particular column. Create a shard mapping as follows:

1. Define a shard mapping, to tell Fabric the kind of sharding mechanism to use.
2. Add a relation between the mapping and a database object, to register the database table and a column in the table with the shard mapping.

Once these operations have been completed, we can describe how the shard mapping should split the tables. This is done while creating the shards.

#### Shards

These are the partitions on the table. Since sharding is done on a database table, using an attribute (column) in the table, the values in the column influence how the shards are created. To explain this further, assume that we have two tables we wish to shard.

- employees
- Salary

Assume further that both tables are to be sharded by the employee ID that is part of their columns. Where a row is placed is based on the value in the employee ID column. Hence, in a range-based sharding scheme, a shard is nothing but a range of employee ID values.

## 3.2.2 Sharding Scenario

In the sections that follow, we take the example of a `employees` table that must be sharded across three MySQL groups. The following procedure lists the sequence of commands to run to perform each step.

### Unsharded Data

Assume that we have an unsharded table named `employees` that contains Employee IDs, on which we want to create the following shards. Each of these ranges is placed in one shard:

- `1-99999`: shard-1
- `100000-199999`: shard-2
- `200000-`: shard-3

In addition to creating the ranges themselves (in a range based sharding scheme) we also must define where this range of values should reside.

### Starting Multiple MySQL Servers

MySQL Servers must be started on the directories that were copied. Assume that MySQL servers are started on the following hosts and ports:

- `localhost:3307`
- `localhost:3308`
- `localhost:3309`
- `localhost:3310`
- `localhost:3311`
- `localhost:3312`
- `localhost:3313`
- `localhost:3314`

### Creating the Groups in Fabric

A logical group in Fabric maps to a shard. So as a first step to sharding we must implement the Fabric groups that store the shards. This can be done as follows:

```
shell> mysqlfabric group create group_id-1
shell> mysqlfabric group create group_id-2
shell> mysqlfabric group create group_id-3
```

The preceding commands create three high-availability groups: `group_id-1`, `group_id-2`, and `group_id-3`. Each group stores a shard.

Then we must define a global group which is used to propagate schema updates to all tables in the sharding setup and updates to global tables throughout the sharding scheme.

```
shell> mysqlfabric group create group_id-global
```

## Registering the Servers

The MySQL servers must be registered with the groups. Each group contains two servers.

*3307, 3308 belong to group\_id-1*

```
shell> mysqlfabric group add group_id-1 localhost:3307
shell> mysqlfabric group add group_id-1 localhost:3308
```

*3309, 3310 belong to group\_id-2*

```
shell> mysqlfabric group add group_id-2 localhost:3309
shell> mysqlfabric group add group_id-2 localhost:3310
```

*3311, 3312 belong to group\_id-3*

```
shell> mysqlfabric group add group_id-3 localhost:3311
shell> mysqlfabric group add group_id-3 localhost:3312
```

*3313, 3314 belong to group\_id-global*

```
shell> mysqlfabric group add group_id-global localhost:3313
shell> mysqlfabric group add group_id-global localhost:3314
```

Then promote one server to master in each group:

```
shell> mysqlfabric group promote group_id-global
shell> mysqlfabric group promote group_id-1
shell> mysqlfabric group promote group_id-2
shell> mysqlfabric group promote group_id-3
```

## Define a Shard Mapping

When we define a shard mapping, we basically do three things:

1. Define the type of sharding scheme we want to use (RANGE or HASH).
2. Define a global group that stores all the updates that must be propagated to all the shards that are part of this sharding scheme.
3. We generate a unique shard mapping id to which we can later associate database objects (tables).

```
shell> mysqlfabric sharding create_definition RANGE group_id-global
Procedure:
{ uuid      = 195bca1e-c552-464b-b4e3-1fa15e9b49d5,
  finished  = True,
  success   = True,
  return    = 1,
  activities =
}
```

## Add Tables to Defined Shard Mappings

Add the database table to the shard mapping defined previously.

```
shell> mysqlfabric sharding add_table 1 employees.employees emp_no
```

## Add Shards for the Registered Tables

```
shell> mysqlfabric sharding add_shard 1 "group_id-1/1, group_id-2/100000, group_id-2/200000" --state=ENABLED
```

## Executing Updates and Queries

Now you can write some example code for querying the sharded system.

```
import random
import mysql.connector
from mysql.connector import fabric

def prepare_synchronization(cur):
    # We need to keep track of what we have executed so far to guarantee
    # that the employees.employees table exists at all shards.
    gtid_executed = None
    cur.execute("SELECT @@global.gtid_executed")
    for row in cur:
        gtid_executed = row[0]
    return gtid_executed

def synchronize(cur, gtid_executed):
    # Guarantee that a slave has created the employees.employees table
    # before reading anything.
    cur.execute(
        "SELECT WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS('%s', 0)" %
        (gtid_executed, )
    )
    cur.fetchall()

def add_employee(conn, emp_no, first_name, last_name, gtid_executed):
    conn.set_property(tables=["employees.employees"], key=emp_no,
                      mode=fabric.MODE_READWRITE)
    cur = conn.cursor()
    synchronize(cur, gtid_executed)
    cur.execute("USE employees")
    cur.execute(
        "INSERT INTO employees VALUES (%s, %s, %s)",
        (emp_no, first_name, last_name)
    )

def find_employee(conn, emp_no, gtid_executed):
    conn.set_property(tables=["employees.employees"], key=emp_no,
                      mode=fabric.MODE_READONLY)
    cur = conn.cursor()
    synchronize(cur, gtid_executed)
    cur.execute("USE employees")
    for row in cur:
        print "Had to synchronize", row, "transactions."
    cur.execute(
        "SELECT first_name, last_name FROM employees "
        "WHERE emp_no = %s", (emp_no, )
    )
    for row in cur:
        print row

def pick_shard_key():
    shard = random.randint(0, 2)
    shard_range = shard * 100000
    shard_range = shard_range if shard != 0 else shard_range + 1
    shift_within_shard = random.randint(0, 99999)
    return shard_range + shift_within_shard

# Address of the Fabric, not the host we are going to connect to.
```



```
conn = mysql.connector.connect(
    fabric={"host" : "localhost", "port" : 32274,
           "username": "admin", "password" : "adminpass"
          },
    user="webuser", password="webpass", autocommit=True
)

conn.set_property(tables=["employees.employees"], scope=fabric.SCOPE_GLOBAL,
                  mode=fabric.MODE_READWRITE)
cur = conn.cursor()
cur.execute("CREATE DATABASE IF NOT EXISTS employees")
cur.execute("USE employees")
cur.execute("DROP TABLE IF EXISTS employees")
cur.execute(
    "CREATE TABLE employees (
      emp_no INT,
      first_name CHAR(40),
      last_name CHAR(40)
    )"
)
gtid_executed = prepare_synchronization(cur)

conn.set_property(scope=fabric.SCOPE_LOCAL)

first_names = ["John", "Buffalo", "Michael", "Kate", "Deep", "Genesis"]
last_names = ["Doe", "Bill", "Jackson", "Bush", "Purple"]

list_emp_no = []
for count in range(10):
    emp_no = pick_shard_key()
    list_emp_no.append(emp_no)
    add_employee(conn, emp_no,
                first_names[emp_no % len(first_names)],
                last_names[emp_no % len(last_names)],
                gtid_executed
    )

for emp_no in list_emp_no:
    find_employee(conn, emp_no, gtid_executed)
```

### Shard Move

If the current set of servers for a shard is not powerful enough, we can move the shard to a more powerful server set.

The shard-move functionality can be used to move a shard from one group to another. These are the steps necessary to move a shard.

1. Set up the required group or groups.

```
shell> mysqlfabric group create group_id-MOVE
shell> mysqlfabric group add group_id-MOVE localhost:3315
shell> mysqlfabric group add group_id-MOVE localhost:3316
shell> mysqlfabric group promote group_id-MOVE
```

2. Execute the shard move.

```
shell> mysqlfabric sharding move_shard 1 group_id-MOVE
```

3. Verify that the move has happened.

```
shell> mysqlfabric sharding lookup_servers employees.employees 4
```

## Shard Split

If the shard becomes overloaded, we may need to split the shard into another group. The `shard-split` feature can be used to split the data in a given shard into another shard. The following demonstrates how to do this.

1. *Set up the required group or groups.*

```
shell> mysqlfabric group create group_id-SPLIT
shell> mysqlfabric group add group_id-SPLIT localhost:3317
shell> mysqlfabric group add group_id-SPLIT localhost:3318
shell> mysqlfabric group promote group_id-SPLIT
```

2. *Execute the shard split.*

```
shell> mysqlfabric sharding split_shard 2 group_id-SPLIT --split_value=150
```

3. *Verify that the shard Split happened.*

```
shell> mysqlfabric sharding lookup_servers employees.employees 152
shell> mysqlfabric sharding lookup_servers employees.employees 103
```

---

# Chapter 4 The `mysqlfabric` Utility

## Table of Contents

4.1 Getting Help .....	29
4.2 Dump Commands .....	30
4.3 Event Commands .....	31
4.4 Group Commands .....	32
4.5 Manage Commands .....	36
4.6 Provider Commands .....	37
4.7 Role Commands .....	38
4.8 Server Commands .....	38
4.9 Sharding Commands .....	42
4.10 Snapshot Commands .....	46
4.11 Statistics Commands .....	47
4.12 Threat Commands .....	47
4.13 User Commands .....	48

This section describes the `mysqlfabric` utility including examples of most commands. For a quick reference guide for all of the commands, see [Chapter 5, \*Fabric Utility Command Matrix\*](#).

Fabric commands are organized in categories that include [dump](#), [event](#), [group](#), [manage](#), [provider](#), [role](#), [server](#), [sharding](#), [snapshot](#), [statistics](#) [threat](#), and [user](#).

## 4.1 Getting Help

- `mysqlfabric help`: Show syntax information and the help commands.
- `mysqlfabric help commands`: List the available commands and their description.
- `mysqlfabric help groups`: List the available groups.
- `mysqlfabric help [group] [command]`: Provide detailed information on a command.

```
shell> mysqlfabric help group create
group create group_id [--description=NONE] [--synchronous]

Create a group.
```

```
shell> mysqlfabric help
Usage: mysqlfabric [--param, --config] <grp> <cmd> [arg, ...].

MySQL Fabric 1.5.6 - MySQL server farm management framework

Options:
  --version          show program's version number and exit
  -h, --help        show this help message and exit
  --param=CONFIG_PARAMS
                    Override a configuration parameter.
  --config=FILE     Read configuration from FILE.

Basic commands:
  help <grp> <cmd> Show help for command
```

```
help commands    List all commands
help groups      List all groups
```

## 4.2 Dump Commands

The dump commands are designed to be used by the connectors to retrieve information on shards, high-availability groups and their servers.

- **Command:** `dump shard_maps`: Return information about all shard mappings matching any of the provided patterns. If no patterns are provided, dump information about all shard mappings.

```
Usage:
mysqlfabric dump shard_maps [--connector_version=CONNECTOR_VERSION]
                [--patterns=]

Options:
--connector_version=CONNECTOR_VERSION
                        The connectors version of the data. By default None.
--patterns=PATTERNS   shard mapping pattern. By default "".
```

- **Command:** `dump servers`: Return information about the index for all mappings matching any of the patterns provided. If no pattern is provided, dump the entire index. The lower\_bound that is returned is a string that is a md-5 hash of the group-id in which the data is stored.

```
Usage:
mysqlfabric dump servers [--connector_version=CONNECTOR_VERSION]
                [--patterns=]

Options:
--connector_version=CONNECTOR_VERSION
                        The connectors version of the data. By default None.
--patterns=PATTERNS   group pattern. By default "".
```

- **Command:** `dump shard_index`: Return information about servers. The servers might belong to any group that matches any of the provided patterns, or all servers if no patterns are provided.

```
Usage:
mysqlfabric dump shard_index [--connector_version=CONNECTOR_VERSION]
                [--patterns=]

Options:
--connector_version=CONNECTOR_VERSION
                        The connectors version of the data. By default None.
--patterns=PATTERNS   group pattern. By default "".
```

- **Command:** `dump sharding_information`: Return information about all shard mappings matching any of the provided patterns. If no patterns are provided, dump information about all shard mappings.

```
Usage:
mysqlfabric dump shard_maps [--connector_version=CONNECTOR_VERSION]
                [--patterns=]

Options:
--connector_version=CONNECTOR_VERSION
                        The connectors version of the data. By default None.
--patterns=PATTERNS   shard mapping pattern. By default "".
```

- **Command:** `dump shard_tables`: Return information about all tables belonging to mappings matching any of the provided patterns. If no patterns are provided, dump information about all tables.

```
Usage:
  mysqlfabric dump shard_tables [--connector_version=CONNECTOR_VERSION]
  [--patterns=]

Description:

Options:
  --connector_version=CONNECTOR_VERSION
                                The connectors version of the data. By default None.
  --patterns=PATTERNS          shard mapping pattern. By default "".
```

- **Command:** `dump fabric_nodes`: Return a list of Fabric Nodes (i.e. addresses). Specifically, the host and port are returned. If the protocol is not specified, it assumes the 'protocol.xmlrpc'. Currently, the 'protocol.xmlrpc' and 'protocol.mysql' are valid options.

```
Usage:
  mysqlfabric dump fabric_nodes [--protocol=PROTOCOL]

Options:
  --protocol=PROTOCOL          MySQL Fabric might support different protocols which
                                have different addresses. By default None.

Return:
  List with existing Fabric Servers. ["host:port", ...]
```

## 4.3 Event Commands

The event commands are used to define events for tailoring the Fabric system to your needs for controlling applications.

- **Command:** `event trigger`: Trigger an event.

```
Usage:
  mysqlfabric event trigger <event> [--locks=LOCKS] [--args=ARGS]
  [--kwargs=KWARGS]

Parameters:
  <event>  Event's identification. Accepted values: String

Options:
  --locks=LOCKS          By default None.
  --args=ARGS            Event's non-keyworded arguments. By default None.
  --kwargs=KWARGS       Event's keyworded arguments. By default None.

Return:
  :class:`CommandResult` instance with UUID of the procedures that were
  triggered.
```

- **Command:** `event wait_for_procedures`: Wait until procedures, which are identified through their uuid in a list and separated by comma, finish their execution. If a procedure is not found an error is returned.

```
Usage:
  mysqlfabric event wait_for_procedures [--proc_uuids=PROC_UUIDS]

Options:
  --proc_uuids=PROC_UUIDS
                                Iterable with procedures' UUIDs. By default None.
```

## 4.4 Group Commands

The group commands are used to define, modify, and control groups.

- **Command:** `activate`: Activate failure detector for server(s) in a group. By default the failure detector is disabled.

```
Usage:
  mysqlfabric group activate <group_id> [--synchronous]

Parameters:
  <group_id>  Group's id.

Options:
  --synchronous=SYNCHRONOUS
                Whether one should wait until the execution finishes or not. By
                default True.

Return:
  Tuple with job's uuid and status.
```

- **Command:** `add`: Add a server into group. If users just want to update the state store and skip provisioning steps such as configuring replication, the `update_only` parameter must be set to true. Note that the current implementation has a simple provisioning step that makes the server point to the master if there is any.

```
Usage:
  mysqlfabric group add <group_id> <address> [--timeout=TIMEOUT]
                                     [--update_only] [--synchronous]

Parameters:
  <group_id>  Group's id.
  <address>   Server's address.

Options:
  --timeout=TIMEOUT  Time in seconds after which an error is reported if one
                    cannot access the server. By default None.
  --update_only=UPDATE_ONLY
                    By default False.
  --synchronous=SYNCHRONOUS
                    Whether one should wait until the execution finishes or
                    not. By default True.

Return:
  Tuple with job's uuid and status.
```

- **Command:** `create`: Create a group.

```
Usage:
  mysqlfabric group create <group_id> [--description=DESCRIPTION]
                                     [--synchronous]

Parameters:
  <group_id>  Group's id.

Options:
  --description=DESCRIPTION
                Group's description. By default None.
  --synchronous=SYNCHRONOUS
                Whether one should wait until the execution finishes or not. By
                default True.
```

Return:  
Tuple with job's uuid and status.

- **Command: deactivate:** Deactivate failure detector for server(s) in a group. By default the failure detector is disabled.

Usage:  
mysqlfabric group deactivate <group\_id> [--synchronous]

Parameters:  
<group\_id> Group's id.

Options:  
--synchronous=SYNCHRONOUS  
Whether one should wait until the execution finishes or not. By default True.

Return:  
Tuple with job's uuid and status.

- **Command: demote:** Demote the current master if there is one. If users just want to update the state store and skip provisioning steps such as configuring replication, the `update_only` parameter must be set to true. Otherwise any write access to the master is blocked, slaves are synchronized with the master, stopped and their replication configuration reset. Note that no slave is promoted as master.

Usage:  
mysqlfabric group demote <group\_id> [--update\_only] [--synchronous]

Parameters:  
<group\_id>

Options:  
--update\_only=UPDATE\_ONLY  
Only update the state store and skip provisioning. By default False.  
--synchronous=SYNCHRONOUS  
Whether one should wait until the execution finishes or not. By default True.

- **Command: description:** Update a group's description.

Usage:  
mysqlfabric group description <group\_id> [--description=DESCRIPTION] [--synchronous]

Parameters:  
<group\_id> Group's id.

Options:  
--description=DESCRIPTION  
Group's description. By default None.  
--synchronous=SYNCHRONOUS  
Whether one should wait until the execution finishes or not. By default True.

Return:  
Tuple with job's uuid and status.

- **Command: destroy:** Remove a group.

Usage:  
mysqlfabric group destroy <group\_id> [--synchronous]

```
Parameters:
  <group_id> Group's id.

Options:
  --synchronous=SYNCHRONOUS
                        Whether one should wait until the execution finishes or not. By
                        default True.

Return:
  Tuple with job's uuid and status.
```

- **Command:** `health`: Check if any server within a group has failed and report health information.

```
Usage:
  mysqlfabric group health <group_id> [--timeout=TIMEOUT]

Parameters:
  <group_id> Timeout value after which a server is considered unreachable. If
             None is provided, it assumes the default value in the
             configuration file.

Options:
  --timeout=TIMEOUT By default None.
```

- **Command:** `lookup_groups`: Return information on existing group(s).

```
Usage:
  mysqlfabric group lookup_groups [--group_id=GROUP_ID]

Options:
  --group_id=GROUP_ID None if one wants to list the existing groups or
                     group's id if one wants information on a group. By
                     default None.

Return:
  List with {"group_id" : group_id, "failure_detector": ON/OFF, "description"
           : description}.
```

- **Command:** `lookup_servers`: Return information on existing server(s) in a group.

```
Usage:
  mysqlfabric group lookup_servers <group_id> [--server_id=SERVER_ID]
                                             [--status=STATUS] [--mode=MODE]

Parameters:
  <group_id> Group's id.

Options:
  --server_id=SERVER_ID
                        None if one wants to list the existing servers in a group
                        or server's id if one wants information on a server in a
                        group. Accepted values: Servers's UUID or HOST:PORT. By
                        default None.
  --status=STATUS      Server's status one is searching for. By default None.
  --mode=MODE          Server's mode one is searching for. By default None.

Return:
  Information on servers. List with [uuid, address, status, mode, weight]
```

- **Command:** `promote`: Promote a server into master.



If users just want to update the state store and skip provisioning steps such as configuring replication, the `update_only` parameter must be set to true. Otherwise, the following happens.

If the master within a group fails, a new master is either automatically or manually selected among the slaves in the group. The process of selecting and setting up a new master after detecting that the current master failed is known as failover.

It is also possible to switch to a new master when the current one is still alive and kicking. The process is known as switchover and may be used, for example, when one wants to take the current master off-line for maintenance.

If a slave is not provided, the best candidate to become the new master is found. Any candidate must have the binary log enabled, should have logged the updates executed through the SQL Thread and both candidate and master must belong to the same group. The smaller the lag between slave and the master the better. So the candidate which satisfies the requirements and has the smaller lag is chosen to become the new master.

In the failover operation, after choosing a candidate, one makes the slaves point to the new master and updates the state store setting the new master.

In the switchover operation, after choosing a candidate, any write access to the current master is disabled and the slaves are synchronized with it. Failures during the synchronization that do not involve the candidate slave are ignored. Then slaves are stopped and configured to point to the new master and the state store is updated setting the new master.

```
Usage:
  mysqlfabric group promote <group_id> [--slave_id=SLAVE_ID] [--update_only]
  [--synchronous]

Parameters:
  <group_id>

Options:
  --slave_id=SLAVE_ID Candidate's UUID or HOST:PORT. By default None.
  --update_only=UPDATE_ONLY
                        Only update the state store and skip provisioning. By
                        default False.
  --synchronous=SYNCHRONOUS
                        Whether one should wait until the execution finishes or
                        not. By default True.
```

- **Command:** `remove`: Remove a server from a group.

```
Usage:
  mysqlfabric group remove <group_id> <server_id> [--synchronous]

Parameters:
  <group_id> Group's id.
  <server_id> Servers's UUID or HOST:PORT.

Options:
  --synchronous=SYNCHRONOUS
                        Whether one should wait until the execution finishes or
                        not. By default True.

Return:
  Tuple with job's uuid and status.
```

## 4.5 Manage Commands

- **Command:** `logging_level`: Set logging level.

```
Usage:
mysqlfabric manage logging_level <module> <level>

Parameters:
<module> Module that will have its logging level changed.
<level>   The logging level that will be set.
```

- **Command:** `ping`: Check whether the Fabric server is running or not.

```
Usage:
mysqlfabric manage ping
```

- **Command:** `setup`: Setup Fabric Storage System. Create a database and necessary objects.

```
Usage:
mysqlfabric manage setup [--read_pw_from_stdin]

Options:
--read_pw_from_stdin=READ_PW_FROM_STDIN
    Whether to read passwords from stdin instead of from the
    controlling tty. By default False.
```

- **Command:** `start`: Start the Fabric server.

```
Usage:
mysqlfabric manage start [--foreground] [--disable_clustering] [--bootstrap]
    [--cluster_seed=CLUSTER_SEED] [--cluster_uuid=CLUSTER_UUID]
    [--cluster_timeout=5] [--read_pw_from_stdin]

Options:
--foreground=FOREGROUND
    Whether it should be started as background process or not. Default
    is False. By default False.
--disable_clustering=DISABLE_CLUSTERING
    Whether it should start with its clustering capabilities disabled
    and never join a cluster. By default False.
--bootstrap=BOOTSTRAP
    Whether the node will be used to bootstrap the cluster or not.
    Default is False. By default False.
--cluster_seed=CLUSTER_SEED
    Information that will be used to discover members of the cluster
    and eventually try to join it. By default None.
--cluster_uuid=CLUSTER_UUID
    Cluster Identifier. Used to ensure that the node does not connect
    to the wrong cluster. By default None.
--cluster_timeout=CLUSTER_TIMEOUT
    Timeout after which the node will stop trying to join the cluster.
    By default 5.
--read_pw_from_stdin=READ_PW_FROM_STDIN
    Whether to read passwords from stdin instead of from the
    controlling tty. By default False.
```

- **Command:** `stop`: Stop the Fabric server.

```
Usage:
mysqlfabric manage stop
```

- **Command:** `teardown`: Tear down Fabric Storage System. Drop database and its objects.



**Note**

A `teardown` removes the backing store contents, therefore all configuration information is lost. It's the contrary of `manage setup`.

```
Usage:
  mysqlfabric manage teardown [--read_pw_from_stdin]

Options:
  --read_pw_from_stdin=READ_PW_FROM_STDIN
      Whether to read passwords from stdin instead of from the
      controlling tty. By default False.
```

## 4.6 Provider Commands

The provider commands are used to manage cloud providers.

- **Command:** `list`: Return information on existing provider(s).

```
Usage:
  mysqlfabric provider list [--provider_id=PROVIDER_ID]

Options:
  --provider_id=PROVIDER_ID
      None if one wants to list the existing providers or provider's id
      if one wants information on a provider. By default None.
```

- **Command:** `register`: Register a provider.

```
Usage:
  mysqlfabric provider register <provider_id> <url> [--tenant=TENANT]
  [--provider_type=OPENSTACK] [--default_image=DEFAULT_IMAGE]
  [--default_flavor=DEFAULT_FLAVOR] [--extra=EXTRA]
  [--synchronous]

Parameters:
  <provider_id>  Provider's Id.
  <url>          URL that is used as an access point.

Options:
  --tenant=TENANT  Tenant's name, i.e. who will access resources in the cloud.
                  By default None.
  --provider_type=PROVIDER_TYPE
                  Provider type. By default OPENSTACK.
  --default_image=DEFAULT_IMAGE
                  By default None.
  --default_flavor=DEFAULT_FLAVOR
                  By default None.
  --extra=EXTRA   Define parameters that are specific to a provider. By
                  default None.
  --synchronous=SYNCHRONOUS
                  Whether one should wait until the execution finishes or
                  not. By default True.

Return:
  Tuple with job's uuid and status.
```

- **Command:** `unregister`: Unregister a provider.

```

Usage:
  mysqlfabric provider unregister <provider_id> [--synchronous]

Parameters:
  <provider_id> Provider's Id.

Options:
  --synchronous=SYNCHRONOUS
                        Whether one should wait until the execution finishes or not. By
                        default True.

Return:
  Tuple with job's uuid and status.

```

## 4.7 Role Commands

The role commands are used to display information about an user's role as description and permissions.

- **Command:** `list`: List roles and associated permissions.

```

Usage:
  mysqlfabric role list [--name=NAME]

Options:
  --name=NAME Role's name. By default None.

```

## 4.8 Server Commands

The server commands are used to get information about servers and set their properties, namely status, mode and weight.

- **Command:** `clone`: Clone the objects of a given server into a destination server.

```

Usage:
  mysqlfabric server clone <group_id> <destn_address> [--source_id=SOURCE_ID]
                        [--timeout=TIMEOUT] [--synchronous]

Parameters:
  <group_id>          The ID of the source group.
  <destn_address>    The address of the destination MySQL Server.

Options:
  --source_id=SOURCE_ID
                        The address or UUID of the source MySQL Server. By
                        default None.
  --timeout=TIMEOUT   Time in seconds after which an error is reported if the
                        destination server is unreachable. By default None.
  --synchronous=SYNCHRONOUS
                        Whether one should wait until the execution finishes or
                        not. By default True.

```

- **Command:** `create`: Create a virtual machine instance.

Usage examples:

```

shell> mysqlfabric server create provider --image name=image-mysql \
      --flavor name=vm-template --meta db=mysql --meta=version=5.6
shell> mysqlfabric server create provider --image name=image-mysql \

```

```
--flavor name=vm-template --security_groups grp_fabric, grp_ham
```

Options that accept a list are defined by providing the same option multiple times in the command-line. The image, flavor, files, meta and scheduler\_hints are those which might be defined multiple times. Note the the security\_groups option might be defined only once but it accept a string with a list of security groups.

Usage:

```
mysqlfabric server create <provider_id> [--image=IMAGE] [--flavor=FLAVOR]
[--number_machines=1] [--availability_zone=AVAILABILITY_ZONE]
[--key_name=KEY_NAME] [--security_groups=SECURITY_GROUPS]
[--private_network=PRIVATE_NETWORK]
[--public_network=PUBLIC_NETWORK] [--userdata=USERDATA]
[--swap=SWAP] [--scheduler_hints=SCHEDULER_HINTS] [--meta=META]
[--datastore=DATASTORE] [--datastore_version=DATASTORE_VERSION]
[--size=SIZE] [--databases=DATABASES] [--users=USERS]
[--configuration=CONFIGURATION] [--security=SECURITY]
[--skip_store] [--wait_spawning] [--synchronous]
```

Parameters:

```
<provider_id> Provider's Id.
```

Options:

```
--image=IMAGE          Image's properties. (e.g. name=image-mysql). Accepted
                        values: list of key/value pairs. By default None.
--flavor=FLAVOR         Flavor's properties (e.g. name=vm-template). Accepted
                        values: list of key/value pairs. By default None.
--number_machines=NUMBER_MACHINES
                        Number of machines to be created. Accepted values:
                        integer. By default 1.
--availability_zone=AVAILABILITY_ZONE
                        Name of availability zone. Accepted values: string. By
                        default None.
--key_name=KEY_NAME     Name of the key previously created. Accepted values:
                        string. By default None.
--security_groups=SECURITY_GROUPS
                        Security groups to have access to the machine(s).
                        Accepted values: string with a list of security groups.
                        By default None.
--private_network=PRIVATE_NETWORK
                        Name of the private network where the machine(s) will
                        be placed to. By default None.
--public_network=PUBLIC_NETWORK
                        Name of the public network which will provide a public
                        address. By default None.
--userdata=USERDATA     Script that to be used to configure the machine(s).
                        Accepted values: path to a file. By default None.
--swap=SWAP             Size of the swap space in megabyte. Accepted values:
                        integer. By default None.
--scheduler_hints=SCHEDULER_HINTS
                        Information on which host(s) the machine(s) will be
                        created in. Accepted values: list of key/value pairs.
                        By default None.
--meta=META            Metadata on the machine(s). Accepted values: list of
                        key/value pairs. By default None.
--datastore=DATASTORE   Database Technology (.e.g. MySQLQ). Accepted values:
                        string. By default None.
--datastore_version=DATASTORE_VERSION
                        Datastore version (.e.g. 5.6). Accepted values: string.
                        By default None.
--size=SIZE            Storage area reserved to the data store. Accepted
                        values: string in Gigabytes. By default None.
--databases=DATABASES  Database objects that will be created. Accepted values:
```

```
--users=USERS          List of strings separated by comma. By default None.
                       By default None.
--configuration=CONFIGURATION
                       Configuration attached to the database. Accepted
                       values: string. By default None.
--security=SECURITY    By default 0.0.0.0/0 is set. Users who want a differnt
                       permission should specify a different value. Accepted
                       values: string. By default None.
--skip_store=SKIP_STORE
                       Do not store information on machine(s) into the state
                       store. Default is False. By default False.
--wait_spawning=WAIT_SPAWNING
                       Whether one should wait until the provider finishes its
                       task or not. By default True.
--synchronous=SYNCHRONOUS
                       Whether one should wait until the execution finishes or
                       not. By default True.
```

- **Command:** `destroy`: Destroy a virtual machine instance.

```
Usage:
  mysqlfabric server destroy <provider_id> <machine_uuid> [--force]
  [--skip_store] [--synchronous]

Parameters:
  <provider_id>  Provider's Id.
  <machine_uuid> Machine's uuid.

Options:
  --force=FORCE  Ignore errors while accessing the cloud provider. By default
                 False.
  --skip_store=SKIP_STORE
                 Proceed anyway if there is no information on the machine in
                 the state store. Default is False. By default False.
  --synchronous=SYNCHRONOUS
                 Whether one should wait until the execution finishes or not.
                 By default True.

Return:
  Tuple with job's uuid and status.
```

- **Command:** `list`: Return information on existing machine(s) created by a provider.

```
Usage:
  mysqlfabric server list <provider_id> [--generic_filters=GENERIC_FILTERS]
  [--meta_filters=META_FILTERS] [--skip_store]

Parameters:
  <provider_id>  Provider's Id.

Options:
  --generic_filters=GENERIC_FILTERS
                 Set of key-value pairs that are used to filter the list of
                 returned machines. By default None.
  --meta_filters=META_FILTERS
                 Set of key-value pairs that are used to filter the list of
                 returned machines. By default None.
  --skip_store=SKIP_STORE
                 Don't check the list of machines from the state store. By default
                 False.
```

- **Command:** `lookup_uuid`: Return server's uuid.

```
Usage:
```

```
mysqlfabric server lookup_uuid <address> [--timeout=TIMEOUT]
```

Parameters:

<address> Server's address.

Options:

--timeout=TIMEOUT Time in seconds after which an error is reported if the UUID is not retrieved. By default None.

Return:

UUID.

- **server set\_mode:** Set a server's mode, which determines whether it can process read\_only, read\_write, or both transaction types. It can also be set to offline meaning that the server does not process any kind of user's request.

Usage:

```
mysqlfabric server set_mode <server_id> <mode> [--synchronous]
```

Parameters:

<server\_id> Servers's UUID or HOST:PORT.  
<mode>

Options:

--synchronous=SYNCHRONOUS  
Whether one should wait until the execution finishes or not. By default True.

- **Command: server set\_status:** Set a server's status.

Any server added into a group has to be alive and kicking and its status is automatically set to SECONDARY. If the failure detector is activate and the server is not reachable, it is automatically set to FAULTY.

Users can also manually change the server's status. Usually, a user may change a slave's mode to SPARE to avoid write and read access and guarantee that it is not chosen when a failover or switchover routine is executed.

By default replication is automatically configured when a server has its status changed. In order to skip this, users must set the update\_only parameter to true. If done so, only the state store is updated with information on the new status.

Usage:

```
mysqlfabric server set_status <server_id> <status> [--update_only]
[--synchronous]
```

Parameters:

<server\_id> Servers's UUID or HOST:PORT.  
<status> Server's status.

Options:

--update\_only=UPDATE\_ONLY  
By default False.  
--synchronous=SYNCHRONOUS  
Whether one should wait until the execution finishes or not. By default True.

- **Command: server set\_weight:** Set a server's weight.

`server set_weight`: Set a server's weight, which helps determine its likelihood of being chosen to process requests or replace a failed master. The value must be greater than 0.0 and lower or equal to 1.0.



### Note

This option was implemented in Fabric 1.5.7.

```
Usage:
  mysqlfabric server set_weight <server_id> <weight> [--synchronous]

Parameters:
  <server_id>  Servers's UUID or HOST:PORT.
  <weight>     Server's weight.

Options:
  --synchronous=SYNCHRONOUS
                Whether one should wait until the execution finishes or not. By
                default True.
```

## 4.9 Sharding Commands

The sharding commands are used to define, modify, and control sharding.

- **Command:** `add_shard`: Add a shard.

```
Usage:
  mysqlfabric sharding add_shard <shard_mapping_id> <groupid_lb_list>
  [--state=DISABLED] [--update_only] [--synchronous]

Parameters:
  <shard_mapping_id>  The unique identification for a shard mapping.
  <groupid_lb_list>   The list of group_id, lower_bounds pairs in the format,
                    group_id/lower_bound, group_id/lower_bound...

Options:
  --state=STATE  Indicates whether a given shard is ENABLED or DISABLED. By
                 default DISABLED.
  --update_only=UPDATE_ONLY
                 Only update the state store and skip adding range checks. By
                 default False.
  --synchronous=SYNCHRONOUS
                 Whether one should wait until the execution finishes or not.
                 By default True.

Return:
  A dictionary representing the current Range specification.
```

- **Command:** `add_table`: Add a table to a shard mapping.

```
Usage:
  mysqlfabric sharding add_table <shard_mapping_id> <table_name>
  <column_name> [--range_check] [--update_only] [--synchronous]

Parameters:
  <shard_mapping_id>  The shard mapping id to which the input table is
                    attached.
  <table_name>       The table being sharded.
  <column_name>     The column whose value is used in the sharding scheme
                    being applied
```



```
Options:
--range_check=RANGE_CHECK
    Indicates if range check should be turned on for this table. By
    default False.
--update_only=UPDATE_ONLY
    Only update the state store and skip adding range checks. By
    default False.
--synchronous=SYNCHRONOUS
    Whether one should wait until the execution finishes or not. By
    default True.
```

- **Command:** `create_definition`: Define a shard mapping.

```
Usage:
mysqlfabric sharding create_definition <type_name> <group_id>
[--synchronous]

Parameters:
<type_name> The type of sharding scheme - RANGE or HASH
<group_id> Every shard mapping is associated with a global group that
stores the global updates and the schema changes for this shard
mapping and dissipates these to the shards.

Options:
--synchronous=SYNCHRONOUS
    Whether one should wait until the execution finishes or not. By
    default True.
```

- **Command:** `disable_shard`: Disable a shard.

```
Usage:
mysqlfabric sharding disable_shard <shard_id> [--synchronous]

Parameters:
<shard_id> The shard ID of the shard that needs to be removed.

Options:
--synchronous=SYNCHRONOUS
    Whether one should wait until the execution finishes or not. By
    default True.
```

- **Command:** `enable_shard`: Enable a shard.

```
Usage:
mysqlfabric sharding enable_shard <shard_id> [--synchronous]

Parameters:
<shard_id> The shard ID of the shard that needs to be removed.

Options:
--synchronous=SYNCHRONOUS
    Whether one should wait until the execution finishes or not. By
    default True.
```

- **Command:** `list_definitions`: Lists all the shard mapping definitions.

```
Usage:
mysqlfabric sharding list_definitions

Return:
A list of shard mapping definitions An Empty List if no shard mapping
definition is found.
```

- **Command:** `list_tables`: Returns all the shard mappings of a particular sharding\_type

```
Usage:
  mysqlfabric sharding list_tables <sharding_type>

Parameters:
  <sharding_type> The sharding type for which the sharding specification
                  needs to be returned.

Return:
  A list of dictionaries of shard mappings that are of the sharding type An
  empty list of the sharding type is valid but no shard mapping definition is
  found An error if the sharding type is invalid.
```

- **Command:** `lookup_servers`: Lookup a shard based on the give sharding key.

```
Usage:
  mysqlfabric sharding lookup_servers <table_name> <key> [--hint=LOCAL]

Parameters:
  <table_name> The table whose sharding specification needs to be looked up.
  <key>        The key value that needs to be looked up

Options:
  --hint=HINT A hint indicates if the query is LOCAL or GLOBAL. By default
              LOCAL.

Return:
  The Group UUID that contains the range in which the key belongs.
```

- **Command:** `lookup_table`: Fetch the shard specification mapping for the given table.

```
Usage:
  mysqlfabric sharding lookup_table <table_name>

Parameters:
  <table_name> The name of the table for which the sharding specification is
              being queried.

Return:
  The a dictionary that contains the shard mapping information for the given
  table.
```

- **Command:** `move_shard`: Move the shard represented by the shard\_id to the destination group.

By default this operation takes a backup, restores it on the destination group and guarantees that source and destination groups are synchronized before pointing the shard to the new group. If users just want to update the state store and skip these provisioning steps, the `update_only` parameter must be set to true.

```
Usage:
  mysqlfabric sharding move_shard <shard_id> <group_id> [--update_only]
                                     [--synchronous]

Parameters:
  <shard_id> The ID of the shard that needs to be moved.
  <group_id> The ID of the group to which the shard needs to be moved.

Options:
  --update_only=UPDATE_ONLY
                  By default False.
  --synchronous=SYNCHRONOUS
                  Whether one should wait until the execution finishes or not. By
```

```
default True.
```

- **Command:** `prune_shard`: Given the table name prune the tables according to the defined sharding specification for the table.

```
Usage:
  mysqlfabric sharding prune_shard <table_name> [--synchronous]

Parameters:
  <table_name>  The table that needs to be sharded.

Options:
  --synchronous=SYNCHRONOUS
                  Whether one should wait until the execution finishes or not. By
                  default True.
```

- **Command:** `remove_definition`: Remove the shard mapping definition represented by the Shard Mapping ID.

```
Usage:
  mysqlfabric sharding remove_definition <shard_mapping_id> [--synchronous]

Parameters:
  <shard_mapping_id>  The shard mapping ID of the shard mapping definition
                      that needs to be removed.

Options:
  --synchronous=SYNCHRONOUS
                  Whether one should wait until the execution finishes or not. By
                  default True.
```

- **Command:** `remove_shard`: Remove a shard.

```
Usage:
  mysqlfabric sharding remove_shard <shard_id> [--synchronous]

Parameters:
  <shard_id>  The shard ID of the shard that needs to be removed.

Options:
  --synchronous=SYNCHRONOUS
                  Whether one should wait until the execution finishes or not. By
                  default True.
```

- **Command:** `remove_table`: Remove the shard mapping represented by the Shard Mapping object.

```
Usage:
  mysqlfabric sharding remove_table <table_name> [--synchronous]

Parameters:
  <table_name>  The name of the table whose sharding specification is being
                removed.

Options:
  --synchronous=SYNCHRONOUS
                  Whether one should wait until the execution finishes or not. By
                  default True.
```

- **Command:** `split_shard`: Split the shard represented by the `shard_id` into the destination group.

By default this operation takes a backup, restores it on the destination group and guarantees that source and destination groups are synchronized before pointing the shard to the new group. If users just want to update the state store and skip these provisioning steps, the `update_only` parameter must be set to true.

```
Usage:
  mysqlfabric sharding split_shard <shard_id> <group_id>
                [--split_value=SPLIT_VALUE] [--update_only] [--synchronous]

Parameters:
  <shard_id>  The shard_id of the shard that needs to be split.
  <group_id>  The ID of the group into which the split data needs to be moved.

Options:
  --split_value=SPLIT_VALUE
                The value at which the range needs to be split. By default None.
  --update_only=UPDATE_ONLY
                By default False.
  --synchronous=SYNCHRONOUS
                Whether one should wait until the execution finishes. By default
                True.
```

## 4.10 Snapshot Commands

The snapshot commands are related to the creation or destruction of machine snapshots.

- **Command:** `create`: Create a snapshot image from a machine.

```
Usage:
  mysqlfabric snapshot create <provider_id> <machine_uuid> [--skip_store]
                [--wait_spawning] [--synchronous]

Parameters:
  <provider_id>  Provider's Id.
  <machine_uuid> Machine's uuid.

Options:
  --skip_store=SKIP_STORE
                Proceed anyway if there is no information on the machine in the
                state store. Default is False. By default False.
  --wait_spawning=WAIT_SPawning
                By default True.
  --synchronous=SYNCHRONOUS
                Whether one should wait until the execution finishes or not. By
                default True.

Return:
  Tuple with job's uuid and status.
```

- **Command:** `destroy`: Destroy snapshot images associated to a machine.

```
Usage:
  mysqlfabric snapshot destroy <provider_id> <machine_uuid> [--skip_store]
                [--synchronous]

Parameters:
  <provider_id>  Provider's Id.
  <machine_uuid> Machine's uuid.

Options:
  --skip_store=SKIP_STORE
                Proceed anyway if there is no information on the machine in the
                state store. Default is False. By default False.
```

```
--synchronous=SYNCHRONOUS
    Whether one should wait until the execution finishes or not. By
    default True.
```

```
Return:
    Tuple with job's uuid and status.
```

## 4.11 Statistics Commands

The statistics commands are used to Retrieve statistics at node, group or procedure level.

- **Command:** `group`: Retrieve statistics on Groups.

```
Usage:
    mysqlfabric statistics group [--group_id=GROUP_ID]

Options:
    --group_id=GROUP_ID  Group one wants to retrieve information on. By default
                        None.
```

- **Command:** `node`: Retrieve statistics on the Fabric node.

```
Usage:
    mysqlfabric statistics node
```

- **Command:** `procedure`: Retrieve statistics on Procedures.

```
Usage:
    mysqlfabric statistics procedure [--procedure_name=PROCEDURE_NAME]

Options:
    --procedure_name=PROCEDURE_NAME
        Procedure one wants to retrieve information on. By default None.
```

## 4.12 Threat Commands

The threat commands are used to report that a server is not working properly for any reason, these commands can be used by external entities (e.g. connectors) and MySQL Fabric itself.

- **Command:** `report_error`: Report a server error. If there are many issues reported by different servers within a period of time, the server is marked as faulty. Should the server be a primary, the failover mechanism is triggered. Users who only want to set the server's status to faulty after getting enough notifications from different clients must set the `update_only` parameter to true. By default its value is false.

```
Usage:
    mysqlfabric threat report_error <server_id> [--reporter=UNKNOWN]
        [--error=UNKNOWN] [--update_only] [--synchronous]

Parameters:
    <server_id>  Servers's UUID or HOST:PORT.

Options:
    --reporter=REPORTER  Who has reported the issue, usually an IP address or a
                        host name. By default unknown.
    --error=ERROR        Error that has been reported. By default unknown.
    --update_only=UPDATE_ONLY
```

```

    Only update the state store and skip provisioning. By
    default False.
--synchronous=SYNCHRONOUS
    By default True.

```

- **Command:** `report_failure`: Report with certainty that a server has failed or is unreachable. Should the server be a primary, the failover mechanism is triggered. Users who only want to set the server's status to faulty must set the `update_only` parameter to True. By default its value is false.

```

Usage:
  mysqlfabric threat report_failure <server_id> [--reporter=UNKNOWN]
  [--error=UNKNOWN] [--update_only] [--synchronous]

Parameters:
  <server_id> Servers's UUID or HOST:PORT.

Options:
  --reporter=REPORTER Who has reported the issue, usually an IP address or a
  host name. By default unknown.
  --error=ERROR Error that has been reported. By default unknown.
  --update_only=UPDATE_ONLY
  Only update the state store and skip provisioning. By
  default False.
  --synchronous=SYNCHRONOUS
  By default True.

```

## 4.13 User Commands

The user commands are used to manage the Fabric user.

- **Command:** `add`: Add a new Fabric user.

```

Usage:
  mysqlfabric user add <username> [--protocol=PROTOCOL] [--roles=ROLES]

Parameters:
  <username> The username account to add.

Options:
  --protocol=PROTOCOL Protocol of the user (for example 'xmlrpc'). By default
  None.
  --roles=ROLES Comma separated list of roles, IDs or names (see `role
  list`). By default None.

```

- **Command:** `delete`: Delete a Fabric user.

```

Usage:
  mysqlfabric user delete <username> [--protocol=PROTOCOL] [--force]

Parameters:
  <username> The username account to delete.

Options:
  --protocol=PROTOCOL Protocol of the user (for example 'xmlrpc'). By default
  None.
  --force=FORCE Do not ask for confirmation. By default False.

```

- **Command:** `list`: List users and their roles.

```

Usage:
  mysqlfabric user list [--name=NAME]

```

```
Options:  
--name=NAME    User's name. By default None.
```

- **Command:** `password`: Change password for a Fabric user.

```
Usage:  
mysqlfabric user password <username> [--protocol=PROTOCOL]  
  
Parameters:  
<username>    The username to change the password of.  
  
Options:  
--protocol=PROTOCOL  Protocol of the user (for example 'xmlrpc'). By default  
                    None.
```

- **Command:** `roles`: Change roles for a Fabric user.

```
Usage:  
mysqlfabric user roles <username> [--protocol=PROTOCOL] [--roles=ROLES]  
  
Parameters:  
<username>    The username to change the roles of.  
  
Options:  
--protocol=PROTOCOL  Protocol of the user (for example 'xmlrpc'). By default  
                    None.  
--roles=ROLES        Comma separated list of roles, IDs or names (see `role  
                    list`). By default None.
```





## Chapter 5 Fabric Utility Command Matrix

The following table lists all of the commands available in the `mysqlfabric` utility. The table is sorted by group and command to make it easier to find things. Each group and command is listed with all available options and parameters. Below each is a short description of the task.

**Table 5.1 Fabric Commands**

Group	Command	Parameters	Options
<i>dump</i>	<i>fabric_nodes</i>		<code>--protocol=PROTOCOL</code>
<i>dump</i>	<i>servers</i>		<code>--connector_version=CONNECTOR</code> <code>--patterns=PATTERNS</code>
<i>dump</i>	<i>shard_index</i>		<code>--connector_version=CONNECTOR</code> <code>--patterns=PATTERNS</code>
<i>dump</i>	<i>shard_maps</i>		<code>--connector_version=CONNECTOR</code> <code>--patterns=PATTERNS</code>
<i>dump</i>	<i>shard_tables</i>		<code>--connector_version=CONNECTOR</code> <code>--patterns=PATTERNS</code>
<i>event</i>	<i>trigger</i>	<i>event</i>	<code>--args=ARGS, --kwargs=KWARGS, --locks=LOCKS</code>
<i>group</i>	<i>activate</i>	<i>group_id</i>	<code>--synchronous=SYNCHRONOUS</code>
<i>group</i>	<i>add</i>	<i>address, group_id</i>	<code>--synchronous=SYNCHRONOUS, --timeout=TIMEOUT, --update_only=UPDATE_ONLY</code>
<i>group</i>	<i>create</i>	<i>group_id</i>	<code>--description=DESCRIPTION, --synchronous=SYNCHRONOUS</code>
<i>group</i>	<i>deactivate</i>	<i>group_id</i>	<code>--synchronous=SYNCHRONOUS</code>
<i>group</i>	<i>demote</i>	<i>group_id</i>	<code>--synchronous=SYNCHRONOUS, --update_only=UPDATE_ONLY</code>
<i>group</i>	<i>description</i>	<i>group_id</i>	<code>--description=DESCRIPTION, --synchronous=SYNCHRONOUS</code>
<i>group</i>	<i>destroy</i>	<i>group_id</i>	<code>--synchronous=SYNCHRONOUS</code>
<i>group</i>	<i>health</i>	<i>group_id</i>	<code>--timeout=TIMEOUT</code>

Group	Command	Parameters	Options
<i>group</i>	<i>lookup_groups</i>		<i>--group_id=GROUP_ID</i>
<i>group</i>	<i>lookup_servers</i>	<i>group_id</i>	<i>--mode=MODE, --server_id=SERVER_ID, --status=STATUS</i>
<i>group</i>	<i>promote</i>	<i>group_id</i>	<i>--slave_id=SLAVE_ID, --synchronous=SYNCHRONOUS, --update_only=UPDATE_ONLY</i>
<i>group</i>	<i>remove</i>	<i>group_id, server_id</i>	<i>--synchronous=SYNCHRONOUS</i>
<i>manage</i>	<i>logging_level</i>	<i>level, module</i>	
<i>manage</i>	<i>ping</i>		
<i>manage</i>	<i>setup</i>		<i>--read_pw_from_stdin=READ_PW_FRO</i>
<i>manage</i>	<i>start</i>		<i>--bootstrap=BOOTSTRAP, --cluster_seed=CLUSTER_SEED, --cluster_timeout=CLUSTER_TIMEOUT, --cluster_uuid=CLUSTER_UUID, --disable_clustering=DISABLE_CLUSTER, --foreground=FOREGROUND, --read_pw_from_stdin=READ_PW_FRO</i>
<i>manage</i>	<i>stop</i>		
<i>manage</i>	<i>teardown</i>		<i>--read_pw_from_stdin=READ_PW_FRO</i>
<i>provider</i>	<i>list</i>		<i>--provider_id=PROVIDER_ID</i>
<i>provider</i>	<i>register</i>	<i>provider_id, url</i>	<i>--default_flavor=DEFAULT_FLAVOR, --default_image=DEFAULT_IMAGE, --extra=EXTRA, --provider_type=PROVIDER_TYPE, --synchronous=SYNCHRONOUS, --tenant=TENANT</i>
<i>provider</i>	<i>unregister</i>	<i>provider_id</i>	<i>--synchronous=SYNCHRONOUS</i>
<i>role</i>	<i>list</i>		<i>--name=NAME</i>
<i>server</i>	<i>clone</i>	<i>destn_address, group_id</i>	<i>--source_id=SOURCE_ID,</i>

Group	Command	Parameters	Options
			-- synchronous=SYNCHRONOUS, --timeout=TIMEOUT
<i>server</i>	<i>create</i>	<i>provider_id</i>	-- availability_zone=AVAILABILITY_ZONE, -- configuration=CONFIGURATION, -- databases=DATABASES, -- datastore=DATASTORE, -- datastore_version=DATASTORE_VERSION, --flavor=FLAVOR, --image=IMAGE, -- key_name=KEY_NAME, --meta=META, -- number_machines=NUMBER_MACHINES, -- private_network=PRIVATE_NETWORK, -- public_network=PUBLIC_NETWORK, -- scheduler_hints=SCHEDULER_HINTS, --security=SECURITY, -- security_groups=SECURITY_GROUPS, --size=SIZE, -- skip_store=SKIP_STORE, --swap=SWAP, -- synchronous=SYNCHRONOUS, --userdata=USERDATA, --users=USERS, -- wait_spawning=WAIT_SPANNING
<i>server</i>	<i>destroy</i>	<i>machine_uuid,</i> <i>provider_id</i>	--force=FORCE, -- skip_store=SKIP_STORE, -- synchronous=SYNCHRONOUS
<i>server</i>	<i>list</i>	<i>provider_id</i>	-- generic_filters=GENERIC_FILTERS, -- meta_filters=META_FILTERS, -- skip_store=SKIP_STORE
<i>server</i>	<i>lookup_uuid</i>	<i>address</i>	--timeout=TIMEOUT
<i>server</i>	<i>set_mode</i>	<i>mode, server_id</i>	-- synchronous=SYNCHRONOUS
<i>server</i>	<i>set_status</i>	<i>server_id, status</i>	-- synchronous=SYNCHRONOUS, -- update_only=UPDATE_ONLY

Group	Command	Parameters	Options
<i>server</i>	<i>set_weight</i>	<i>server_id, weight</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>add_shard</i>	<i>groupid_lb_list, shard_mapping_id</i>	-- <i>state=STATE, --synchronous=SYNCHRONOUS, --update_only=UPDATE_ONLY</i>
<i>sharding</i>	<i>add_table</i>	<i>column_name, column_name, shard_mapping_id, table_name</i>	-- <i>range_check=RANGE_CHECK, --synchronous=SYNCHRONOUS, --update_only=UPDATE_ONLY</i>
<i>sharding</i>	<i>create_definition</i>	<i>group_id, type_name</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>disable_shard</i>	<i>shard_id</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>enable_shard</i>	<i>shard_id</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>list_definitions</i>		
<i>sharding</i>	<i>list_tables</i>	<i>sharding_type</i>	
<i>sharding</i>	<i>lookup_servers</i>	<i>key, table_name</i>	-- <i>hint=HINT</i>
<i>sharding</i>	<i>lookup_table</i>	<i>table_name</i>	
<i>sharding</i>	<i>move_shard</i>	<i>group_id, shard_id</i>	-- <i>synchronous=SYNCHRONOUS, --update_only=UPDATE_ONLY</i>
<i>sharding</i>	<i>prune_shard</i>	<i>table_name</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>remove_definition</i>	<i>shard_mapping_id</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>remove_shard</i>	<i>shard_id</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>remove_table</i>	<i>table_name</i>	-- <i>synchronous=SYNCHRONOUS</i>
<i>sharding</i>	<i>split_shard</i>	<i>group_id, shard_id</i>	-- <i>split_value=SPLIT_VALUE, --synchronous=SYNCHRONOUS, --update_only=UPDATE_ONLY</i>
<i>snapshot</i>	<i>create</i>	<i>machine_uuid, provider_id</i>	-- <i>skip_store=SKIP_STORE, --synchronous=SYNCHRONOUS, --wait_spawning=WAIT_SPAWNING</i>

<b>Group</b>	<b>Command</b>	<b>Parameters</b>	<b>Options</b>
<i>snapshot</i>	<i>destroy</i>	<i>machine_uuid,</i> <i>provider_id</i>	-- <i>skip_store=SKIP_STORE,</i> -- <i>synchronous=SYNCHRONOUS</i>
<i>statistics</i>	<i>group</i>		-- <i>group_id=GROUP_ID</i>
<i>statistics</i>	<i>node</i>		
<i>statistics</i>	<i>procedure</i>		-- <i>procedure_name=PROCEDURE_</i>
<i>threat</i>	<i>report_error</i>	<i>server_id</i>	-- <i>error=ERROR, --</i> <i>reporter=REPORTER, --</i> <i>synchronous=SYNCHRONOUS,</i> -- <i>update_only=UPDATE_ONLY</i>
<i>threat</i>	<i>report_failure</i>	<i>server_id</i>	-- <i>error=ERROR, --</i> <i>reporter=REPORTER, --</i> <i>synchronous=SYNCHRONOUS,</i> -- <i>update_only=UPDATE_ONLY</i>
<i>user</i>	<i>add</i>	<i>username</i>	-- <i>protocol=PROTOCOL,</i> <i>roles=ROLES</i>
<i>user</i>	<i>delete</i>	<i>username</i>	-- <i>force=FORCE, --</i> <i>protocol=PROTOCOL</i>
<i>user</i>	<i>list</i>		-- <i>name=NAME</i>
<i>user</i>	<i>password</i>	<i>username</i>	-- <i>protocol=PROTOCOL</i>
<i>user</i>	<i>roles</i>	<i>username</i>	-- <i>protocol=PROTOCOL,</i> <i>roles=ROLES</i>



---

# Chapter 6 Backing Store

## Table of Contents

6.1 Backing Store Tables ..... 57  
6.2 Protecting the Backing Store ..... 60

The backing store feature requires a MySQL instance. This server should be the same version as your other servers and MySQL version 5.6.10 or later. This section explains how to set up the backing store and provides information about some of the tables created.

To set up the backing store, use the `mysqlfabric` command. The `--param` options specify the user and password we created in [Section 2.3.1, “Create the Associated MySQL Users”](#) so that the utility can connect to the backing store and create the database and tables. We show the resulting tables in the new `fabric` database below.

```
shell> mysqlfabric manage setup --param=storage.user=fabric --param=storage.password=secret
[INFO] 1379444563.457977 - MainThread - Initializing persister:
user (fabric), server (localhost:3306), database (fabric).
shell> mysqlshow -ufabric -psecret fabric
+-----+
| Tables |
+-----+
| checkpoints |
| error_log |
| group_replication |
| groups |
| permissions |
| role_permissions |
| roles |
| servers |
| shard_maps |
| shard_ranges |
| shard_tables |
| shards |
| user_roles |
| users |
+-----+
```



### Note

The tables described here are subject to change in future versions of Fabric.

## 6.1 Backing Store Tables

The `checkpoints` table stores information on procedures' executions and is used to safely resume executing a procedure after a crash and recovery:

**Table 6.1 checkpoints**

Column	Type	Comment
proc_uuid	varchar(40)	The procedure's unique identification
lockable_objects	blob	Set of objects locked by the procedure
job_uuid	varchar(60)	The job's unique identification
sequence	int(11)	The job's sequence in the execution

Column	Type	Comment
action_fqn	text	Reference to the fully qualified name of the function to be executed on behalf of the job
param_args	blob	Positional arguments to the job's function
param_kwargs	blob	Keyword arguments to the job's function
started	double(16,6)	When the job started
finished	double(16,6)	When the job finished

The [error\\_log](#) table contains information on servers' errors reported.

**Table 6.2 error\_log**

Column	Type	Comment
server_uuid	varchar(40)	The <a href="#">server_uuid</a> value from the server that has raised an error
reported	timestamp(6)	When the error was reported
reporter	varchar(64)	Who reported the error: IP address, host name
error	text	Error message or code reported

The [group\\_replication](#) table defines replication among global groups and groups used in shards. They are used primarily for shard splitting, moving, and global updates.

**Table 6.3 group\_replication**

Column	Type
master_group_id	varchar(64)
slave_group_id	varchar(64)

The [groups](#) table contains information about the groups being managed.

**Table 6.4 groups**

Column	Type	Comment
group_id	varchar(64)	The identifier for the group
description	varchar(256)	A description of the group
master_uuid	varchar(40)	The <a href="#">server_uuid</a> value from the master server
master_defined	timestamp(6)	When the current master was defined.
status	bit(1)	1 if the group is being monitored, 0 otherwise

The [permissions](#) table contains information on rights to access the different sub-systems in Fabric. Currently, a [core](#) sub-system is formally defined:

**Table 6.5 permissions**

Column	Type	Comment
permission_id	int(10) unsigned	Permission's ID



Column	Type	Comment
subsystem	varchar(60)	Sub-system identification
component	varchar(60)	Sub-system component
function	varchar(60)	Sub-system function. Currently, this is not used
description	varchar(1000)	Description

The `roles` table contains information on possible roles a user may have and by consequence his/her permissions:

**Table 6.6 roles**

Column	Type	Comment
role_id	int(10) unsigned	Roles' ID
name	varchar(80)	Role's name
description	varchar(1000)	Role's description

The `role_permissions` table associates roles and permissions:

**Table 6.7 role\_permissions**

Column	Type	Comment
role_id	int(10) unsigned	Roles' ID
permission_id	int(10) unsigned	Permission's ID

The `servers` table contains a list of all servers managed by Fabric.

**Table 6.8 servers**

Column	Type	Comment
server_uuid	varchar(40)	UUID of the server
server_address	varchar(128)	Address of the server
mode	int(11)	Mode of the server (OFFLINE, READ_ONLY, WRITE_ONLY, READ_WRITE)
status	int(11)	Status of the server (FAULTY, SPARE, SECONDARY, PRIMARY)
weight	float	Likelihood of receiving a request
group_id	varchar(64)	Group the server belongs to

The `shard_maps` table contains the names and properties of the shard maps.

**Table 6.9 shard\_maps**

Column	Type	Comment
shard_mapping_id	int(11)	Shard map identifier
type_name	enum('RANGE','HASH')	Shard map type
global_group	varchar(64)	Name of the global group (likely to go away in the next revision)

The `shard_ranges` table is the sharding index and is used to map a sharding key to a shard.

**Table 6.10** `shard_ranges`

Column	Type	Comment
<code>shard_mapping_id</code>	<code>int(11)</code>	Shard map identifier
<code>lower_bound</code>	<code>varbinary(16)</code>	Lower bound for the range encoded as a binary string
<code>shard_id</code>	<code>int(11)</code>	Shard identifier (a number)

The `shard_tables` table lists all tables that are sharded and what sharding map each belongs to. It also names the column by which it is sharded.

**Table 6.11** `shard_tables`

Column	Type	Comment
<code>shard_mapping_id</code>	<code>int(11)</code>	Shard map identifier
<code>table_name</code>	<code>varchar(64)</code>	Fully qualified table name
<code>column_name</code>	<code>varchar(64)</code>	Column name that is the sharding key

The `shards` table names the groups where each shard identifier is stored.

**Table 6.12** `shards`

Column	Type	Comment
<code>shard_id</code>	<code>int(11)</code>	Shard identifier
<code>group_id</code>	<code>varchar(64)</code>	Group identifier (a dotted-name)
<code>state</code>	<code>enum('DISABLED','ENABLED')</code>	Status of the shard; <code>DISABLED</code> means that it is not available for use

The `users` table identifies the users that might have permission to access the functions in the different sub-systems:

**Table 6.13** `users`

Column	Type	Comment
<code>user_id</code>	<code>int(10) unsigned</code>	User's internal ID
<code>username</code>	<code>varchar(100)</code>	User's name
<code>protocol</code>	<code>varchar(200)</code>	Protocol that the user is allowed to use to access Fabric and its sub-systems
<code>password</code>	<code>varchar(128)</code>	Hashed user's password

## 6.2 Protecting the Backing Store

The backing store is very important to Fabric. You should take steps to ensure the database is backed up and the server where it resides is stable and well maintained. For the purposes of backup, it is sufficient to make periodic backups using either the `mysqldump` client or `mysqldbexport` utility in MySQL Utilities.

---

# Chapter 7 Using MySQL Fabric with Pacemaker and Corosync

## Table of Contents

7.1 Introduction .....	61
7.2 Pre-requisites .....	62
7.3 Target Configuration .....	62
7.4 Setting up and testing your system .....	63
7.4.1 Configure Network .....	63
7.4.2 Install all packages .....	63
7.4.3 Configure DRBD .....	64
7.4.4 Configure MySQL Server .....	66
7.4.5 Configure MySQL Fabric .....	67
7.4.6 Configure Corosync and Pacemaker .....	68
7.5 Key administrative tasks .....	71

There are a number of ways to make the MySQL Fabric node and its status and configuration data highly available; this section describes one such approach, but other alternatives are possible. This section can be used as a set of guidelines to build your own framework to add fault tolerance.



### Note

The described Pacemaker setup is not currently a solution that has been through QA.

## 7.1 Introduction

A MySQL Fabric instance is composed of a MySQL Fabric process together with a MySQL server. The MySQL Fabric node contains the protocols and the executor, but is in itself *stateless* in that if for some reason crashes, re-starting permits it to continue where it left off. The configuration and state information about the farm is instead stored in a separate MySQL server.

To provide a highly available solution, both components must be redundant. From a failure viewpoint, these two components are treated as a single unit, meaning that each pair is collocated on a machine, and if one of them fails, the stack on that machine fails.

There are two MySQL Fabric instances working in active and stand-by mode. The data stored in the MySQL server is replicated through the Distributed Replicated Block Device, or simply **DRBD**. The MySQL Fabric process is stateless though and must simply be started in the stand-by node in the event of a failure.

**Pacemaker and CoroSyc** are used to monitor whether the instances are running properly, and to automate the failover and switchover operations. Pacemaker is responsible for monitoring components, such as the MySQL Fabric Process, MySQL server, and DRBD, and also for executing the failover and switchover operations. CoroSyc is the communication infrastructure used by Pacemaker to exchange messages between the two nodes.

Applications accessing this cluster do so through a Virtual IP address that is assigned to the active node, specifically to the MySQL Fabric process, and is automatically migrated to the stand-by node during a failover or switchover operation.

This section aims to describe how to set up this highly available cluster.

## 7.2 Pre-requisites

Two servers or virtual machines with:

- A Linux distribution. This guide is based on Ubuntu 14.04, but any other distribution should work equally well.
- Unpartitioned space on the local disk to create a DRBD partition.
- Network connectivity
- Both hosts must be accessible through `ssh`.
- User "mysql" and group "mysql" that have the same ids at the different nodes.

Linux is used because Pacemaker, Corosync, and DRBD are commonly available on this platform.



### Note

Pacemaker, Corosync, and DRBD are not include with MySQL Fabric and need to be installed separately for the target platform.

If Virtual Machines are used, make sure they run in different physical servers to avoid a single point of failure. If possible, also make the network connectivity redundant.

## 7.3 Target Configuration

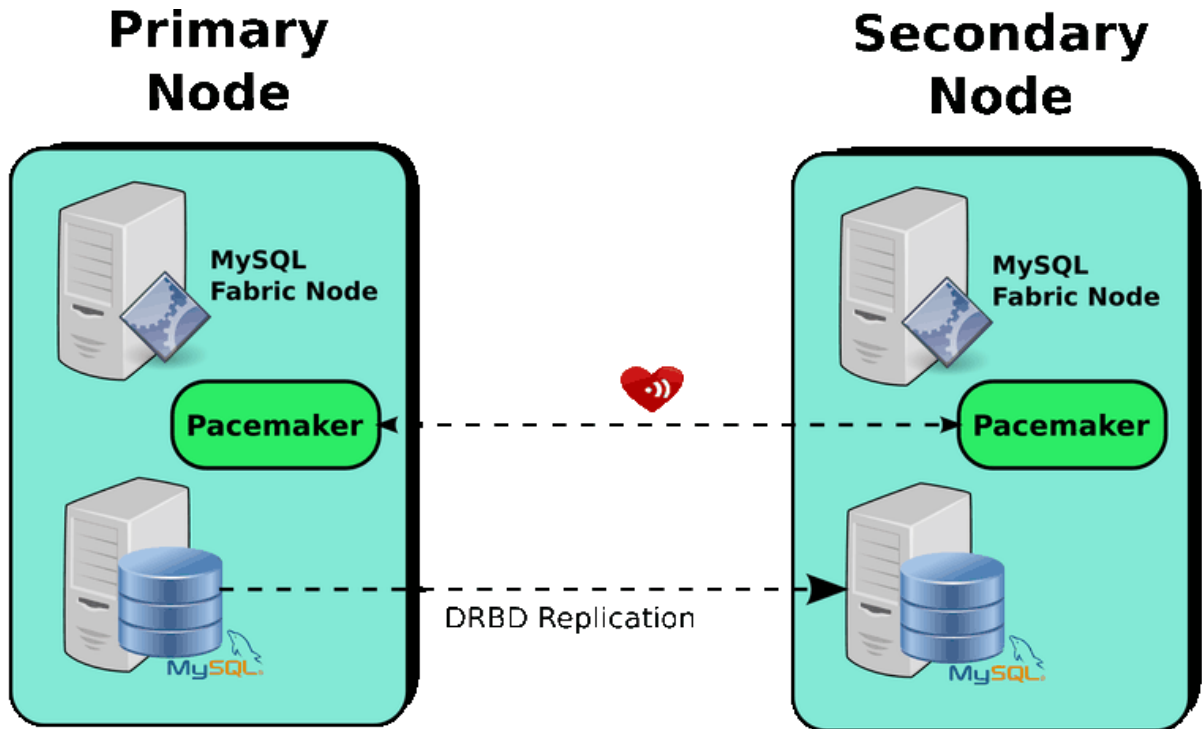
The two physical hosts are `host1.localdomain` (`192.168.1.101`) and `host2.localdomain` (`192.168.1.102`). It is recommended that you do not rely on an external DNS service (as that is an additional point of failure) and so these mappings should be configured on each host in the `/etc/hosts` file.

A single Virtual IP (VIP) is shown in the figure (`192.168.1.200`) and this is the address that the application connects to when accessing the MySQL Fabric. Pacemaker is responsible for migrating this between the two hosts.

One of the final steps in configuring Pacemaker is to add network connectivity monitoring in order to attempt to have an isolated host stop its MySQL services to avoid a split-brain scenario. This is achieved by having each host ping an external (not one part of the cluster) IP addresses - in this case the network router (`192.168.1.1`).

All the necessary software (i.e. binaries) must be installed in a regular partition, independent on each node. MySQL socket (`mysql.sock`) and process-id (`mysql.pid`) files are stored in a regular partition as well. The MySQL Server configuration file (`my.cnf`), the database files (`data/*`) and the MySQL Fabric configuration file (`fabric.cfg`) are stored in a DRBD controlled file system that at any point in time is only available on one of the two hosts.

Figure 7.1 MySQL Fabric Setup using DRBD and Pacemaker



## 7.4 Setting up and testing your system

### 7.4.1 Configure Network

It is recommended that you do not rely on DNS to resolve host names and so the following configuration files should be updated:

#### Example 7.1 /etc/hosts (Host 1)

```
127.0.0.1 localhost localhost.localdomain
::1 localhost localhost.localdomain
192.168.1.102 host2 host2.localdomain
```

#### Example 7.2 /etc/hosts (Host 2)

```
127.0.0.1 localhost localhost.localdomain
::1 localhost localhost.localdomain
192.168.1.101 host1 host1.localdomain
```

### 7.4.2 Install all packages

Install the necessary packages through the apt-get repositories:

```
[root@host1]# apt-get install drbd8-utils corosync pacemaker sysv-rc-conf libaio1
[root@host2]# apt-get install drbd8-utils corosync pacemaker sysv-rc-conf libaio1
```

Download and install the common, server, and client components on both hosts. Our example downloads and installs the bundled binaries from [dev.mysql.com](http://dev.mysql.com). Download the latest [MySQL Server release](#) bundle and install it on *both* machines using the following commands:

```
[root@host1]# dpkg -i mysql-common_ubuntu14.04_amd64.deb
[root@host1]# dpkg -i mysql-community-server_ubuntu14.04_amd64.deb
[root@host1]# dpkg -i mysql-community-client_ubuntu14.04_amd64.deb

[root@host2]# dpkg -i mysql-common_ubuntu14.04_amd64.deb
[root@host2]# dpkg -i mysql-community-server_ubuntu14.04_amd64.deb
[root@host2]# dpkg -i mysql-community-client_ubuntu14.04_amd64.deb
```

Next, install MySQL Fabric. Download MySQL Fabric by downloading the [MySQL Utilities](#) and install it using the following commands on each machine:

```
shell> unzip mysql-utilities-*.zip
shell> cd mysql-utilities-*
shell> python setup.py install
```

The script required to run MySQL Fabric with Pacemaker is not distributed with the packages and you need to manually download and install the script on each machine:

```
[root@host1]# cp mysql-fabric /usr/lib/ocf/resource.d/heartbeat/.
[root@host1]# chmod 755 /usr/lib/ocf/resource.d/heartbeat/mysql-fabric

[root@host2]# cp mysql-fabric /usr/lib/ocf/resource.d/heartbeat/.
[root@host2]# chmod 755 /usr/lib/ocf/resource.d/heartbeat/mysql-fabric
```

### 7.4.3 Configure DRBD

If your nodes do not already have an empty partition that you plan to use for the DRBD, then create one. If you are using a Virtual Machine, you can add a new storage to your machine. These details go beyond the scope of this guide.

This partition is used as a resource, managed (and synchronized between nodes by DRBD); in order for DRBD to be able to do this a new configuration file (in this case called `clusterdb_res.res`) must be created in the `/etc/drbd.d/` directory; the contents should look similar to:

```
resource clusterdb_res {
    protocol C;
    handlers {
        pri-on-incon-degr "/usr/lib/drbd/notify-pri-on-incon-degr.sh; /usr/lib/drbd/notify-emergency-reboot.sh; ec
        pri-lost-after-sb "/usr/lib/drbd/notify-pri-lost-after-sb.sh; /usr/lib/drbd/notify-emergency-reboot.sh; ec
        local-io-error "/usr/lib/drbd/notify-io-error.sh; /usr/lib/drbd/notify-emergency-shutdown.sh; echo o > /pr
        fence-peer "/usr/lib/drbd/crm-fence-peer.sh";
    }
    startup {
        degr-wfc-timeout 120;    # 2 minutes
        outdated-wfc-timeout 2; # 2 seconds
    }
    disk {
        on-io-error detach;
    }
    net {
        cram-hmac-alg "sha1";
        shared-secret "clusterdb";
        after-sb-0pri disconnect;
        after-sb-1pri disconnect;
    }
}
```

```

    after-sb-2pri disconnect;
    rr-conflict disconnect;
  }
  syncer {
    rate 10M;
    al-extents 257;
    on-no-data-accessible io-error;
  }
  on host1 {
    device /dev/drbd0;
    disk /dev/sdb;
    address 192.168.1.101:7788;
    flexible-meta-disk internal;
  }
  on host2 {
    device /dev/drbd0;
    disk /dev/sdb;
    address 192.168.1.102:7788;
    meta-disk internal;
  }
}

```

The IP addresses and disk locations should be specific to the hosts that the cluster are using. In this example the device that DRBD creates is located at `/dev/drbd0` - it is this device that is swapped back and forth between the hosts by DRBD. This resource configuration file should be copied over to the same location on the second host:

```
[root@host1]# scp clusterdb_res.res host2:/etc/drbd.d/
```

The configuration file previously presented uses DRBD 8.3 dialect. Although DRBD 8.4 is the newest version, some distributions might still contain DRBD 8.3. If you have installed DRBD 8.4 do not worry though because it understands the DRBD 8.3 configuration file.

Before starting the DRBD daemon, meta data must be created for the new resource (`clusterdb_res`) on each host using the command:

```
[root@host1]# drbdadm create-md clusterdb_res
```

```
[root@host2]# drbdadm create-md clusterdb_res
```

It is now possible to start the DRBD daemon on each host:

```
[root@host1]# /etc/init.d/drbd start
```

```
[root@host2]# /etc/init.d/drbd start
```

At this point the DRBD service is running on both hosts but neither host is the "primary" and so the resource (block device) cannot be accessed on either host; this can be confirmed by querying the status of the service:

```
[root@host1]# /etc/init.d/drbd status
```

```
[root@host2]# /etc/init.d/drbd status
```

In order to create the file systems (and go on to storing useful data in it), one of the hosts must be made the primary for the `clusterdb_res` resource, so execute the following on **host1**

```
[root@host1]# drbdadm -- --overwrite-data-of-peer primary all
[root@host1]# /etc/init.d/drbd status
```

The status output also shows the progress of the block-level syncing of the device from the new primary (host1) to the secondary (host2). This initial sync can take some time but it should not be necessary to wait for it to complete in order to complete the other steps.

Now that the device is available on **host1**, it is possible to create a file system on it:

```
[root@host1]# mkfs -t ext4 /dev/drbd0
```



### Note

The above does not need to be repeated on the second host as DRBD handles the syncing of the raw disk data

In order for the DRBD file system to be mounted, the `/var/lib/mysql_drbd` directory should be created on both hosts:

```
[root@host1]# mkdir /var/lib/mysql_drbd
[root@host1]# chown mysql /var/lib/mysql_drbd
[root@host1]# chgrp mysql /var/lib/mysql_drbd

[root@host2]# mkdir /var/lib/mysql_drbd
[root@host2]# chown mysql /var/lib/mysql_drbd
[root@host2]# chgrp mysql /var/lib/mysql_drbd
```

On just the one (DRBD active) host, the DRBD file system must be temporarily mounted:

```
[root@host1]# mount /dev/drbd0 /var/lib/mysql_drbd
```

## 7.4.4 Configure MySQL Server

First, we have to stop the MySQL Server on both hosts and update configuration files and create new data files. First, stop the MySQL server on both hosts using the command:

```
shell> /etc/init.d/mysql stop
```



### Note

If you are using an Ubuntu distribution you need change the `/etc/apparmor.d/usr.sbin.mysql` on both hosts according to the following diff:

```
@@ -40,8 +40,8 @@
 /usr/share/mysql/** r,
 # Allow data dir access
 - /var/lib/mysql/ r,
 - /var/lib/mysql/** rwk,
 + /var/lib/mysql_drbd/ r,
 + /var/lib/mysql_drbd/** rwk,
 # Allow log file access
 /var/log/mysql/ r,
```

If you do not do that, the MySQL server may not be able to access files in the new location and you may get strange errors since [AppArmor](#) prevents reading and writing from the new locations.



To avoid any mismatches, the configuration file can be copied from **host1** to **host2** :

```
shell> scp /etc/apparmor.d/usr.sbin.mysqlld host2:/etc/apparmor.d/usr.sbin.mysqlld
```

Then restart AppArmor on both hosts using:

```
shell> /etc/init.d/apparmor restart
```

Edit the `/etc/mysql/my.cnf` file and set `datadir /var/lib/mysql_drbd/data` in the `[mysqld]` section on both hosts.

To avoid any mismatches, the configuration file can be copied from **host1** to **host2** :

```
shell> scp /etc/mysql/my.cnf host2:/etc/mysql/my.cnf
```

Now the configuration file can be copied and the default database files populated on **host1** using:

```
shell> cp /etc/mysql/my.cnf /var/lib/mysql_drbd/my.cnf
shell> mkdir /var/lib/mysql_drbd/data
shell> mysql_install_db --no-defaults --datadir=/var/lib/mysql_drbd/data --user=mysql
```

Configure some permissions on **host1** :

```
shell> chmod -R uog+rw /var/lib/mysql_drbd
shell> chown -R mysql /var/lib/mysql_drbd
shell> chmod 644 /var/lib/mysql_drbd/my.cnf
```

Start MySQL Server and configure users and access:

```
shell> /etc/init.d/mysql start
```

```
shell> mysql -u root -e "GRANT ALL ON *.* to 'root'@'%';"
shell> mysql -u root -e "CREATE USER 'fabric'@'localhost' IDENTIFIED BY 'secret';"
shell> mysql -u root -e "GRANT ALL ON fabric.* TO 'fabric'@'localhost';"
```

## 7.4.5 Configure MySQL Fabric

On just the one (DRBD active) host, do the following:

```
shell> cp /etc/mysql/fabric.cfg /var/lib/mysql_drbd/fabric.cfg
shell> chmod 600 /var/lib/mysql_drbd/fabric.cfg
shell> chown root:root /var/lib/mysql_drbd/fabric.cfg
```

Edit the `/var/lib/mysql_drbd/fabric.cfg`:

1. Set **address** to `192.168.1.200:32274` in the `[protocol.xmlrpc]` section
2. Set **password** to `password` in the `[protocol.xmlrpc]` section
3. Set **address** to `192.168.1.200:32275` in the `[protocol.mysql]` section
4. Set the **password** to `password` in the `[protocol.mysql]` section
5. Set the **password** to `secret` in the `[storage]` section

**Warning**

Do *not* change the address in the `[storage]` section.

Take the opportunity to set the other options if you need/want to do so, specially the user/password in the `[servers]` and `[client]` sections. Finally, create MySQL Fabric's state store as follows:

```
[root@host1]# mysqlfabric --config /var/lib/mysql_drbd/fabric.cfg \
--param protocol.xmlrpc.address=localhost:32274 \
--param protocol.mysql.address=localhost:32275 manage setup
```

From this point onwards all resources are managed by the clustering software so they have to be stopped:

```
[root@host1]# /etc/init.d/mysql stop
[root@host1]# umount /var/lib/mysql_drbd
[root@host1]# drbdadm secondary clusterdb_res
```

Then DRBD should be stopped as well:

```
[root@host1]# /etc/init.d/drbd stop
[root@host2]# /etc/init.d/drbd stop
```

## 7.4.6 Configure Corosync and Pacemaker

At this point, the DRBD file system is configured and initialized and both MySQL Fabric and MySQL Server has been installed and the required files set up on the replicated DRBD file system. Pacemaker and Corosync are installed but they are not yet managing the MySQL Fabric Process, MySQL Server and DRBD resources to provide a clustered solution - the next step is to set that up.

Firstly, set up some network-specific parameters from the Linux command line and also in the Corosync configuration file. The multi-cast address should be unique in your network but the port can be left at 5405. The IP address should be based on the IP addresses being used by the servers but should take the form of XX.YY.ZZ.0.

Copy an example to make your life easier:

```
shell> cp /etc/corosync/corosync.conf.example /etc/corosync/corosync.conf
```

After editing it, it should have a content similar to what follows:

```
totem {
    version: 2
    crypto_cipher: none
    crypto_hash: none
    interface {
        ringnumber: 0
        bindnetaddr: 192.168.1.0
        mcastaddr: 239.255.1.1
        mcastport: 5405
        ttl: 1
    }
}
logging {
    to_syslog: yes
}
quorum {
    provider: corosync_votequorum
```

```
two_node: 1
wait_for_all: 1
}
nodelist {
  node {
    ring0_addr: 192.168.1.101
    nodeid: 1
  }
  node {
    ring0_addr: 192.168.1.102
    nodeid: 2
  }
}
```

Be careful while setting up the network address that the Corosync binds to. For example, according to the Corosync documentation, if the local interface is 192.168.5.92 with netmask 255.255.255.0, set `bindnetaddr` to 192.168.5.0. If the local interface is 192.168.5.92 with netmask 255.255.255.192, set `bindnetaddr` to 192.168.5.64, and so forth.

This makes Corosync automatically pick the network interface based on the network address provided. It is also possible to set up a specific address, such as 192.168.5.92, but in this case the configuration file is different per machine.

Create the `/etc/corosync/service.d/pcmk` file to tell the Corosync to load the Pacemaker plug-in:

```
service {
  # Load the Pacemaker Cluster Resource Manager
  name: pacemaker
  ver: 1
}
```

Change the `/etc/default/corosync` file as follows:

```
# start corosync at boot [yes|no]
START=yes
```

To avoid any mismatches, the configuration file can be copied across by using these commands on **host1**:

```
shell> scp /etc/corosync/corosync.conf host2:/etc/corosync/corosync.conf
shell> scp /etc/corosync/service.d/pcmk host2:/etc/corosync/service.d/pcmk
shell> scp /etc/default/corosync host2:/etc/default/corosync
```

Start Corosync on both hosts using:

```
shell> /etc/init.d/corosync start
```

Run `tcpdump` to check whether Corosync is working or not:

```
shell> tcpdump -i eth0 -n port 5405
```

To start the Pacemaker on **host1**, execute the following command:

```
shell> /etc/init.d/pacemaker start
```

Run Pacemaker's cluster resource monitoring command on **host1** to view the status of the cluster:

```
shell> crm_mon --one-shot -V
```

As we are configuring a cluster made up of just 2 hosts, when one host fails (or loses contact with the other) there is no node majority (quorum) left and so by default the surviving node (or both if they are still running but isolated from each other) would be shut down by Pacemaker. This is not the desired behavior as it does not offer High Availability and so that default should be overridden (we later add an extra behavior whereby each node shuts itself down if it cannot ping a 3 node that is external to the cluster, thus preventing a split brain situation):

```
[root@host1]# crm configure property no-quorum-policy=ignore
```

We turn STONITH (Shoot The Other Node In The Head) off as this solution relies on each node shutting itself down in the event that it loses connectivity with the independent host:

```
[root@host1]# crm configure property stonith-enabled=false
```

Roughly speaking, STONITH refers to one node trying to kill another in the even that it believes the other has partially failed and should be stopped in order to avoid any risk of a split-brain scenario. To prevent a healthy resource from being moved around the cluster when a node is brought back on-line, Pacemaker has the concept of resource stickiness which controls how much a service prefers to stay running where it is.

```
[root@host1]# crm configure rsc_defaults resource-stickiness=100
```

In the next steps, we describe how to configure the different resources as a cluster:

```
[root@host1]# crm configure edit
```

This opens your default text editor, and you should use it to add the following lines into the cluster configuration:

```
primitive p_drbd_mysql ocf:linbit:drbd \  
  params drbd_resource="clusterdb_res" \  
  op monitor interval="15s" \  
primitive p_fabric_mysql ocf:heartbeat:mysql-fabric \  
  params binary="/usr/local/bin/mysqlfabric" \  
  config="/var/lib/mysql_drbd/fabric.cfg" \  
  op start timeout="120s" interval="0" \  
  op stop timeout="120s" interval="0" \  
  op monitor interval="20s" timeout="30s" \  
primitive p_fs_mysql ocf:heartbeat:Filesystem \  
  params device="/dev/drbd0" directory="/var/lib/mysql_drbd" \  
  fstype="ext4" \  
primitive p_ip_mysql ocf:heartbeat:IPaddr2 \  
  params ip="192.168.1.200" cidr_netmask="24" nic="eth0" \  
primitive p_mysql ocf:heartbeat:mysql \  
  params binary="/usr/sbin/mysqld" \  
  config="/var/lib/mysql_drbd/my.cnf" \  
  datadir="/var/lib/mysql_drbd/data" \  
  pid="/var/run/mysqld/mysqld.pid" \  
  socket="/var/run/mysqld/mysqld.sock" \  
  user="mysql" group="mysql" \  
  additional_parameters="--bind-address=localhost" \  
  op start timeout="120s" interval="0" \  
  op stop timeout="120s" interval="0" \  
  op monitor interval="20s" timeout="30s" \  
group g_mysql p_fs_mysql p_ip_mysql p_mysql p_fabric_mysql \  
ms ms_drbd_mysql p_drbd_mysql \  
  meta master-max="1" master-node-max="1" clone-max="2" \  
  clone-node-max="1" notify="true" \  
colocation c_mysql_on_drbd inf: g_mysql ms_drbd_mysql:Master
```

```
order o_drbd_before_mysql inf: ms_drbd_mysql:promote g_mysql:start
primitive p_ping ocf:pacemaker:ping params name="ping" \
    multiplier="1000" host_list="192.168.1.1" \
    op monitor interval="15s" timeout="60s" start timeout="60s"
clone cl_ping p_ping meta interleave="true"
location l_drbd_master_on_ping ms_drbd_mysql rule $role="Master" \
    -inf: not_defined ping or ping number:lte 0
```

As the MySQL service (group) has a dependency on the host it is running on being the DRBD master, that relationship is added by defining a co-location and an ordering constraint to ensure that the MySQL group is co-located with the DRBD master and that the DRBD promotion of the host to the master must happen before the MySQL group can be started:

```
colocation c_mysql_on_drbd inf: g_mysql ms_drbd_mysql:Master
order o_drbd_before_mysql inf: ms_drbd_mysql:promote g_mysql:start
```

In order to prevent a split-brain scenario in the event of network partitioning, Pacemaker can ping independent network resources (such as a network router) and then prevent the host from being the DRBD master in the event that it becomes isolated:

```
primitive p_ping ocf:pacemaker:ping params name="ping" multiplier="1000" \
    host_list="192.168.1.1" \
    op monitor interval="15s" timeout="60s" start timeout="60s"
clone cl_ping p_ping meta interleave="true"
location l_drbd_master_on_ping ms_drbd_mysql rule $role="Master" \
    -inf: not_defined ping or ping number:lte 0
```

Check if everything is running fine using the following command:

```
[root@host1]# crm_mon --one-shot -V
```

## Ensure the correct daemons are started at system boot

At this point, a reliable MySQL service is in place but it is also important to check that the correct cluster services are started automatically as part of the servers' system startup. It is necessary for the Linux startup to start the Corosync and Pacemaker services but not DRBD or MySQL Process and MySQL Server as those services are started on the correct server by Pacemaker. To this end, execute the following commands on each host:

```
[root@host1] sysv-rc-conf drbd off
[root@host1] sysv-rc-conf corosync on
[root@host1] sysv-rc-conf mysql off
[root@host1] sysv-rc-conf pacemaker on

[root@host2] sysv-rc-conf drbd off
[root@host2] sysv-rc-conf corosync on
[root@host2] sysv-rc-conf mysql off
[root@host2] sysv-rc-conf pacemaker on
```



### Note

MySQL Fabric is not installed as a service so there is nothing to do here for it.

## 7.5 Key administrative tasks

The cluster management tool can then be used to migrate the resources between machines:

```
[root@host1 ~]# crm resource migrate g_mysql host2
```

Specifying the `g_mysql` group migrates all resources in the group and implicitly any colocated resources as well. If for any reason a resource cannot be properly started up or shut down, it becomes unmanaged. In this case, we have to manually put it back to a managed state. For example, this could mean that we would have to fix an issue that blocked the shutdown, kill or stop some processes, and run the following command:

```
[root@host1 ~]# crm resource cleanup 'resource'
```

The components of this stack are designed to cope with component failures but there may be cases where a sequence of multiple failures could result in DRBD not being confident that the data on the two hosts is consistent. In the event that this happens DRBD breaks the connection. Should this happen, we need to identify which of the two hosts has the correct data and then have DRBD resynchronize the data; for the steps below, it is assumed that host1 has the correct data:

```
[root@host2]# drbdadm secondary clusterdb_res  
[root@host2]# drbdadm -- --discard-my-data connect clusterdb_res  
  
[root@host1]# drbdadm primary clusterdb_res  
[root@host1]# drbdadm connect clusterdb_res
```

Before executing these steps, it is advised to check the error log(s) and run the following command to identify the DRBD's status:

```
shell> /etc/init.d/drbd status
```

---

# Chapter 8 Using Connector/Python with MySQL Fabric

## Table of Contents

8.1 Installing Connector/Python with MySQL Fabric Support .....	74
8.2 Requesting a Fabric Connection .....	74
8.3 Providing Information to Choose a MySQL Server .....	76

MySQL Fabric provides data distribution and high-availability features for a set of MySQL database servers.

Developers using Connector/Python can take advantage of its features to work with a set of servers managed by MySQL Fabric. Connector/Python supports the following MySQL Fabric capabilities:

- Automatic node selection based on application-provided *shard* information (tables and key)
- Read/write splitting within a MySQL Fabric *high-availability group*

More specifically, Connector/Python Fabric support provides these features:

- Requesting a connection to a MySQL server managed by Fabric is as transparent as possible to users already familiar with Connector/Python.
- Connector/Python is able to get a MySQL server connection given a high-availability group and a mode specifying whether the connection is read-only or also permits updates (read-write).
- Connector/Python supports sharding and is able to find the correct MySQL server for a given table or tables and key based on scope (local or global) and mode (read-only or read-write). *RANGE* and *HASH* mechanisms are supported transparently to the user.
- Among secondary MySQL servers in the same group, read-only connections are load balanced. Load balancing is based on a weight set for each MySQL server, using a Weighted Round-Robin algorithm.
- Faulty MySQL servers are reported to Fabric, and failover is supported when failure occurs for a server in a group.
- To speed up operations, Connector/Python caches information obtained from Fabric, such as group and sharding information. Each time Connector/Python looks up data, it first checks its cache. When information in the cache is too old, it is invalidated. By default, the time-to-live (TTL) for cached information is 1 minute. However, Fabric itself can provide this TTL for its clients and this value is used instead if greater than zero.

Cache information is also invalidated when failures occur, such as when a connection to a MySQL server fails (invalidation applies to the group to which the server belongs).

- Fabric support applies to versions of Python supported by Connector/Python itself (see [Connector/Python Versions](#)). In particular, you can use Connector/Python with Python 3.1 and later to establish Fabric connections, even though Fabric does not support Python 3.

Connector/Python support for Fabric comprises the following module and classes:

- Module `mysql.connector.fabric`: All classes, functions, and constants related to MySQL Fabric.
- Class `fabric.MySQLFabricConnection`: Similar to `MySQLConnection`, it creates a connection to the MySQL server based on application-provided information.

- Class `fabric.Fabric`: Manages the connection with a MySQL Fabric node; used by `MySQLFabricConnection`.
- Other helper classes for caching information.

## 8.1 Installing Connector/Python with MySQL Fabric Support

Fabric support in Connector/Python requires version 1.2.0 or greater. Downloads are available at <http://dev.mysql.com/downloads/connector/python/> in various packages such as Zip archives, compressed tar archives, RPM packages, Debian packages, and Windows Installer packages.

Using the compressed tar package, you can install MySQL Connector/Python as follows:

```
shell> tar xzf mysql-connector-python-1.2.3.tar.gz
shell> cd mysql-connector-python-1.2.3
shell> python setup.py install
```

For more information, see [Connector/Python Installation](#).

## 8.2 Requesting a Fabric Connection

The modules related to Fabric are located under `mysql.connector.fabric`. Importing `fabric` provides access to everything needed to use Fabric:

```
import mysql.connector
from mysql.connector import fabric
```

Traditionally, a MySQL connection is set up using the `mysql.connector.connect()` method using the connection arguments described at [Connector/Python Connection Arguments](#), and the connection is established immediately.

A request for a Fabric connection, by contrast, does not immediately connect. Instead, pass a `fabric` argument to the `connect()` call. This argument must be a dictionary. When Fabric connects to the MySQL server that it provides, it uses the connection arguments other than the `fabric` argument (except that the `unix_socket` connection argument is not supported).

To prepare a connection with Fabric, do this:

```
fabric_config = {
    'host': 'fabric.example.com',
}
fcnx = mysql.connector.connect(fabric=fabric_config, user='webuser',
                               password='webpass', database='employees')
```

If you prefer to pass a dictionary to `connect()`, do this:

```
config = {
    'fabric': {
        'host': 'fabric.example.com',
    },
    'user': 'webuser',
    'password': 'webpass',
    'database': 'employees',
}
fcnx = mysql.connector.connect(**config)
```

The `fabric` dictionary argument permits these values:



- `host`: The host to connect to (required).
- `port`: The TCP/IP port number to connect to on the specified host (optional; default 32274).
- `username`: The user name of the account to use (optional).
- `password`: The password of the account to use (optional).
- `connect_attempts`: The number of connection attempts to make before giving up (optional; default 3).
- `connect_delay`: The delay in seconds between attempts (optional; default 1).
- `report_errors`: Whether to report errors to Fabric while accessing a MySQL instance. (optional; default `False`).
- `ssl_ca`: The file containing the SSL certificate authority (optional).
- `ssl_cert`: The file containing the SSL certificate file (optional).
- `ssl_key`: The file containing the SSL key (optional).
- `protocol`: The connection protocol to use (optional; default `xmlrpc`). Permitted values are `xmlrpc` (use XML-RPC protocol) and `mysql` (use MySQL client/server protocol). If a value of `mysql` is specified, the default port becomes 32275, although that can be changed with an explicit `port` value.

The `username`, `password`, `report_errors`, `ssl_ca`, `ssl_cert` and `ssl_key` options were added in Connector/Python 1.2.1. It is possible to establish an SSL connection using only the `ssl_ca` argument. The `ssl_key` and `ssl_cert` arguments are optional. However, when either is given, both must be given or an `AttributeError` is raised.

The `protocol` option was added in Connector/Python2 2.1.2.

It is also possible to pass a `Fabric()` object instance as the `fabric` argument:

```
fabric_config = {
    'host': 'fabric.example.com',
}
fabinst = Fabric(**fabric_config)
fcnx = mysql.connector.connect(fabric=fabinst, user='webuser',
                               password='webpass', database='employees')
```

Or:

```
fabric_config = {
    'host': 'fabric.example.com',
}
fabinst = Fabric(**fabric_config)
config = {
    'fabric': fabinst,
    'user': 'webuser',
    'password': 'webpass',
    'database': 'employees',
}
fcnx = mysql.connector.connect(**config)
```

Once a `Fabric()` object is used, it is cached and reused.

Another (less preferred) way to establish a Fabric connection is pass configuration information to the `MySQLFabricConnection` class to create a connection with a Fabric node. This is similar to using the

`mysql.connector.connect()` method or `MySQLConnection()` class with the addition of the required `fabric` argument:

```
config = {
    'fabric': {
        'host': 'fabric.example.com',
    },
    'user': 'webuser',
    'password': 'webpass',
}

fcnx = fabric.MySQLFabricConnection(**config)
```

## Error Reporting

Connector/Python can report errors to Fabric that occur while accessing a MySQL instance. The information can be used to update the backing store and trigger a failover operation, provided that the instance is a primary server and Fabric has received a sufficient number of problem reports from different connectors.

- The `fabric` dictionary argument to the `connect()` method accepts a `report_errors` value. Its default value is `False`; pass a value of `True` to enable error reporting to Fabric.
- To define which errors to report, use the `extra_failure_report()` function:

```
from mysql.connector.fabric import extra_failure_report
extra_failure_report([error_code_0, error_code_1, ...])
```

## 8.3 Providing Information to Choose a MySQL Server

If you create a Fabric connection without providing any information about which data to access, the connection cannot function. To access a database, you must provide the driver with either of these types of information:

- The name of a *high-availability group* known by the MySQL Fabric instance to which you've connected. In such a group, one server is the master (the primary) and the others are slaves (secondaries).
- A *shard table*, and optionally a *shard key*, to guide Connector/Python in selecting a high-availability group.

The following discussion describes both ways of providing information. You do this by setting one or more properties of the Fabric connection object using its `set_property()` method, so the discussion begins by setting forth the sharding-related properties.. In the examples, `fcnx` represents the Fabric connection object, created as shown in [Section 8.2, “Requesting a Fabric Connection”](#).



### Note

`set_property()` does not connect. The connection is opened when a cursor is requested from the Fabric connection object or when its `cmd_query()` or `cmd_query_iter()` method is invoked.

These `set_property()` arguments are shard-related:

- `group`: A high-availability group name
- `tables`: The sharding table or tables

- `mode`: Whether operations are read/write or read only
- `scope`: Whether operations are local or global
- `key`: The key that identifies which row to affect

`group` and `tables` are mutually exclusive, so you specify only one of them. Applicability of the remaining arguments depends on which of `group` or `tables` you use:

If you specify `group`:

- `mode` is optional. The default is `fabric.MODE_READWRITE` if this property is omitted.
- `scope` is inapplicable. Do not specify it.
- `key` is inapplicable. Do not specify it.

If you specify `tables`:

- `mode` is optional. The default is `fabric.MODE_READWRITE` if this property is omitted.
- `scope` is optional. The default is `fabric.SCOPE_LOCAL` if this property is omitted.
- `key`: If `scope` is `fabric.SCOPE_LOCAL`, `key` is required to indicate which row to affect. If `scope` is `fabric.SCOPE_GLOBAL`, `key` is inapplicable; do not specify it.

When the `mode` argument is applicable, these values are permitted:

- `fabric.MODE_READWRITE`: Connect to a master server. This is the default.
- `fabric.MODE_READONLY`: Connect to a slave if one is available, to the master otherwise. If there are multiple secondary MySQL servers, load balancing is used to obtain the server information.

When the `scope` argument is applicable, these values are permitted:

- `fabric.SCOPE_LOCAL`: Local operation that affects the row with a given key. This is the default.
- `fabric.SCOPE_GLOBAL`: Global operation that affects all rows.

Providing the name of a high-availability group specifies that we know exactly the set of database servers that with which to interact. To do this, set the `group` property using the `set_property()` method:

```
fcnx.set_property(group='myGroup')
```

Providing `shard` information avoids the need to choose a high-availability group manually and permits Connector/Python to do so based on information from the MySQL Fabric server.

Whether operations use `RANGE` or `HASH` is transparent to the user. The information is provided by Fabric and Connector/Python uses the correct mechanism automatically.

To specify `shard tables` and `shard keys`, use the `tables` and `key` attributes of the `set_property()` method.

The format of each `shard table` is usually given as `'db_name.tbl_name'`. Because one or more tables can be specified, the `tables` argument to `set_property()` is specified as a tuple or list:

```
fcnx.set_property(tables=['employees.employees'], key=40)
```

```
cur = fcnx.cursor()
# do operations for employee with emp_no 40
fcnx.close()
```

By default, operations occur in local scope, or the `scope` property can be given to specify local or global scope explicitly. For local operations (as in the preceding example), the `key` argument must be specified to indicate which row to use. For global operations, do not specify the `key` attribute because the operation is performed on all rows in the table:

```
fcnx.set_property(tables=['employees.employees'], scope=fabric.SCOPE_GLOBAL)
cur = fcnx.cursor()
cur.execute("UPDATE employees SET last_name = UPPER(last_name)")
cnx.commit()
fcnx.close()
```

The default mode is read/write, so the driver connects to the master. The `mode` property can be given to specify read/write or read-only mode explicitly:

```
fcnx.set_property(group='myGroup', mode=fabric.MODE_READWRITE)
cur = fcnx.cursor()
cur.execute("UPDATE employees SET last_name = UPPER(last_name)")
cnx.commit()
fcnx.close()
```

Applications for which read-only mode is sufficient can specify a `mode` attribute of `fabric.MODE_READONLY`. In this case, a connection is established to a slave if one is available, or to the master otherwise.

Connector/Python 2.0.1 or later supports `RANGE_STRING` and `RANGE_DATETIME` as sharding types. These are similar to the regular `RANGE` sharding type, but instead of an integer key, require a value of a different type:

- For `RANGE_STRING`, a UTF-8 encoded string key is required. For example:

```
cnx.set_property(tables=["employees.employees"],
                 key=u'employee_name', mode=fabric.MODE_READONLY)
```

Only Unicode strings are supported. Any other type given when using a shard defined using `RANGE_STRING` causes a `ValueError` to be raised.

- For `RANGE_DATETIME`, a datetime or date object key is required. For example, to get the shard which holds employees hired after the year 2000, you could do the following, with lower bounds set as "group1/1980-01-01, group2/2000-01-01":

```
cnx.set_property(tables=["employees.employees"],
                 key=datetime.date(2000, 1, 1), mode=fabric.MODE_READONLY)
```

If the lower bounds included a time, it would have been like this:

```
cnx.set_property(tables=["employees.employees"],
                 key=datetime.datetime(2000, 1, 1, 12, 0, 0),
                 mode=fabric.MODE_READONLY)
```

Only `datetime.datetime` and `datetime.date` values are supported. Any other type given when using a shard defined using `RANGE_DATETIME` causes a `ValueError` to be raised.

---

# Chapter 9 Using Connector/J with MySQL Fabric

## Table of Contents

9.1 Installing Connector/J with MySQL Fabric Support .....	79
9.2 Loading the Driver and Requesting a Fabric Connection .....	80
9.3 Providing Information to Choose a MySQL Server .....	80
9.4 MySQL Fabric Configuration for Running Samples .....	82
9.5 Running Tests .....	83
9.6 Running Demonstration Programs .....	84
9.7 A Complete Example: Working with Employee Data .....	84
9.8 How Connector/J Chooses a MySQL Server .....	88
9.9 Using Hibernate with MySQL Fabric .....	89
9.10 Connector/J Fabric Support Reference .....	92
9.10.1 Connection Properties .....	92
9.10.2 FabricMySQLConnection API .....	93

MySQL Fabric provides data distribution and high-availability features for a set of MySQL database servers.

Developers using Connector/J can take advantage of its features to work with a set of servers managed by MySQL Fabric. Connector/J supports the following MySQL Fabric capabilities:

- Automatic node selection based on application-provided *shard* information (table and key)
- Read/write splitting within a MySQL Fabric *server group*
- Reporting errors to the Fabric node as part of the distributed failure detector

The Fabric Client library for Java, which is included with Connector/J, is comprised of the following packages:

- `src/com/mysql/fabric/xmlrpc`: Classes for core implementation of the XML-RPC protocol
- `src/com/mysql/fabric`: Classes for interacting with the MySQL Fabric management system using the XML-RPC protocol
- `src/com/mysql/fabric/jdbc`: Classes for JDBC access to MySQL servers based on shard information
- `src/com/mysql/fabric/hibernate`: `FabricMultiTenantConnectionProvider.java` class enabling integration with Hibernate
- `testsuite/fabric`: JUnit tests
- `src/demo/fabric`: Usage samples

## 9.1 Installing Connector/J with MySQL Fabric Support

Fabric support is available in Connector/J 5.1.30 and later. Please refer to Connector/J documentation for installation instructions.

## 9.2 Loading the Driver and Requesting a Fabric Connection

When using Connector/J with MySQL Fabric, you must provide the host name and port of the MySQL Fabric server instead of the database server. The connection string must follow the same form it normally does, with the addition of the `fabric` keyword, as follows:

```
jdbc:mysql:fabric://fabrihost:32274/database
```

The user name and password provided to the connection are used for authentication with the individual database servers. Fabric authentication parameters can be given in the URL using the `fabricUsername` and `fabricPassword` properties. e.g.

```
jdbc:mysql:fabric://fabrihost:32274/database?fabricUsername=admin&fabricPassword=secret
```



### Note

If the username and password to authenticate to the Fabric node are omitted, no authentication is used. This should only be done when authentication has been disabled on the Fabric node.



### Note

If you are using Java 5, you must manually load the driver class before attempting to connect.

```
Class.forName("com.mysql.fabric.jdbc.Driver");
```

Connection now proceeds normally.

```
Connection conn = DriverManager.getConnection(
    "jdbc:mysql:fabric://fabrihost:32274/database",
    user,
    password);
```

To use the Connector/J APIs that support MySQL Fabric, you must cast the `Connection` object to the public interface that provides the necessary methods. The interfaces are:

- `com.mysql.fabric.jdbc.FabricMySQLConnection`: JDBC3 interface, compatible with Java 5 and later
- `com.mysql.fabric.jdbc.JDBC4FabricMySQLConnection`: JDBC4 interface, compatible with Java 6 and later. This interface must be used to access JDBC4 APIs on the connection.

## 9.3 Providing Information to Choose a MySQL Server

If you create a Fabric connection without providing any information about which data to access, the connection cannot function. To access a database, you must provide the driver with one of the following:

- The name of a high-availability group known by the MySQL Fabric instance to which you've connected. In such a group, one server is the master (the primary) and the others are slaves (secondaries).
- A shard table, and optionally a shard key, to guide Connector/J in selecting a high-availability group.
- One or more query tables to guide the connector in selecting a server group. Query tables can reference only a single shard mapping or the query is rejected. A shard mapping can include several tables but they must be sharded on the same index.

The following discussion describes both ways of providing information. In the examples, `conn` represents the Fabric connection object, created as shown in [Section 9.2, “Loading the Driver and Requesting a Fabric Connection”](#).

Providing the name of a high-availability group specifies that we know exactly the set of database servers with which to interact. We can do this in two ways.

- The simplest method is to include the name of the server group in the connection string. This is useful if a connection needs to access data only in that server group. It is also possible to set the name of the server group in this way initially and to change it programmatically later.

```
// provide the server group as a connection property
Connection conn = DriverManager.getConnection(
    "jdbc:mysql:fabric://fabrichost:32274/database?fabricServerGroup=myGroup");
```

- If we connect without specifying a server group, or want to change it later, we can use the `JDBC4FabricMySQLConnection` interface to set the **server group name**.

```
JDBC4FabricMySQLConnection conn;
// connection initialization here
conn.setServerGroupName("myGroup");
```

Providing shard information avoids the need to choose a high-availability group manually and permits Connector/J to do so based on information from the MySQL Fabric server.

- Shard tables and shard keys can also be given as connection properties if desirable. Here we say that we want to access the `employees` table. The driver chooses a server group based on this shard table.



#### Note

The format of the shard table is usually given as `db_name.tbl_name`.

```
// provide the shard table as a connection property
Connection conn = DriverManager.getConnection(
    "jdbc:mysql:fabric://fabrichost:32274/database?fabricShardTable=employees.employees");
```

- Alternatively, query tables can be provided to the connection before creating a statement. The following example sets up the connection to perform a join between the `employees` and `departments` tables. Details on how to provide the shard key are given in the next step.

```
JDBC4FabricMySQLConnection conn;
// provide the set of query tables to the connection
conn.addQueryTable("departments");
conn.addQueryTable("employees");
```

- In many cases, you want to work with different sets of data at different times. You can specify the shard key to change the set of data to be accessible.

```
JDBC4FabricMySQLConnection conn;
// connection initialization here
conn.setShardKey("40"); // work with data related to shard key = 40
```

In summary, it is necessary to provide the name of a server group or a shard table and possibly shard key to access a database server.

## 9.4 MySQL Fabric Configuration for Running Samples

To run JUnit tests from `testsuite/fabric` or demonstration examples from `src/demo/fabric`, you must configure the MySQL Fabric test environment as follows.

### 1. Set up MySQL servers.

- `mysql-fabric-config`: The backing store for Fabric configuration. Used internally by Fabric to store the server list, shard mappings, and so forth. You set up this server instance during the MySQL Fabric setup procedure.
- `mysql-global`: The only server in the “global” group. Used to send DDL commands and update any data not in the shard data set.
  - Default location: 127.0.0.1:3401
  - Config properties: `com.mysql.fabric.testsuite.global.host`,  
`com.mysql.fabric.testsuite.global.port`
- `mysql-shard1`: First shard of sharded data set
  - Default location: 127.0.0.1:3402
  - Config properties: `com.mysql.fabric.testsuite.shard1.host`,  
`com.mysql.fabric.testsuite.shard1.port`
- `mysql-shard2`: Second shard of sharded data set
  - Default location: 127.0.0.1:3403
  - Config properties: `com.mysql.fabric.testsuite.shard2.host`,  
`com.mysql.fabric.testsuite.shard2.port`

All except `mysql-fabric-config` should have `server-id` set to a distinct value and the following entries added to `my.cnf`:

```
log-bin = mysql-bin
log-slave-updates = true
enforce-gtid-consistency = true
gtid-mode = on
```

### 2. Set up sharding. The user name and password of the account used to manage Fabric ([Section 2.3.1, “Create the Associated MySQL Users”](#)) must be specified in Fabric’s configuration file ([Section 2.3.2, “Configuration File”](#)).

- Create the global group:

```
shell> mysqlfabric group create fabric_test1_global
shell> mysqlfabric group add fabric_test1_global 127.0.0.1:3401
shell> mysqlfabric group promote fabric_test1_global
```

- Create shard groups:



```

shell> mysqlfabric group create fabric_test1_shard1
shell> mysqlfabric group add fabric_test1_shard1 127.0.0.1:3402
shell> mysqlfabric group promote fabric_test1_shard1

shell> mysqlfabric group create fabric_test1_shard2
shell> mysqlfabric group add fabric_test1_shard2 127.0.0.1:3403
shell> mysqlfabric group promote fabric_test1_shard2

```

- Create the sharding definition:

```
shell> mysqlfabric sharding create_definition RANGE fabric_test1_global
```

Notice the return value in the command output; for example, return = 1. This is the `$MAPPING_ID` used in the following commands

```
shell> mysqlfabric sharding add_table $MAPPING_ID employees.employees emp_no
```

- Create the shard index:

```
shell> mysqlfabric sharding add_shard $MAPPING_ID \
    fabric_test1_shard1/0,fabric_test1_shard2/10000 --state=ENABLED
```

## 9.5 Running Tests

The `test-fabric` target in the `build.xml` file runs JUnit tests on these servers. It requires only the setup described in [Section 9.4, “MySQL Fabric Configuration for Running Samples”](#). All necessary tables and data are created during test run.

The parameters for the servers must be provided to Ant to verify that the correct information is received from the Fabric node. This includes server host names and ports. This data can be provided on the command line with `-D` arguments to Ant or in a `build.properties` file. This file should be placed in the root of the source directory, where `build.xml` is located. Based on the information given so far, this file would contain the following entries:

```

com.mysql.fabric.testsuite.hostname=localhost
com.mysql.fabric.testsuite.port=32274

com.mysql.fabric.testsuite.fabricUsername=admin
com.mysql.fabric.testsuite.fabricPassword=secret

com.mysql.fabric.testsuite.username=root
com.mysql.fabric.testsuite.password=
com.mysql.fabric.testsuite.database=employees
com.mysql.fabric.testsuite.global.host=127.0.0.1
com.mysql.fabric.testsuite.global.port=3401
com.mysql.fabric.testsuite.shard1.host=127.0.0.1
com.mysql.fabric.testsuite.shard1.port=3402
com.mysql.fabric.testsuite.shard2.host=127.0.0.1
com.mysql.fabric.testsuite.shard2.port=3403

```

Sample Ant calls are shown below. If the parameters are specified in your `build.properties` file, it is not necessary to include them on the command line.

```

shell> JAVA_HOME=/opt/jdk1.5/ ant \
    -Dcom.mysql.fabric.testsuite.password=pwd \
    -Dcom.mysql.fabric.testsuite.global.port=3401 \
    -Dcom.mysql.fabric.testsuite.shard1.port=3402 \
    -Dcom.mysql.fabric.testsuite.shard2.port=3403 \

```

```
test-fabric
```

## 9.6 Running Demonstration Programs

To run the demo programs, you must set up the MySQL Fabric environment as described in [Section 9.4, “MySQL Fabric Configuration for Running Samples”](#). After that, you can use the `demo-fabric-*` Ant targets.

```
shell> JAVA_HOME=/opt/jdk1.5/ ant \
-Dcom.mysql.fabric.testsuite.password=pwd \
demo-fabric
```

These targets invoke all demo programs except Hibernate demos. You should use the `demo-fabric-hibernate` target instead:

```
shell> JAVA_HOME=/opt/jdk1.6/ ant \
-Dcom.mysql.fabric.testsuite.password=pwd \
demo-fabric-hibernate
```



### Note

You need Java 6+ to use Hibernate Fabric integration. Multi-tenancy is a feature specific to Hibernate 4 and higher which requires Java 6+. That is why we do not provide the `FabricMultiTenantConnectionProvider` class or related demos compatible with Java 5.

## 9.7 A Complete Example: Working with Employee Data

This document demonstrates two possible ways of working with sharded data relating to employees. To run this program, you must set up a shard mapping for the employees table in MySQL Fabric as described in [Section 9.4, “MySQL Fabric Configuration for Running Samples”](#).

This code can be found in the distribution package in `src/demo/fabric/EmployeesJdbc.java`.

```
/*
 Copyright (c) 2013, 2014, Oracle and/or its affiliates. All rights reserved.

 The MySQL Connector/J is licensed under the terms of the GPLv2
 <http://www.gnu.org/licenses/old-licenses/gpl-2.0.html>, like most MySQL Connectors.
 There are special exceptions to the terms and conditions of the GPLv2 as it is applied to
 this software, see the FLOSS License Exception
 <http://www.mysql.com/about/legal/licensing/foss-exception.html>.

 This program is free software; you can redistribute it and/or modify it under the terms
 of the GNU General Public License as published by the Free Software Foundation; version 2
 of the License.

 This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY;
 without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.
 See the GNU General Public License for more details.

 You should have received a copy of the GNU General Public License along with this
 program; if not, write to the Free Software Foundation, Inc., 51 Franklin St, Fifth
 Floor, Boston, MA 02110-1301 USA

 */

package demo.fabric;

import java.sql.Connection;
```

```

import java.sql.DriverManager;
import java.sql.PreparedStatement;
import java.sql.Statement;
import java.sql.ResultSet;

import com.mysql.fabric.jdbc.FabricMySQLConnection;

/**
 * Demonstrate working with employee data in MySQL Fabric with Connector/J and the JDBC APIs.
 */
public class EmployeesJdbc {
    public static void main(String args[]) throws Exception {

        String hostname = System.getProperty("com.mysql.fabric.testsuite.hostname");
        String port = System.getProperty("com.mysql.fabric.testsuite.port");
        String database = System.getProperty("com.mysql.fabric.testsuite.database");
        String user = System.getProperty("com.mysql.fabric.testsuite.username");
        String password = System.getProperty("com.mysql.fabric.testsuite.password");

        String baseUrl = "jdbc:mysql:fabric://" + hostname + ":" + Integer.valueOf(port) + "/";

        // Load the driver if running under Java 5
        if (!com.mysql.jdbc.Util.isJdbc4()) {
            Class.forName("com.mysql.fabric.jdbc.FabricMySQLDriver");
        }

        // 1. Create database and table for our demo
        Connection rawConnection = DriverManager.getConnection(
            baseUrl + "mysql?fabricServerGroup=fabric_test1_global",
            user,
            password);
        Statement statement = rawConnection.createStatement();
        statement.executeUpdate("create database if not exists employees");
        statement.close();
        rawConnection.close();

        // We should connect to the global group to run DDL statements,
        // they will be replicated to the server groups for all shards.

        // The 1-st way is to set its name explicitly via the
        // "fabricServerGroup" connection property
        rawConnection = DriverManager.getConnection(
            baseUrl + database + "?fabricServerGroup=fabric_test1_global",
            user,
            password);
        statement = rawConnection.createStatement();
        statement.executeUpdate("create database if not exists employees");
        statement.close();
        rawConnection.close();

        // The 2-nd way is to get implicitly connected to global group
        // when the shard key isn't provided, ie. set "fabricShardTable"
        // connection property but don't set "fabricShardKey"
        rawConnection = DriverManager.getConnection(
            baseUrl + "employees" + "?fabricShardTable=employees.employees",
            user,
            password);
        // At this point, we have a connection to the global group for
        // the 'employees.employees' shard mapping.
        statement = rawConnection.createStatement();
        statement.executeUpdate("drop table if exists employees");
        statement.executeUpdate("create table employees (emp_no int not null," +
            "first_name varchar(50), last_name varchar(50)," +
            "primary key (emp_no))");

        // 2. Insert data
    }
}

```

```

// Cast to a Fabric connection to have access to specific methods
FabricMySQLConnection connection = (FabricMySQLConnection)rawConnection;

// example data used to create employee records
Integer ids[] = new Integer[] {1, 2, 10001, 10002};
String firstNames[] = new String[] {"John", "Jane", "Andy", "Alice"};
String lastNames[] = new String[] {"Doe", "Doe", "Wiley", "Wein"};

// insert employee data
PreparedStatement ps = connection.prepareStatement(
    "INSERT INTO employees.employees VALUES (?, ?, ?)");
for (int i = 0; i < 4; ++i) {
    // choose the shard that handles the data we interested in
    connection.setShardKey(ids[i].toString());

    // perform insert in standard fashion
    ps.setInt(1, ids[i]);
    ps.setString(2, firstNames[i]);
    ps.setString(3, lastNames[i]);
    ps.executeUpdate();
}

// 3. Query the data from employees
System.out.println("Querying employees");
System.out.format("%7s | %-30s | %-30s%n", "emp_no", "first_name", "last_name");
System.out.println("-----+-----+-----");
ps = connection.prepareStatement(
    "select emp_no, first_name, last_name from employees where emp_no = ?");
for (int i = 0; i < 4; ++i) {

    // we need to specify the shard key before accessing the data
    connection.setShardKey(ids[i].toString());

    ps.setInt(1, ids[i]);
    ResultSet rs = ps.executeQuery();
    rs.next();
    System.out.format("%7d | %-30s | %-30s%n", rs.getInt(1), rs.getString(2), rs.getString(3));
    rs.close();
}
ps.close();

// 4. Connect to the global group and clean up
connection.setServerGroupName("fabric_test1_global");
statement.executeUpdate("drop table if exists employees");
statement.close();
connection.close();
}
}

```

Here is an alternative using the DataSource API:

```

/*
Copyright (c) 2015, Oracle and/or its affiliates. All rights reserved.

The MySQL Connector/J is licensed under the terms of the GPLv2
<http://www.gnu.org/licenses/old-licenses/gpl-2.0.html>, like most MySQL Connectors.
There are special exceptions to the terms and conditions of the GPLv2 as it is applied to
this software, see the FOSS License Exception
<http://www.mysql.com/about/legal/licensing/foss-exception.html>.

This program is free software; you can redistribute it and/or modify it under the terms
of the GNU General Public License as published by the Free Software Foundation; version 2
of the License.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY;

```

without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

```
*/
package demo.fabric;

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.PreparedStatement;
import java.sql.ResultSet;
import java.sql.Statement;

import com.mysql.fabric.jdbc.FabricMySQLConnection;
import com.mysql.fabric.jdbc.FabricMySQLDataSource;

/**
 * Demonstrate working with employee data in MySQL Fabric with Connector/J and the JDBC APIs via a DataSource
 */
public class EmployeesDataSource {
    public static void main(String args[]) throws Exception {

        String hostname = System.getProperty("com.mysql.fabric.testsuite.hostname");
        String port = System.getProperty("com.mysql.fabric.testsuite.port");
        String database = System.getProperty("com.mysql.fabric.testsuite.database");
        // credentials to authenticate with the SQL nodes
        String user = System.getProperty("com.mysql.fabric.testsuite.username");
        String password = System.getProperty("com.mysql.fabric.testsuite.password");
        // credentials to authenticate to the Fabric node
        String fabricUsername = System.getProperty("com.mysql.fabric.testsuite.fabricUsername");
        String fabricPassword = System.getProperty("com.mysql.fabric.testsuite.fabricPassword");

        // setup the Fabric datasource to create connections
        FabricMySQLDataSource ds = new FabricMySQLDataSource();
        ds.setServerName(hostname);
        ds.setPort(Integer.valueOf(port));
        ds.setDatabaseName(database);
        ds.setFabricUsername(fabricUsername);
        ds.setFabricPassword(fabricPassword);

        // Load the driver if running under Java 5
        if (!com.mysql.jdbc.Util.isJdbc4()) {
            Class.forName("com.mysql.fabric.jdbc.FabricMySQLDriver");
        }

        // 1. Create database and table for our demo
        ds.setDatabaseName("mysql"); // connect to the `mysql` database before creating our `employees` database
        ds.setFabricServerGroup("fabric_test1_global"); // connect to the global group
        Connection rawConnection = ds.getConnection(user, password);
        Statement statement = rawConnection.createStatement();
        statement.executeUpdate("create database if not exists employees");
        statement.close();
        rawConnection.close();

        // We should connect to the global group to run DDL statements, they will be replicated to the server

        // The 1-st way is to set its name explicitly via the "fabricServerGroup" datasource property
        ds.setFabricServerGroup("fabric_test1_global");
        rawConnection = ds.getConnection(user, password);
        statement = rawConnection.createStatement();
        statement.executeUpdate("create database if not exists employees");
        statement.close();
        rawConnection.close();
    }
}
```

```
// The 2-nd way is to get implicitly connected to global group when the shard key isn't provided, ie.
// don't set "fabricShardKey"
ds.setFabricServerGroup(null); // clear the setting in the datasource for previous connections
ds.setFabricShardTable("employee.employees");
rawConnection = ds.getConnection(user, password);
// At this point, we have a connection to the global group for the 'employees.employees' shard mapping
statement = rawConnection.createStatement();
statement.executeUpdate("drop table if exists employees");
statement.executeUpdate("create table employees (emp_no int not null, first_name varchar(50), last_name varchar(50))");

// 2. Insert data

// Cast to a Fabric connection to have access to Fabric-specific methods
FabricMySQLConnection connection = (FabricMySQLConnection) rawConnection;

// example data used to create employee records
Integer ids[] = new Integer[] { 1, 2, 10001, 10002 };
String firstNames[] = new String[] { "John", "Jane", "Andy", "Alice" };
String lastNames[] = new String[] { "Doe", "Doe", "Wiley", "Wein" };

// insert employee data
PreparedStatement ps = connection.prepareStatement("INSERT INTO employees.employees VALUES (?, ?, ?)");
for (int i = 0; i < 4; ++i) {
    // choose the shard that handles the data we interested in
    connection.setShardKey(ids[i].toString());

    // perform insert in standard fashion
    ps.setInt(1, ids[i]);
    ps.setString(2, firstNames[i]);
    ps.setString(3, lastNames[i]);
    ps.executeUpdate();
}

// 3. Query the data from employees
System.out.println("Querying employees");
System.out.format("%7s | %-30s | %-30s\n", "emp_no", "first_name", "last_name");
System.out.println("-----+-----+-----");
ps = connection.prepareStatement("select emp_no, first_name, last_name from employees where emp_no = ?");
for (int i = 0; i < 4; ++i) {

    // we need to specify the shard key before accessing the data
    connection.setShardKey(ids[i].toString());

    ps.setInt(1, ids[i]);
    ResultSet rs = ps.executeQuery();
    rs.next();
    System.out.format("%7d | %-30s | %-30s\n", rs.getInt(1), rs.getString(2), rs.getString(3));
    rs.close();
}
ps.close();

// 4. Connect to the global group and clean up
connection.setServerGroupName("fabric_test1_global");
statement.executeUpdate("drop table if exists employees");
statement.close();
connection.close();
}
}
```

## 9.8 How Connector/J Chooses a MySQL Server

Before a database server can be chosen, a server group must be chosen. The following values are taken into account:

- *server group* name: If a server group name is specified directly, it is used.
- *shard table*: If a shard table is given, the global group for the shard mapping governing the given shard table is used.
- *shard table + shard key*: If a shard key is given the shard mapping is used to determine which server group handles data for the shard key.
- *query tables*: If the query table set contains a sharded table, the shard mapping for that table is used.

The server group name can be accessed by `con.getCurrentServerGroup().getName()`.

Once a server group is chosen, an individual database server is chosen based on the read-only state of the connection. A read-only server is chosen if one is available. Otherwise a read-write server is chosen. The server *weight* is not currently taken into account.

None of these values can be changed while a transaction is in progress.

## 9.9 Using Hibernate with MySQL Fabric

It is possible to use [Hibernate 4's multi-tenancy support](#) to work with a set of database servers managed by MySQL Fabric.

### APIs necessary to implement MultiTenantConnectionProvider

We can use internal APIs included with Connector/J with MySQL Fabric support to implement Hibernate's [MultiTenantConnectionProvider](#).

The following implementation is included in the package as the `com.mysql.fabric.hibernate.FabricMultiTenantConnectionProvider` class. An example of how to use it is included as the class `demo.fabric.HibernateFabric`.

To implement [MultiTenantConnectionProvider](#), we use the `com.mysql.fabric.FabricConnection` class. This class connects to the MySQL Fabric manager to obtain information about servers and data sharding. This is an internal API and subject to change. The following [FabricConnection](#) methods can be used:

- `FabricConnection(String url, String username, String password)` throws `FabricCommunicationException`

Construct a new instance of a MySQL Fabric client and initiate the connection.

- `ServerGroup getServerGroup(String serverGroupName)` throws `FabricCommunicationException`

Retrieve an object representing the named server group. This includes the list of servers in the group including their mode (read-only or read-write).

- `ShardMapping getShardMapping(String database, String table)` throws `FabricCommunicationException`

Retrieve an object representing a shard mapping for the given database and table. The [ShardMapping](#) indicates the global group and the individual shards.

The following additional methods are used:

- `Set<Server> ServerGroup.getServers()`

Return the servers in this group.

- `String Server.getHostname()`

Return the server host name.

- `int Server.getPort()`

Return the server port.

## Implementing MultiTenantConnectionProvider

To begin, we declare the class with members to keep necessary information for the connection and the constructor:

```
public class FabricMultiTenantConnectionProvider implements MultiTenantConnectionProvider {
    // a connection to the MySQL Fabric manager
    private FabricConnection fabricConnection;
    // the database and table of the sharded data
    private String database;
    private String table;
    // user and password for Fabric manager and MySQL servers
    private String user;
    private String password;
    // shard mapping for `database.table`
    private ShardMapping shardMapping;
    // global group for the shard mapping
    private ServerGroup globalGroup;

    public FabricMultiTenantConnectionProvider(
        String fabricUrl, String database, String table,
        String user, String password) {
        try {
            this.fabricConnection = new FabricConnection(fabricUrl, user, password);
            this.database = database;
            this.table = table;
            this.user = user;
            this.password = password;
            // eagerly retrieve the shard mapping and server group from the Fabric manager
            this.shardMapping = this.fabricConnection.getShardMapping(this.database, this.table);
            this.globalGroup = this.fabricConnection.
                getServerGroup(this.shardMapping.getGlobalGroupName());
        } catch (FabricCommunicationException ex) {
            throw new RuntimeException(ex);
        }
    }
}
```

Next, create a method to create connections:

```
/**
 * Find a server with mode READ_WRITE in the given server group and create a JDBC connection to it.
 *
 * @returns a {@link Connection} to an arbitrary MySQL server
 * @throws SQLException if connection fails or a READ_WRITE server is not contained in the group
 */
private Connection getReadWriteConnectionFromServerGroup(ServerGroup serverGroup)
    throws SQLException {
    // iterate the list of servers in the given group until we find a r/w server
    for (Server s : serverGroup.getServers()) {
        if (ServerMode.READ_WRITE.equals(s.getMode())) {
            // create a connection to the server using vanilla JDBC

```



```
        String jdbcUrl = String.format("jdbc:mysql://%s:%s/%s",
            s.getHostname(), s.getPort(), this.database);
        return DriverManager.getConnection(jdbcUrl, this.user, this.password);
    }
}
// throw an exception if we are unable to make the connection
throw new SQLException(
    "Unable to find r/w server for chosen shard mapping in group " + serverGroup.getName());
}
```

To implement the interface, the following methods must be implemented:

- `Connection getAnyConnection()` throws `SQLException`

This method should obtain a connection to the global group. We can implement it like this:

```
/**
 * Get a connection that be used to access data or metadata not specific to any shard/tenant.
 * The returned connection is a READ_WRITE connection to the global group of the shard mapping
 * for the database and table association with this connection provider.
 */
public Connection getAnyConnection() throws SQLException {
    return getReadWriteConnectionFromServerGroup(this.globalGroup);
}
```

- `Connection getConnection(String tenantIdentifier)` throws `SQLException`

This method must use the `tenantIdentifier` to determine which server to access. We can look up the `ServerGroup` from the `ShardMapping` like this:

```
/**
 * Get a connection to access data association with the provided `tenantIdentifier' (or shard
 * key in Fabric-speak). The returned connection is a READ_WRITE connection.
 */
public Connection getConnection(String tenantIdentifier) throws SQLException {
    String serverGroupName = this.shardMapping.getGroupNameForKey(tenantIdentifier);
    try {
        ServerGroup serverGroup = this.fabricConnection.getServerGroup(serverGroupName);
        return getReadWriteConnectionFromServerGroup(serverGroup);
    } catch (FabricCommunicationException ex) {
        throw new RuntimeException(ex);
    }
}
```

Finally, our trivial implementation to release connections:

```
/**
 * Release a non-shard-specific connection.
 */
public void releaseAnyConnection(Connection connection) throws SQLException {
    connection.close();
}

/**
 * Release a connection specific to `tenantIdentifier'.
 */
public void releaseConnection(String tenantIdentifier, Connection connection)
    throws SQLException {
    releaseAnyConnection(connection);
}

/**
 * We don't track connections.
 */
```

```

* @returns false
*/
public boolean supportsAggressiveRelease() {
    return false;
}

```

And finally to implement the [Wrapped](#) role:

```

public boolean isUnwrappableAs(Class unwrapType) {
    return false;
}

public <T> T unwrap(Class<T> unwrapType) {
    return null;
}

```

## Using a custom MultiTenantConnectionProvider

The [SessionFactory](#) can be created like this:

```

// create a new instance of our custom connection provider supporting MySQL Fabric
FabricMultiTenantConnectionProvider connProvider =
    new FabricMultiTenantConnectionProvider(
        fabricUrl, "employees", "employees", username, password);
// create a service registry with the connection provider to construct the session factory
ServiceRegistryBuilder srb = new ServiceRegistryBuilder();
srb.addService(
    org.hibernate.service.jdbc.connections.spi.MultiTenantConnectionProvider.class,
    connProvider);
srb.applySetting("hibernate.dialect", "org.hibernate.dialect.MySQLInnoDBDialect");

// create the configuration and build the session factory
Configuration config = new Configuration();
config.setProperty("hibernate.multiTenancy", "DATABASE");
config.addResource("com/mysql/fabric/demo/employee.hbm.xml");
return config.buildSessionFactory(srb.buildServiceRegistry());

```

## Using Hibernate multi-tenancy

Once you have created a [SessionFactory](#) with your custom [MultiTenantConnectionProvider](#), it is simple to use. Provide the [shard key](#) to the [SessionFactory](#) when creating the session:

```

// access data related to shard key = 40
Session session = sessionFactory.withOptions().tenantIdentifier("40").openSession();

```

Each [Session](#) is given a [shard key](#) (tenant identifier in Hibernate-speak) and uses it to obtain a connection to an appropriate server. This cannot be changed for the duration of the [Session](#).

## 9.10 Connector/J Fabric Support Reference

### 9.10.1 Connection Properties

The following connection properties are recognized by Connector/J for dealing with MySQL Fabric:

- [fabricShardKey](#)

The initial shard key used to determine which server group to send queries to. The [fabricShardTable](#) property must also be specified.

- `fabricShardTable`  
The initial shard mapping used to determine a server group to send queries to.
- `fabricServerGroup`  
The initial server group to direct queries to.
- `fabricProtocol`  
Protocol used to communicate with the Fabric node. XML-RPC over HTTP is currently the only supported protocol and is specified with a value of "http".
- `fabricUsername`  
Username used to authenticate with the Fabric node.
- `fabricPassword`  
Password used to authenticate with the Fabric node.
- `fabricReportErrors` (default=false)  
Determines whether or not errors are reported to Fabric's distributed failure detector. Only connection errors, those with an SQL state beginning with "08", are reported.

## 9.10.2 FabricMySQLConnection API

The following methods are available in the `com.mysql.fabric.jdbc.FabricMySQLConnection` interface.

- `void clearServerSelectionCriteria()`  
Clear all the state that is used to determine which server to send queries to.
- `void setShardKey(String shardKey) throws SQLException`  
Set the shard key for the data being accessed.
- `String getShardKey()`  
Get the shard key for the data being accessed.
- `void setShardTable(String shardTable) throws SQLException`  
Set the table being accessed. Can be a table name or a database and table name pair in the form `db_name.tbl_name`. The table must be known by Fabric as a sharded table.
- `String getShardTable()`  
Get the name of the table being accessed.
- `void setServerGroupName(String serverGroupName) throws SQLException`  
Set the server group name to connect to. Direct server group selection is mutually exclusive of sharded data access.
- `String getServerGroupName()`

Get the server group name when using direct server group selection.

- `ServerGroup getCurrentServerGroup()`

Get the current server group if sufficient server group selection has been provided. Otherwise null.

- `void clearQueryTables() throws SQLException`

Clear the list of tables for the last query. This also clears the shard mapping/table and must be given again for the next query via `setShardTable()` or `addQueryTable()`.

- `void addQueryTable(String tableName) throws SQLException`

Specify that the given table is intended to be used in the next query.

- `Set<String> getQueryTables()`

Get the list of tables intended to be used in the next query.

---

# Chapter 10 Using Connector/Net with MySQL Fabric

## Table of Contents

10.1 System Requirements .....	95
10.2 Set up the MySQL Fabric Plugin .....	95
10.3 Using MySQL Fabric Groups .....	97
10.4 Using Ranged Sharding .....	99

Connector/Net supports MySQL Fabric as a Replication/Load balancing plugin.



### Note

This feature was added in MySQL Connector/Net 6.9.4.

Support for this feature was removed in MySQL Connector/Net 6.10.2 with the introduction of InnoDB cluster, which combines various MySQL technologies to enable you to create highly available clusters of MySQL instances.

The following steps are required to use MySQL Fabric with Connector/Net:

## 10.1 System Requirements

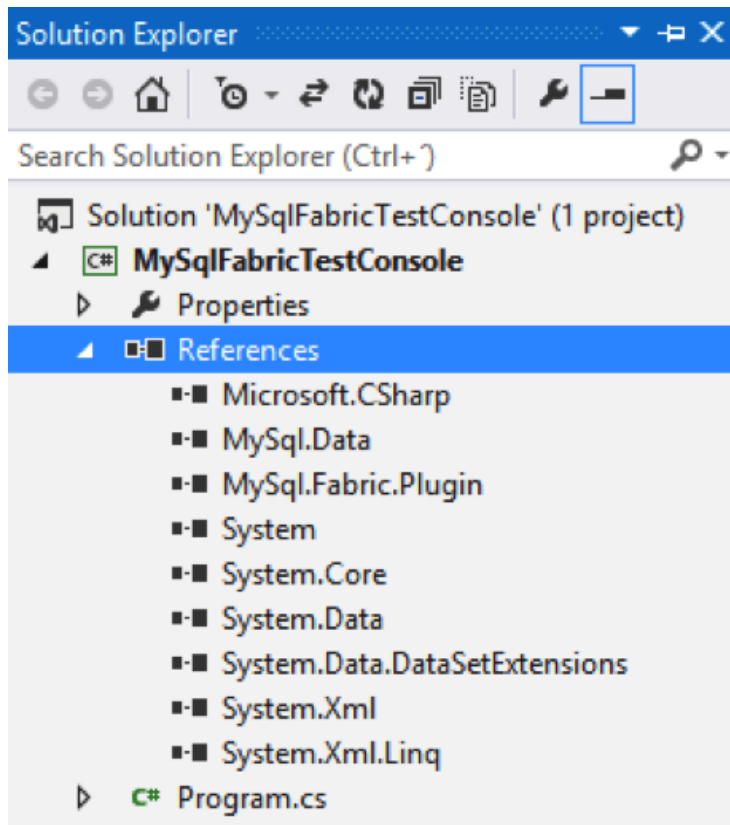
Confirm that you have the required Connector/Net and Fabric versions installed:

- Connector/Net 6.9.4 or newer
- MySQL Fabric 1.5.0 or newer

## 10.2 Set up the MySQL Fabric Plugin

First, add **MySql.Data** and **MySql.Fabric.Plugin** to the project references:

Figure 10.1 MySQL Fabric Project References



Second, add a configuration section with the Fabric connection to the `App.config` configuration file. For example:

```
<configuration>
  <configSections>
    <section name="MySQL" type="MySql.Data.MySqlClient.MySqlConfiguration, MySql.Data,
      Version=6.9.4.0, Culture=neutral, PublicKeyToken=c5687fc88969c44d"/>
  </configSections>
  <MySQL>
    <Replication>
      <ServerGroups>
        <Group name="Fabric" groupType="MySql.Fabric.FabricServerGroup, MySql.Fabric.Plugin">
          <Servers>
            <Server name="fabric" connectionString="server=localhost;port=32275;uid=admin;password=adminpass;"/>
          </Servers>
        </Group>
      </ServerGroups>
    </Replication>
  </MySQL>
</configuration>
```

Notice that the Fabric connection is set in the Server node:

```
<Server name="fabric" connectionString="server=localhost;port=32275;uid=admin;password=adminpass;"/>
```

Connector/Net only supports the MySQL protocol for connecting to MySQL Fabric, so the correct port must be used.

## 10.3 Using MySQL Fabric Groups

The MySQL Fabric group is used with a **MySQLConnection** that contains the server name specified in the `App.config` file, and a username and password for connecting to the servers defined in the group.

A Fabric extension method is used to specify the group and mode:

```
MySQLConnection conn = new MySQLConnection(connectionString);
conn.SetFabricProperties(groupId: "my_group", mode: FabricServerModeEnum.Read_Write);
```

The following example shows how to store and retrieve information in a specific Fabric group:



### Note

The initial MySQL Fabric configuration for this example is defined in the MySQL Fabric documentation at [Section 3.1, "Example: Fabric and Replication"](#).

```
using System;
using MySql.Data.MySqlClient;
using MySql.Fabric;

namespace FabricTest
{
    class Program
    {
        public const string connectionString = "server=fabric;uid=appuser;password=pass;";

        static void Main(string[] args)
        {
            RunFabricTest();
        }

        static string AddEmployee(MySqlConnection conn, int emp_no, string first_name, string last_name)
        {
            conn.SetFabricProperties(groupId: "my_group", mode: FabricServerModeEnum.Read_Write);

            MySqlCommand cmd = new MySqlCommand("USE employees", conn);
            cmd.ExecuteNonQuery();

            cmd.CommandText = "INSERT INTO employees VALUES (@emp_no, @first_name, @last_name)";
            cmd.Parameters.Add("emp_no", emp_no);
            cmd.Parameters.Add("first_name", first_name);
            cmd.Parameters.Add("last_name", last_name);
            cmd.ExecuteNonQuery();

            cmd.CommandText = "SELECT @@global.gtid_executed";
            cmd.Parameters.Clear();
            using (MySqlDataReader reader = cmd.ExecuteReader())
            {
                while (reader.Read())
                {
                    Console.WriteLine("Transactions executed on the master " + reader.GetValue(0));
                }
                return reader.GetString(0);
            }
        }

        static void FindEmployee(MySqlConnection conn, int emp_no, string gtid_executed)
        {
```

```

conn.SetFabricProperties(groupId: "my_group", mode: FabricServerModeEnum.Read_only);

MySQLCommand cmd = new MySQLCommand("", conn);
cmd.CommandText = "SELECT WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS(@gtid_executed, 0)";
cmd.Parameters.Add("gtid_executed", gtid_executed);
using (MySQLDataReader reader = cmd.ExecuteReader())
{
    while (reader.Read())
    {
        Console.WriteLine("Had to synchronize " + reader.GetValue(0) + " transactions.");
    }
}

cmd.CommandText = "USE employees";
cmd.Parameters.Clear();
cmd.ExecuteNonQuery();

cmd.CommandText = "SELECT first_name, last_name FROM employees ";
cmd.CommandText += " WHERE emp_no = @emp_no";
cmd.Parameters.Clear();
cmd.Parameters.Add("emp_no", emp_no);
using (MySQLDataReader reader = cmd.ExecuteReader())
{
    while (reader.Read())
    {
        object[] values = new object[reader.FieldCount];
        reader.GetValues(values);
        Console.WriteLine("Retrieved {0}", string.Join(", ", values));
    }
}
}

static void RunFabricTest()
{
    using (MySQLConnection conn = new MySQLConnection(connectionString))
    {
        string gtid_executed;
        conn.SetFabricProperties(groupId: "my_group", mode: FabricServerModeEnum.Read_Write);
        conn.Open();

        MySQLCommand cmd = new MySQLCommand("", conn);
        cmd.CommandText = "CREATE DATABASE IF NOT EXISTS employees;";
        cmd.ExecuteNonQuery();

        cmd.CommandText = "USE employees;";
        cmd.ExecuteNonQuery();

        cmd.CommandText = "DROP TABLE IF EXISTS employees;";
        cmd.ExecuteNonQuery();

        cmd.CommandText = "CREATE TABLE employees(";
        cmd.CommandText += "    emp_no INT, ";
        cmd.CommandText += "    first_name CHAR(40), ";
        cmd.CommandText += "    last_name CHAR(40)";
        cmd.CommandText += ");";
        cmd.ExecuteNonQuery();

        gtid_executed = AddEmployee(conn, 12, "John", "Doe");
        FindEmployee(conn, 12, gtid_executed);
    }
}
}
}

```



## 10.4 Using Ranged Sharding

Sharding with Connector/Net requires you to specify the table, key, and scope for each executed query.

```

MySQLConnection con = new MySQLConnection(connectionString);

con.SetFabricProperties(table: "employees.employees", key: empId.ToString(),
    mode: FabricServerModeEnum.Read_Write, scope: FabricScopeEnum.Local);

MySQLCommand cmd = new MySQLCommand(
    string.Format("insert into employees(emp_no, first_name, last_name) values ({0}, '{1}', '{2}']",
        empId, firstName, lastName), con);
cmd.ExecuteNonQuery();

```

You can use the following MySQL Fabric configuration to execute the code example:



### Note

For related MySQL Fabric documentation, see [Section 3.2.2, “Sharding Scenario”](#).

```

using System;
using System.Collections.Generic;
using MySql.Data.MySqlClient;
using MySql.Fabric;

namespace FabricTest
{
    class ShardTest
    {
        public const string connectionString = "server=fabric;uid=appuser;password=pass;";

        public static void test_shard_range()
        {
            using (MySQLConnection con = new MySQLConnection(connectionString))
            {
                con.SetFabricProperties(groupId: "group_id-global", mode: FabricServerModeEnum.Read_Write,
                    scope: FabricScopeEnum.Global);
                con.Open();

                MySQLCommand cmd = new MySQLCommand("create database if not exists employees", con);
                cmd.ExecuteNonQuery();

                cmd.CommandText = "use employees";
                cmd.ExecuteNonQuery();

                cmd.CommandText = "drop table if exists employees";
                cmd.ExecuteNonQuery();

                cmd.CommandText =
@"create table employees (
    emp_no int,
    first_name char( 40 ),
    last_name char( 40 )
)";
                cmd.ExecuteNonQuery();

                string gtid = prepare_synchronization(con);

                string[] first_names = { "John", "Buffalo", "Michael", "Kate", "Deep", "Genesis" };
                string[] last_names = { "Doe", "Bill", "Jackson", "Bush", "Purple" };
                List<int> list_emp_no = new List<int>();

```

```

con.SetFabricProperties(scope: FabricScopeEnum.Local);

for (int i = 0; i < 10; i++)
{
    int empId = pick_shard_key();
    list_emp_no.Add(empId);
    add_employee(con, empId, first_names[empId % first_names.Length], last_names[empId % last_names.Length]
}

for (int i = 0; i < list_emp_no.Count; i++)
{
    int empId = list_emp_no[ i ];
    find_employee(con, empId, gtid);
}
}

public static int pick_shard_key()
{
    Random r = new Random();
    int shard = r.Next(0, 2);
    int shard_range = shard * 10000;
    shard_range = (shard != 0) ? shard_range : shard_range + 1;
    int shift_within_shard = r.Next(0, 99999);
    return shard_range + shift_within_shard;
}

public static void add_employee(MySqlConnection con, int empId, string firstName, string lastName, string gtid)
{
    con.SetFabricProperties( table: "employees.employees", key: empId.ToString(), mode: FabricServerModeEnum.Read_write);
    synchronize(con, gtid);
    MySqlCommand cmd = new MySqlCommand(
        string.Format("insert into employees( emp_no, first_name, last_name ) values ( {0}, '{1}', '{2}' )",
            empId, firstName, lastName), con);
    cmd.ExecuteNonQuery();
}

public static void find_employee(MySqlConnection con, int empId, string gtid)
{
    con.SetFabricProperties(table: "employees.employees", key: empId.ToString(),
        mode: FabricServerModeEnum.Read_only);
    synchronize(con, gtid);
    MySqlCommand cmd = new MySqlCommand(string.Format("
        select first_name, last_name from employees where emp_no = {0}", empId), con);
    using (MySqlDataReader r = cmd.ExecuteReader())
    {
        while (r.Read())
        {
            Console.WriteLine("( {0}, {1} )", r.GetString(0), r.GetString(1));
        }
    }
}

public static string prepare_synchronization(MySqlConnection con)
{
    string gtid_executed = "";
    MySqlCommand cmd = new MySqlCommand("select @@global.gtid_executed", con);
    gtid_executed = ( string )cmd.ExecuteScalar();
    return gtid_executed;
}

public static void synchronize(MySqlConnection con, string gtid_executed)
{
    MySqlCommand cmd = new MySqlCommand( string.Format( "SELECT WAIT_UNTIL_SQL_THREAD_AFTER_GTIDS('{0}', 0)"
        gtid_executed ), con);
    cmd.ExecuteNonQuery();
}

```

```
}  
}
```



# Chapter 11 MySQL Workbench and MySQL Fabric Integration

Browse, view status, and connect to any MySQL instance in a Fabric Cluster.



### Note


MySQL Workbench 6.3.9 no longer supports this feature.

### Feature Support:

- MySQL Workbench 6.3.0 - 6.3.8: Support for this feature requires Connector/Python and MySQL Fabric 1.5, including the Python module. Fabric 1.5 support was added in MySQL Workbench 6.3, and due to incompatible protocol changes, Fabric 1.4 support was dropped.
- MySQL Workbench 6.2.0 - 6.2.5: Support for this feature requires MySQL Fabric 1.4.

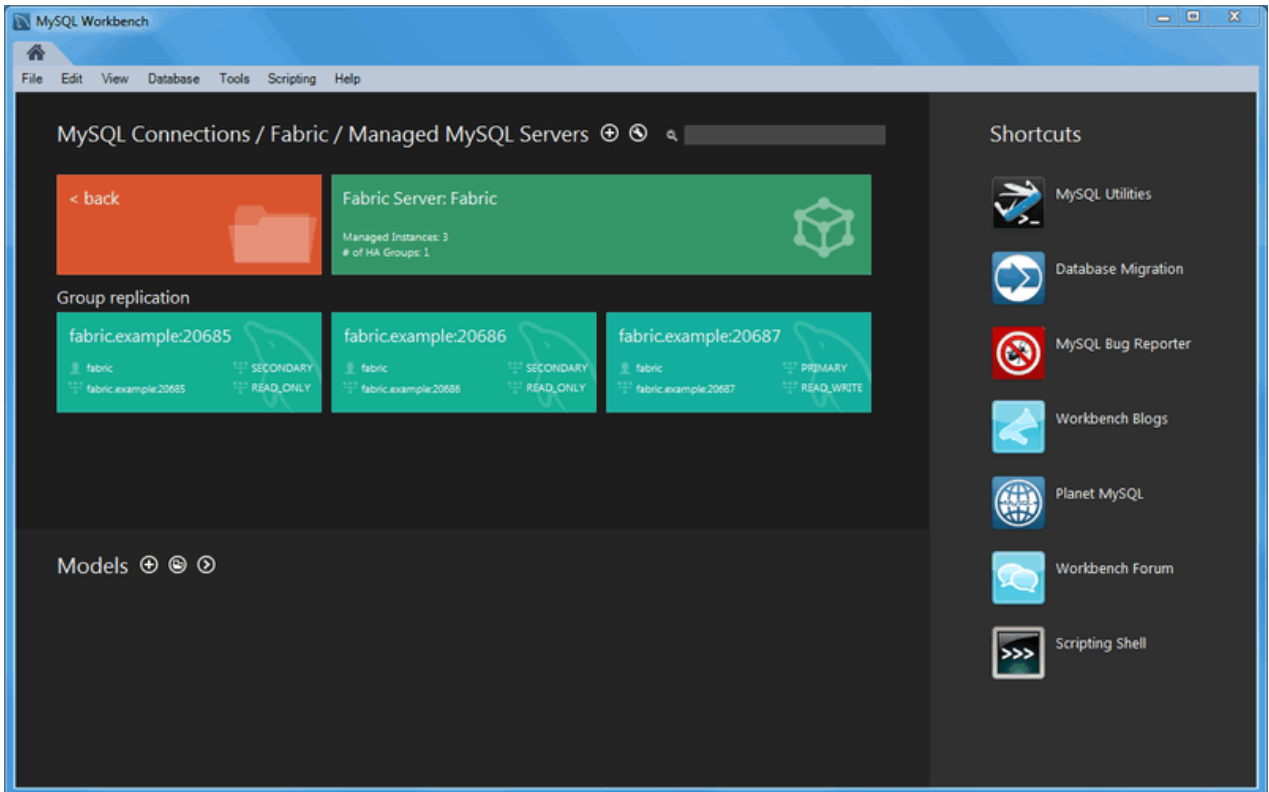
To set up a managed Fabric connection, create a new MySQL connection with the new **MySQL Fabric Management Node** connection method. The connection tiles have a different look:

**Figure 11.1 Fabric Connection Group Tile**



Clicking the new fabric group tile shows the managed connections:

Figure 11.2 Fabric Connection Group Tiles



---

# Chapter 12 MySQL Fabric Frequently Asked Questions

FAQ Categories

- [General Questions](#)
- [High-Availability Questions](#)
- [Sharding Questions](#)
- [Consistency Questions](#)

## General

12.1 What is MySQL Fabric? .....	105
12.2 Is it necessary to use a MySQL Fabric-specific Storage Engine? .....	105
12.3 What versions of MySQL are supported by MySQL Fabric? .....	105
12.4 What connectors support MySQL Fabric? .....	105
12.5 Are transactions ACID? .....	105
12.6 How many machines are needed in order to use MySQL Fabric? .....	105
12.7 Do I need to run an agent for each MySQL Server? .....	106
12.8 What interface is available to manage MySQL Fabric and its server farm? .....	106
12.9 How does MySQL Fabric Compare with MySQL Cluster? .....	106
12.10 How is MySQL Fabric licensed? .....	106
12.11 What if MySQL Fabric doesn't do what I need it to? .....	106

### 12.1. What is MySQL Fabric?

MySQL Fabric is a framework for managing groups of MySQL Servers and using those servers to provide services. It is designed to be extensible so that over time many different services can be added. In the current version the services provided are High Availability (built on top of MySQL Replication) and scale-out (by sharding the data).

MySQL Fabric is implemented as a MySQL Fabric node/process (which performs management functions) and Fabric-aware connectors that are able to route queries and transactions directly to the most appropriate MySQL Server. The MySQL Fabric node stores state and routing information in its State Store (which is a MySQL database).

### 12.2. Is it necessary to use a MySQL Fabric-specific Storage Engine?

**No.** The MySQL Servers that are being managed by MySQL Fabric continue to use InnoDB (and in the future NDB/MySQL Cluster may also be supported).

### 12.3. What versions of MySQL are supported by MySQL Fabric?

Currently MySQL 5.6. New MySQL releases will be fully supported as they reach General Availability status.

### 12.4. What connectors support MySQL Fabric?

Java, PHP, Python, and .NET. In addition the Hibernate and Doctrine Object-Relational Mappings frameworks are also supported. Connector/C 6.2 also adds fabric support as a labs release.

### 12.5. Are transactions ACID?

**Yes.** Because each transaction is local to a single MySQL Server, all of the ACID behavior of the InnoDB storage engine is experienced.

### 12.6. How many machines are needed in order to use MySQL Fabric?

---

For development, the MySQL Fabric node and all of the managed MySQL Servers can be hosted on a single machine.

For deployment, the minimal HA configuration would need 3 or more machines:

- 2 to host MySQL Servers
- 1 to host the MySQL Fabric process (that machine could also be running application code).

**12.7.** Do I need to run an agent for each MySQL Server?

**No.** The MySQL Fabric node is the only additional process and does not need to be co-located with any of the MySQL Servers that are being managed.

**12.8.** What interface is available to manage MySQL Fabric and its server farm?

A Command Line Interface (CLI) is provided as well as an XML/RPC API that can be used by connectors and/or applications to make management changes or retrieve the routing information - in this way, an application could use MySQL Fabric without a Fabric-aware connector.

**12.9.** How does MySQL Fabric Compare with MySQL Cluster?

[MySQL Cluster](#) is a mature, well proven solution for providing very high levels of availability and scaling out of both reads and writes. Some of the main extra capabilities that [MySQL Cluster](#) has over MySQL Fabric are:

- Synchronous replication
- Faster (automated) fail-over (resulting in higher availability)
- Transparent sharding
- Cross-shard joins and Foreign Keys
- In-memory, real-time performance

MySQL Fabric on the other hand, allows the application to stick with the InnoDB storage engine which is better suited to many applications.

**12.10.** How is MySQL Fabric licensed?

MySQL Fabric is available for use under the GPL v2 Open Source license or it can be commercially licensed as part of MySQL Enterprise Edition or MySQL Cluster Carrier Grade Edition.

**12.11.** What if MySQL Fabric doesn't do what I need it to?

There are a number of options:

- Raise feature requests or bug reports
- Modify the code to customize the current services. MySQL Fabric is written in Python and is designed to be easy to extend.
- Implement new modules that bind into the MySQL Fabric framework to implement new services.

**High-Availability**

12.1 How is High Availability achieved with MySQL Fabric? .....	107
12.2 How are MySQL Server failures detected? .....	107



---

12.3	What happens when the primary (master) MySQL Server fails? .....	107
12.4	Does my application need to do anything as part of the failover? .....	107
12.5	Is a recovered MySQL Server automatically put back into service? .....	107
12.6	Does MySQL Fabric work with semisynchronous replication? .....	107
12.7	Do I have to use MySQL Replication for implementing HA? .....	108
12.8	Is the MySQL Fabric node itself fault tolerant? What happens when the MySQL Fabric node is not available? .....	108

**12.1.** How is High Availability achieved with MySQL Fabric?

MySQL Fabric manages one or more HA-Groups where each HA-Group contains one or more MySQL Servers. For High Availability, a HA Group contains a Primary and one or more Secondary MySQL Servers. The Primary is currently a MySQL Replication master which replicates to each of the secondaries (MySQL Replication slaves) within the group.

By default, the Fabric-aware connectors route writes to the Primary and load balance reads across the available secondaries.

Should the Primary fail, MySQL Fabric will promote one of the Secondaries to be the new Primary (automatically promoting the MySQL Server to be the replication Master and updating the routing performed by the Fabric-aware connectors).

**12.2.** How are MySQL Server failures detected?

The MySQL Fabric node has a built-in monitoring function that checks on the status of the master. In addition, the Fabric-aware connectors report to MySQL Fabric when the Primary becomes unavailable to them. The administrator can configure how many problems need to be reported (and over what time period) before the failover is initiated.

**12.3.** What happens when the primary (master) MySQL Server fails?

The MySQL Fabric node will orchestrate the promotion of one of the Secondaries to be the new Primary. This involves two main actions:

- Promoting the Secondary to be the replication master (and any other surviving Secondaries will become slaves to the new master)
- Updating the routing information such that Fabric-aware connectors will no longer send any queries or transactions to the failed Primary and instead send all writes to the new Primary.

**12.4.** Does my application need to do anything as part of the failover?

**No.** The failover is transparent to the application as the Fabric-aware connectors will automatically start routing transactions and queries based on the new server topology. The application does need to handle the failure of a number of transactions when the Primary has failed but before the new Primary is in place but this should be considered part of normal MySQL error handling.

**12.5.** Is a recovered MySQL Server automatically put back into service?

No, the user must explicitly invoke MySQL Fabric to return a recovered MySQL Server to a HA Group. This is intentional so that the user can ensure that the server really is ready to take on an active role again.

**12.6.** Does MySQL Fabric work with semisynchronous replication?

In this version, MySQL Fabric sets up the HA Group to use asynchronous replication. If the user prefers to use semisynchronous replication then they can activate it manually after MySQL Fabric has created the replication relationships.

---

## 12.7. Do I have to use MySQL Replication for implementing HA?

At present, HA Groups are formed using MySQL Replication; future releases may offer further options such as MySQL Cluster or DRBD.

## 12.8. Is the MySQL Fabric node itself fault tolerant? What happens when the MySQL Fabric node is not available?

There is currently only a single instance of the MySQL Fabric node. If that process should fail then it can be restarted on that or another machine and the state and routing information read from the existing state store (a MySQL database) or a replicated copy of the state store.

While the MySQL Fabric node is unavailable, Fabric-aware connectors continue to route queries and transactions to the correct MySQL Servers based on their cached copies of the routing data. However, should a Primary fail, automated failover will not happen until the MySQL Fabric node is returned to service and so it's important to recover the process as quickly as possible.

## Sharding

12.1 How is scaling achieved with MySQL Fabric? .....	108
12.2 Does scaling apply to both reads and writes? .....	108
12.3 What if I have table data that needs to be in every shard? .....	108
12.4 How many MySQL Servers can I scale to? .....	109
12.5 Can MySQL Fabric introduce contention or deadlock? .....	109
12.6 What happens when my data set or usage grows and a shard grows too big? .....	109
12.7 Is there extra latency when using MySQL Fabric? .....	109
12.8 Why does MySQL Fabric route using connector logic rather than via a proxy? .....	109
12.9 What is the difference between a shard key and a shard identifier? .....	109
12.10 Does my application need to change when a shard is moved to a different MySQL Server or split into multiple shards? .....	109
12.11 Is it possible to perform cross-shard unions or joins? .....	110
12.12 Is the routing of queries and transactions transparent to my application? .....	110

### 12.1. How is scaling achieved with MySQL Fabric?

Horizontal scaling is achieved by partitioning (sharding) the data from a table across multiple MySQL Servers or HA Groups. In that way, each server or group will contain a subset of the rows from a specific table.

The user specifies what column from the table(s) should be used as the shard key as well as indicating whether to use a HASH or RANGE partitioning scheme for that key; if using RANGE based sharding then the user should also specify which ranges map to which shards. Currently the sharding key must be numeric.

When accessing the database, the application specifies the sharding key which the Fabric-aware connector will then map to a shard ID (using the mapping data it has retrieved and cached from MySQL Fabric) and route the query or transaction to the correct MySQL Server instance.

Within a HA group, the Fabric-aware connector is able to direct writes to the Primary and then spread the read queries across all available Secondaries (and optionally the Primary).

### 12.2. Does scaling apply to both reads and writes?

**Yes.** Both reads and writes scale linearly as more HA groups are added. Reads can also be scaled independently by adding more Secondary servers to a HA Group.

### 12.3. What if I have table data that needs to be in every shard?

---

A special group can be created called the *Global Group* which holds the Global Tables. Any table whose data should appear in its entirety in all HA Groups should be handled as a Global Table. For a Global Table, all writes are sent to the Global Group and then replicated to all of the HA Groups. An example might be the department table from an employee database - the contents of the department table being small enough to be stored in every server and where that data could be referenced by any record from one of the sharded employee tables.

Similarly, any schema changes would be sent to the Global Group where they can be replicated to all of the HA Groups.

**12.4.** How many MySQL Servers can I scale to?

There is no limit—either to the number of HA Groups or the number of servers within a HA group.

**12.5.** Can MySQL Fabric introduce contention or deadlock?

**No.** A single transaction can only access data from a single shard (+ Global Table data) and all writes are sent to the Primary server within that shard's HA Group. In other words, all writes to a specific row will be handled by the same MySQL Server and so InnoDB's row-based locking will work as normal.

**12.6.** What happens when my data set or usage grows and a shard grows too big?

MySQL Fabric provides the ability to either:

- Move a shard to a new HA group containing larger or more powerful servers
- Split an existing shard into two shards where the new shard will be stored in a new HA Group. In the future, different levels of granularity may be supported for shard splitting.

**12.7.** Is there extra latency when using MySQL Fabric?

**No.** Because the routing is handled within the connector there is no need for any extra "hops" to route the request via a proxy process.

**12.8.** Why does MySQL Fabric route using connector logic rather than via a proxy?

One reason is to reduce complexity; rather than having a pool of proxy processes (a single proxy would represent a single point of failure) the logic is just part of each connector instance. The second reason is to avoid the latency involved in all operations being diverted via a proxy process (which is likely to be on a different machine).

**12.9.** What is the difference between a shard key and a shard identifier?

The shard key is simply the value of a column from one or more tables. The shard key does not change if a row is migrated from one shard or server to another. The shard key is mapped to a shard id (using either a HASH or RANGE based mapping scheme); the shard id represents the shard itself.

As an example, if an existing shard were split in two then some of the rows would map to one shard's shard id and the rest to the other's; any given row's shard key would *\*not\** change as part of that split.

Very importantly, shard keys are known to the application while shard ids are not and so any changes to the topology of the collection of servers is completely transparent to the application.

**12.10.** Does my application need to change when a shard is moved to a different MySQL Server or split into multiple shards?

---

**No.** Because the application deals in shard keys and shard keys do not change during shard moves or splits.

**12.11.** Is it possible to perform cross-shard unions or joins?

Not at present; all queries are limited to the data within a single shard + the Global Table data. If data from multiple shards is required then it is currently the application's responsibility to collect and aggregate the data.

**12.12.** Is the routing of queries and transactions transparent to my application?

Partially.

For HA, the application simply needs to specify whether the operations are read-only or involve writes (or consistent reads).

For sharding, the application must specify the sharding key (a column from one or more tables) \*but\* this is independent of the topology of the MySQL Servers and where the data is held and it is unaffected when data is moved from one server to another.

**Consistency**

12.1 What do I do if I need immediately-consistent reads? ..... 110

**12.1.** What do I do if I need immediately-consistent reads?

Because replication from the Primary (master) to the Secondaries (slaves) is not synchronous, you cannot guarantee that you will retrieve the most current data when reading from a secondary. To force a read to be sent to the Primary, the application may set the \*mode\* property for the connection to read/write rather than read.