

MySQL Connector/J 8.0 Developer Guide

Abstract

This manual describes how to install, configure, and develop database applications using MySQL Connector/J 8.0, a JDBC and X DevAPI driver for communicating with MySQL servers.

MySQL Connector/J 8.0 is highly recommended for use with MySQL Server 8.0, 5.7, 5.6, and 5.5. Please upgrade to MySQL Connector/J 8.0.

For notes detailing the changes in each release of Connector/J 8.0, see [MySQL Connector/J 8.0 Release Notes](#).

For legal information, including licensing information, see the [Preface and 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.

Document generated on: 2018-07-18 (revision: 58142)

Table of Contents

Preface and Legal Notices	v
1 Overview of MySQL Connector/J	1
2 Connector/J Versions, and the MySQL and Java Versions They Support	3
3 What's New in Connector/J 8.0?	5
4 Connector/J Installation	7
4.1 Installing Connector/J from a Binary Distribution	7
4.2 Installing the Driver and Configuring the <code>CLASSPATH</code>	7
4.3 Upgrading from an Older Version	9
4.3.1 Upgrading to MySQL Connector/J 8.0	9
4.4 Installing from the Development Source Tree	13
4.5 Testing Connector/J	14
5 Connector/J Examples	17
6 Connector/J Reference	19
6.1 Driver/Datasource Class Name	19
6.2 Connection URL Syntax	19
6.3 Configuration Properties	22
6.4 JDBC API Implementation Notes	53
6.5 Java, JDBC, and MySQL Types	55
6.6 Using Character Sets and Unicode	58
6.7 Connecting Securely Using SSL	59
6.8 Connecting Using Unix Domain Sockets	63
6.9 Connecting Using Named Pipes	63
6.10 Connecting Using PAM Authentication	64
6.11 Using Master/Slave Replication with <code>ReplicationConnection</code>	64
6.12 Mapping MySQL Error Numbers to JDBC <code>SQLState</code> Codes	64
7 JDBC Concepts	73
7.1 Connecting to MySQL Using the JDBC <code>DriverManager</code> Interface	73
7.2 Using JDBC <code>Statement</code> Objects to Execute SQL	74
7.3 Using JDBC <code>CallableStatements</code> to Execute Stored Procedures	75
7.4 Retrieving <code>AUTO_INCREMENT</code> Column Values through JDBC	78
8 Connection Pooling with Connector/J	83
9 Multi-Host Connections	87
9.1 Configuring Server Failover	87
9.2 Configuring Client-Side Failover when using the X Protocol	90
9.3 Configuring Load Balancing with Connector/J	90
9.4 Configuring Master/Slave Replication with Connector/J	93
9.5 Advanced Load-balancing and Failover Configuration	96
10 Using the Connector/J Interceptor Classes	99
11 Using Connector/J with Tomcat	101
12 Using Connector/J with JBoss	103
13 Using Connector/J with Spring	105
13.1 Using <code>JdbcTemplate</code>	106
13.2 Transactional JDBC Access	108
13.3 Connection Pooling with Spring	109
14 Troubleshooting Connector/J Applications	111
15 Known Issues and Limitations	119
16 Connector/J Support	121
16.1 Connector/J Community Support	121
16.2 How to Report Connector/J Bugs or Problems	121
Index	123

Preface and Legal Notices

This manual describes how to install, configure, and develop database applications using MySQL Connector/J, the JDBC driver for communicating with MySQL servers.

Licensing information. This product may include third-party software, used under license. If you are using a *Commercial* release of MySQL Connector/J 8.0, see the [MySQL Connector/J 8.0 Commercial License Information User Manual](#) 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 Connector/J 8.0, see the [MySQL Connector/J 8.0 Community License Information User Manual](#) for licensing information, including licensing information relating to third-party software that may be included in this Community release.

Legal Notices

Copyright © 1998, 2018, 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.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

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 Overview of MySQL Connector/J

MySQL provides connectivity for client applications developed in the Java programming language with MySQL Connector/J. Connector/J implements the [Java Database Connectivity \(JDBC\) API](#), as well as a number of value-adding extensions of it. It also supports the new X DevAPI.

MySQL Connector/J is a JDBC Type 4 driver. Different versions are available that are compatible with the [JDBC 3.0](#) and [JDBC 4.2](#) specifications (see [Chapter 2, Connector/J Versions, and the MySQL and Java Versions They Support](#)). The Type 4 designation means that the driver is a pure Java implementation of the MySQL protocol and does not rely on the MySQL client libraries.

For large-scale programs that use common design patterns of data access, consider using one of the popular persistence frameworks such as [Hibernate](#), [Spring's JDBC templates](#) or [MyBatis SQL Maps](#) to reduce the amount of JDBC code for you to debug, tune, secure, and maintain.

Key Topics

- For installation instructions for Connector/J, see [Chapter 4, Connector/J Installation](#).
- For help with connection strings, connection options, and setting up your connection through JDBC, see [Chapter 6, Connector/J Reference](#).
- For information on connection pooling, see [Chapter 8, Connection Pooling with Connector/J](#).
- For information on multi-host connections, see [Chapter 9, Multi-Host Connections](#).

Chapter 2 Connector/J Versions, and the MySQL and Java Versions They Support

There are currently two MySQL Connector/J versions available:

- Connector/J 8.0 (formerly Connector/J 6.0; see [Changes in MySQL Connector/J 8.0.7](#) for an explanation of the version number change) is a Type 4 pure Java JDBC 4.2 driver for the Java 8 platform. It provides compatibility with all the functionality of MySQL 5.5, 5.6, 5.7, and 8.0. Connector/J 8.0 provides ease of development features, including auto-registration with the Driver Manager, standardized validity checks, categorized SQLExceptions, support for large update counts, support for local and offset date-time variants from the `java.time` package, support for JDBC-4.x XML processing, support for per connection client information, and support for the `NCHAR`, `NVARCHAR` and `NCLOB` data types.
- Connector/J 5.1 is also a Type 4 pure Java JDBC driver that conforms to the JDBC 3.0, 4.0, 4.1, and 4.2 specifications. It provides compatibility with all the functionality of MySQL 5.5, 5.6, 5.7, and 8.0. Connector/J 5.1 is covered by [its own manual](#).

The following table summarizes the Connector/J versions available, along with the details of JDBC driver type, versions of the JDBC API supported, versions of MySQL Server supported, JRE supported, JDK required for building, and the support status for each of the Connector/J versions:

Table 2.1 Summary of Connector/J Versions

Connector/J version	JDBC version	MySQL Server version	JRE Supported	JDK Required for Compilation	Status
8.0	4.2	5.5, 5.6, 5.7, 8.0	1.8.x	1.8.x	General Availability
5.1	3.0, 4.0, 4.1, 4.2	5.5, 5.6*, 5.7*, 8.0*	1.5.x, 1.6.x, 1.7.x, 1.8.x*	1.5.x and 1.8.x	General Availability

* JRE 1.8.x is required for Connector/J 5.1 to connect to MySQL 5.6, 5.7, and 8.0 with SSL/TLS when using some cipher suites.

Chapter 3 What's New in Connector/J 8.0?

Note

Connector/J 8.0 was formerly known as Connector/J 6.0; see [Changes in MySQL Connector/J 8.0.7](#) for an explanation of the version number change. This section describes differences between Connector/J 8.0 and Connector/J 5.1.

Here is a summary of major new features of Connector/J 8.0 (for details on the differences between the Connector/J 5.1 and 8.0 and for instructions on migrating, see [Section 4.3.1, “Upgrading to MySQL Connector/J 8.0”](#)):

- It supports MySQL 5.5, 5.6, 5.7, and 8.0.
- It supports the JDBC 4.2 specification.
- It is a MySQL driver for the Java 8 platform. For Java 7 or earlier, use Connector/J 5.1 instead.
- It supports the new X DevAPI, through which native support by MySQL 5.7 and 8.0 for JSON, NoSQL, document collection, and other features are provided to Java applications. See [Using MySQL as a Document Store](#) and the [X DevAPI User Guide](#) for details.

Chapter 4 Connector/J Installation

Table of Contents

4.1 Installing Connector/J from a Binary Distribution	7
4.2 Installing the Driver and Configuring the <code>CLASSPATH</code>	7
4.3 Upgrading from an Older Version	9
4.3.1 Upgrading to MySQL Connector/J 8.0	9
4.4 Installing from the Development Source Tree	13
4.5 Testing Connector/J	14

MySQL Connector/J is distributed as a `.zip` or `.tar.gz` archive, available for download from the [Connector/J Download page](#). The archive contains the sources and the JAR archive named `mysql-connector-java-version.jar`.

You can install the Connector/J package using either the binary or source distribution. The binary distribution provides the easiest method for installation; the source distribution lets you customize your installation further. With either solution, you manually add the Connector/J location to your Java `CLASSPATH`.

If you are upgrading from a previous version, read the upgrade information in [Section 4.3, “Upgrading from an Older Version”](#) before continuing.

Connector/J is also available as part of the Maven project. For more information and to download the Connector/J JAR files, see the [Maven repository](#).

Important

You may also need to install the following third-party libraries on your system for Connector/J 8.0 to work:

- Protocol Buffers (only required for using the X DevAPI)
- Javassist (only required for building Connector/J 8.0 from source)

4.1 Installing Connector/J from a Binary Distribution

For the easiest method of installation, use the binary distribution of the Connector/J package. Extract the JAR archive from the tar/gzip or zip archive to a suitable location, then optionally make the information about the JAR archive available by changing your `CLASSPATH` (see [Section 4.2, “Installing the Driver and Configuring the `CLASSPATH`”](#)).

Use the appropriate graphical or command-line utility to extract the distribution (for example, WinZip for the `.zip` archive, and `tar` for the `.tar.gz` archive). Because there are potentially long file names in the distribution, we use the GNU tar archive format. Use GNU tar (or an application that understands the GNU tar archive format) to unpack the `.tar.gz` variant of the distribution.

You can also use Maven dependencies as an alternative installation method. In that case, the Connector/J binaries are automatically downloaded from The [Maven Central Repository](#) by default and managed locally by your Maven tool

4.2 Installing the Driver and Configuring the `CLASSPATH`

Once you have extracted the distribution archive, you can install the driver by placing `mysql-connector-java-version.jar` in your classpath, either by adding the full path to it to your `CLASSPATH` environment variable, or by directly specifying it with the command line switch `-cp` when starting the JVM.

To use the driver with the JDBC `DriverManager`, use `com.mysql.cj.jdbc.Driver` as the class that implements `java.sql.Driver`.

You can set the `CLASSPATH` environment variable under Unix, Linux, or OS X either locally for a user within the user's `.profile`, `.login` or other login file, or you can also set it globally by editing the global `/etc/profile` file.

For example, add the Connector/J driver to your `CLASSPATH` using one of the following forms, depending on your command shell:

```
# Bourne-compatible shell (sh, ksh, bash, zsh):
shell> export CLASSPATH=/path/mysql-connector-java-ver.jar:$CLASSPATH

# C shell (csh, tcsh):
shell> setenv CLASSPATH /path/mysql-connector-java-ver.jar:$CLASSPATH
```

For Windows platforms, you set the environment variable through the System Control Panel.

To use the X DevAPI features in Connector/J, you also need the external library `protobuf-java-2.6.0.jar`, which you can download from any official Maven repository, for example <https://repo1.maven.org/maven2/com/google/protobuf/protobuf-java/2.6.0/>, and add it to the `CLASSPATH`.

If you prefer, you can use Maven dependencies manager to install and configure the Connector/J library in your project. Connector/J is published in The [Maven Central Repository](#) with "`GroupId: mysql`" and "`ArtifactId: mysql-connector-java`," and can be linked to your project by adding the following dependency in your `pom.xml` file:

```
<dependency>
  <groupId>mysql</groupId>
  <artifactId>mysql-connector-java</artifactId>
  <version>x.y.z</version>
</dependency>
```

Note that if you choose use Maven to manage your project dependencies, you do not need to explicitly refer to the library `protobuf-java` as it is resolved by dependency transitivity. However, if you do not want to use the X DevAPI features, you may also want to add a dependency exclusion to avoid linking the unneeded sub-library. For example:

```
<dependency>
  <groupId>mysql</groupId>
  <artifactId>mysql-connector-java</artifactId>
  <version>x.y.z</version>
  <exclusions>
    <exclusion>
      <groupId>com.google.protobuf</groupId>
      <artifactId>protobuf-java</artifactId>
    </exclusion>
  </exclusions>
</dependency>
```

To use MySQL Connector/J with an application server such as GlassFish, Tomcat, or JBoss, read your vendor's documentation for more information on how to configure third-party class libraries, as most application servers ignore the `CLASSPATH` environment variable. For configuration examples for some J2EE application servers, see [Chapter 8, Connection Pooling with Connector/J, Section 9.3](#),

“[Configuring Load Balancing with Connector/J](#)”, and [Section 9.5, “Advanced Load-balancing and Failover Configuration](#)”. However, the authoritative source for JDBC connection pool configuration information for your particular application server is the documentation for that application server.

If you are developing servlets or JSPs, and your application server is J2EE-compliant, you can put the driver's `.jar` file in the `WEB-INF/lib` subdirectory of your webapp, as this is a standard location for third party class libraries in J2EE web applications.

You can also use the `MySQLDataSource` or `MySQLConnectionPoolDataSource` classes in the `com.mysql.cj.jdbc` package, if your J2EE application server supports or requires them. The `javax.sql.XADataSource` interface is implemented using the `com.mysql.cj.jdbc.MySQLXADataSource` class, which supports XA distributed transactions.

The various `MySQLDataSource` classes support the following parameters (through standard set mutators):

- `user`
- `password`
- `serverName` (see the previous section about failover hosts)
- `databaseName`
- `port`

4.3 Upgrading from an Older Version

This section has information for users who are upgrading from one version of Connector/J to another, or to a new version of the MySQL server that supports a more recent level of JDBC. A newer version of Connector/J might include changes to support new features, improve existing functionality, or comply with new standards.

4.3.1 Upgrading to MySQL Connector/J 8.0

Upgrading an application developed for Connector/J 5.1 to use Connector/J 8.0 might require certain changes to your code or the environment in which it runs. Here are some changes for Connector/J going from 5.1 to 8.0, for which adjustments might be required:

4.3.1.1 Running on the Java 8 Platform

Connector/J 8.0 is created specifically to run on the Java 8 platform. While Java 8 is known to be strongly compatible with earlier Java versions, incompatibilities do exist, and code designed to work on Java 7 might need to be adjusted before being run on Java 8. Developers should refer to the [incompatibility information](#) provided by Oracle.

4.3.1.2 Changes in Connection Properties

A complete list of Connector/J 8.0 connection properties are available in [connector-j-reference-set-config](#). The following are connection properties that have been changed (removed, added, have their names changed, or have their default values changed) going from Connector/J 5.1 to 8.0.

Properties that have been removed (do not use them during connection):

- `useDynamicCharsetInfo`

- `useBlobToStoreUTF8OutsideBMP`, `utf8OutsideBmpExcludedColumnNamePattern`, and `utf8OutsideBmpIncludedColumnNamePattern`: MySQL 5.5 and later supports the `utf8mb4` character set, which is the character set that should be used by Connector/J applications for supporting characters beyond the Basic Multilingual Plane (BMP) of Unicode Version 3.
- `useJvmCharsetConverters`: JVM character set conversion is now used in all cases
- The following date and time properties:
 - `dynamicCalendars`
 - `noTzConversionForTimeType`
 - `noTzConversionForDateType`
 - `cacheDefaultTimezone`
 - `useFastIntParsing`
 - `useFastDateParsing`
 - `useJDBCCompliantTimezoneShift`
 - `useLegacyDatetimeCode`
 - `useSSPSCompatibleTimezoneShift`
 - `useTimezone`
 - `useGmtMillisForDatetimes`
- `dumpMetadataOnColumnNotFound`
- `relaxAutoCommit`
- `strictFloatingPoint`
- `runningCTS13`
- `retainStatementAfterResultSetClose`
- `nullNamePatternMatchesAll` (removed since release 8.0.9)

Properties that have been added:

- `mysqlx.useAsyncProtocol`

Property that has its name changed:

- `com.mysql.jdbc.faultInjection.serverCharsetIndex` changed to `com.mysql.cj.testsuite.faultInjection.serverCharsetIndex`
- `loadBalanceEnableJMX` to `ha.enableJMX`
- `replicationEnableJMX` to `ha.enableJMX`

Properties that have their default values changed:

- `nullCatalogMeansCurrent` is now `false` by default

4.3.1.3 Changes in the Connector/J API

This section describes the changes to the Connector/J API going from version 5.1 to 8.0. You might need to adjust your API calls accordingly:

- The name of the class that implements `java.sql.Driver` in MySQL Connector/J has changed from `com.mysql.jdbc.Driver` to `com.mysql.cj.jdbc.Driver`. The old class name has been deprecated.
- The names of these commonly-used interfaces have also been changed:
 - `ExceptionInterceptor`: from `com.mysql.jdbc.ExceptionInterceptor` to `com.mysql.cj.exceptions.ExceptionInterceptor`
 - `StatementInterceptor`: from `com.mysql.jdbc.StatementInterceptorV2` to `com.mysql.cj.interceptors.QueryInterceptor`
 - `ConnectionLifecycleInterceptor`: from `com.mysql.jdbc.ConnectionLifecycleInterceptor` to `com.mysql.cj.jdbc.interceptors.ConnectionLifecycleInterceptor`
 - `AuthenticationPlugin`: from `com.mysql.jdbc.AuthenticationPlugin` to `com.mysql.cj.protocol.AuthenticationPlugin`
 - `BalanceStrategy`: from `com.mysql.jdbc.BalanceStrategy` to `com.mysql.cj.jdbc.ha.BalanceStrategy`.

4.3.1.4 Changes for Build Properties

A number of Ant properties for [building Connector/J from source](#) have been renamed; see [Table 4.1, “Changes with the Build Properties from Connector/J 5.1 to 8.0”](#)

Table 4.1 Changes with the Build Properties from Connector/J 5.1 to 8.0

Old name	New name
<code>com.mysql.jdbc.extra.libs</code>	<code>com.mysql.cj.extra.libs</code>
<code>com.mysql.jdbc.jdk</code>	<code>com.mysql.cj.build.jdk</code>
<code>debug.enable</code>	<code>com.mysql.cj.build.addDebugInfo</code>
<code>com.mysql.jdbc.noCleanBetweenCompiles</code>	<code>com.mysql.cj.build.noCleanBetweenCompiles</code>
<code>com.mysql.jdbc.commercialBuild</code>	<code>com.mysql.cj.build.commercial</code>
<code>com.mysql.jdbc.filterLicense</code>	<code>com.mysql.cj.build.filterLicense</code>
<code>com.mysql.jdbc.noCryptoBuild</code>	<code>com.mysql.cj.build.noCrypto</code>
<code>com.mysql.jdbc.noSources</code>	<code>com.mysql.cj.build.noSources</code>
<code>com.mysql.jdbc.noMavenSources</code>	<code>com.mysql.cj.build.noMavenSources</code>
<code>major_version</code>	<code>com.mysql.cj.build.driver.version.major</code>
<code>minor_version</code>	<code>com.mysql.cj.build.driver.version.minor</code>
<code>subminor_version</code>	<code>com.mysql.cj.build.driver.version.subminor</code>
<code>version_status</code>	<code>com.mysql.cj.build.driver.version.status</code>
<code>extra_version</code>	<code>com.mysql.cj.build.driver.version.extra</code>
<code>snapshot.version</code>	<code>com.mysql.cj.build.driver.version.snapshot</code>
<code>version</code>	<code>com.mysql.cj.build.driver.version</code>

Old name	New name
full.version	com.mysql.cj.build.driver.version.full
prodDisplayName	com.mysql.cj.build.driver.displayName
prodName	com.mysql.cj.build.driver.name
fullProdName	com.mysql.cj.build.driver.fullName
buildDir	com.mysql.cj.build.dir
buildDriverDir	com.mysql.cj.build.dir.driver
mavenUploadDir	com.mysql.cj.build.dir.maven
distDir	com.mysql.cj.dist.dir
toPackage	com.mysql.cj.dist.dir.prepare
packageDest	com.mysql.cj.dist.dir.package
com.mysql.jdbc.docs.sourceDir	com.mysql.cj.dist.dir.prebuilt.docs

4.3.1.5 Change for Test Properties

A number of Ant properties for testing Connector/J have been renamed or removed; see [Table 4.2, “Changes with the Test Properties from Connector/J 5.1 to 8.0”](#)

Table 4.2 Changes with the Test Properties from Connector/J 5.1 to 8.0

Old name	New name
buildTestDir	com.mysql.cj.testsuite.build.dir
junit.results	com.mysql.cj.testsuite.junit.results
com.mysql.jdbc.testsuite.jvm	com.mysql.cj.testsuite.jvm
test	com.mysql.cj.testsuite.test.class
methods	com.mysql.cj.testsuite.test.methods
com.mysql.jdbc.testsuite.url	com.mysql.cj.testsuite.url
com.mysql.jdbc.testsuite.admin-url	com.mysql.cj.testsuite.url.admin
com.mysql.jdbc.testsuite.ClusterUrl	com.mysql.cj.testsuite.url.cluster
com.mysql.jdbc.testsuite.url.sha256default	com.mysql.cj.testsuite.url.openssl
com.mysql.jdbc.testsuite.cantGrant	com.mysql.cj.testsuite.cantGrant
com.mysql.jdbc.testsuite.no-multi-hosts-tests	com.mysql.cj.testsuite.disable.multihost.tests
com.mysql.jdbc.test.ds.host	com.mysql.cj.testsuite.ds.host
com.mysql.jdbc.test.ds.port	com.mysql.cj.testsuite.ds.port
com.mysql.jdbc.test.ds.db	com.mysql.cj.testsuite.ds.db
com.mysql.jdbc.test.ds.user	com.mysql.cj.testsuite.ds.user
com.mysql.jdbc.test.ds.password	com.mysql.cj.testsuite.ds.password
com.mysql.jdbc.test.tabletype	com.mysql.cj.testsuite.loadstoreperf.tabletype
com.mysql.jdbc.testsuite.loadstoreperf.useBigResults	com.mysql.cj.testsuite.loadstoreperf.useBigResults
com.mysql.jdbc.testsuite.MiniAdminTest.runShutdown	com.mysql.cj.testsuite.miniAdminTest.runShutdown
com.mysql.jdbc.testsuite.noDebugOutput	com.mysql.cj.testsuite.noDebugOutput

Old name	New name
<code>com.mysql.jdbc.testsuite.retainArtifacts</code>	<code>com.mysql.cj.testsuite.retainArtifacts</code>
<code>com.mysql.jdbc.testsuite.runLongTests</code>	<code>com.mysql.cj.testsuite.runLongTests</code>
<code>com.mysql.jdbc.test.ServerControllerBased</code>	<code>com.mysql.cj.testsuite.serverControllerBased</code>
<code>com.mysql.jdbc.ReplicationConnection.isSlave</code>	<code>com.mysql.cj.testsuite.replicationConnection.isSlave</code>
<code>com.mysql.jdbc.test.isLocalHostnameReplacement</code>	Removed
<code>com.mysql.jdbc.testsuite.driver</code>	Removed
<code>com.mysql.jdbc.testsuite.url.default</code>	Removed. No longer needed, as multi-JVM tests have been removed from the test suite.

4.3.1.6 Other Changes

Here are other changes with Connector/J 8.0:

- Removed `ReplicationDriver`. Instead of using a separate driver, you can now obtain a connection for a replication setup just by using the `jdbc:mysql:replication://` scheme.
- See [Chapter 4, Connector/J Installation](#) for [third-party libraries required for Connector/J 8.0 to work](#).

4.4 Installing from the Development Source Tree

Caution

Read this section only if you are interested in helping us test our new code. To just get MySQL Connector/J up and running on your system, use a standard binary release distribution.

To install MySQL Connector/J from the development source tree, make sure that you have the following software on your system:

- A Git client, to check out the sources from our GitHub repository (available from <http://git-scm.com/downloads>).
- Apache Ant version 1.8.2 or newer (available from <http://ant.apache.org/>).
- JDK 1.8.x.
- JUnit 4.1.2 (available from <https://github.com/junit-team/junit/wiki/Download-and-Install>).
- Javaassist 3.19 or newer (available from <http://jboss-javassist.github.io/javassist/>).
- Protocol Buffers Java API 2.6.0 or newer (available from, for example, the [Maven Central Repository](#)).

To check out and compile MySQL Connector/J, follow these steps:

1. Check out the code from the source code repository for MySQL Connector/J located on GitHub at <https://github.com/mysql/mysql-connector-j>. The latest release of the Connector/J 8.0 series is on the `release/8.0` branch; use the following command to check it out:

```
shell> git clone --branch release/8.0 https://github.com/mysql/mysql-connector-j.git
```

Under the current directory, the commands create a `mysql-connector-j` subdirectory, which contains the code you want.

2. Make sure that you have JDK 1.8.x installed.
3. Place the required `junit.jar`, `javaassist.jar`, and `protobuf-java-x.y.z.jar` files in a separate directory—for example, `/home/username/ant-extralibs`.
4. Change your current working directory to the `mysql-connector-j` directory created in step 1 above.
5. In the directory, create a file named `build.properties` to indicate to Ant the locations of the root directories for your JDK 1.8.x installation, as well as the location of the extra libraries. The file should contain the following property settings, with the “`path_to_*`” parts replaced by the appropriate filepaths:

```
com.mysql.cj.build.jdk=path_to_jdk_1.8
com.mysql.cj.extra.libs=path_to_folder_for_extra_libraries
```

Alternatively, you can set the values of those properties through the Ant `-D` options.

6. Issue the following command to compile the driver and create a `.jar` file for Connector/J:

```
shell> ant dist
```

This creates a `build` directory in the current directory, where all the build output goes. A directory is created under the `build` directory, whose name includes the version number of the release you are building. That directory contains the sources, the compiled `.class` files, and a `.jar` file for deployment. For more information and other possible targets, including those that create a fully packaged distribution, issue the following command:

```
shell> ant -projecthelp
```

7. Install the newly created `.jar` file for the JDBC driver as you would install a binary `.jar` file you download from MySQL by following the instructions given in [Section 4.2, “Installing the Driver and Configuring the CLASSPATH”](#).

Note that a package containing both the binary and source code for Connector/J 8.0 can also be downloaded from the [Connector/J Download page](#).

Note

Going from Connector/J 5.1 to 8.0, a number of Ant properties for building Connector/J have been renamed or removed; see [Section 4.3.1.4, “Changes for Build Properties”](#) for details.

4.5 Testing Connector/J

The Connector/J source code repository or packages that are shipped with source code include an extensive test suite, containing test cases that can be executed independently. The test cases are divided into the following categories:

- *Unit tests*: They are methods located in packages aligning with the classes that they test.
- *Functional tests*: Classes from the package `testsuite.simple`. Include test code for the main features of Connector/J.
- *Performance tests*: Classes from the package `testsuite.perf`. Include test code to make measurements for the performance of Connector/J.

- *Regression tests*: Classes from the package `testsuite.regression`. Includes code for testing bug and regression fixes.
- *X DevAPI and X Protocol tests*: Classes from the package `testsuite.x` for testing X DevAPI and X Protocol functionality.

The bundled Ant build file contains targets like `test`, which can facilitate the process of running the Connector/J tests; see the target descriptions in the build file for details. Besides the requirements for building Connector/J from the source code described in [Section 4.4, “Installing from the Development Source Tree”](#), a number of the tests also require the File System Service Provider 1.2 for the Java Naming and Directory Interface (JNDI), available at <http://www.oracle.com/technetwork/java/javasebusiness/downloads/java-archive-downloads-java-plat-419418.html>—place the jar files downloaded from there into the `lib` directory or in the directory pointed to by the property `com.mysql.cj.extra.libs`.

To run the test using Ant, in addition to the properties required for [Section 4.4, “Installing from the Development Source Tree”](#), you must set the following properties in the `build.properties` file or through the Ant `-D` options:

- `com.mysql.cj.testsuite.jvm`: the JVM to be used for the tests. If the property is not set, the JVM supplied with `com.mysql.cj.build.jdk` will be used.
- `com.mysql.cj.testsuite.url`: it specifies the JDBC URL for connection to a MySQL test server; see [Section 6.2, “Connection URL Syntax”](#).
- `com.mysql.cj.testsuite.url.openssl`: it specifies the JDBC URL for connection to a MySQL test server compiled with OpenSSL; see [Section 6.2, “Connection URL Syntax”](#).
- `com.mysql.cj.testsuite.mysqlx.url`: it specifies the X DevAPI URL for connection to a MySQL test server; see [Section 6.2, “Connection URL Syntax”](#).
- `com.mysql.cj.testsuite.mysqlx.url.openssl`: it specifies the X DevAPI URL for connection to a MySQL test server compiled with OpenSSL; see [Section 6.2, “Connection URL Syntax”](#).

After setting these parameters, run the tests with Ant in the following ways:

- Building the `test` target with `ant test` runs all test cases by default on a single server instance. If you want to run a particular test case, put the test’s fully qualified class names in the `com.mysql.cj.testsuite.test.class` variable; for example:

```
shell > ant -Dcom.mysql.cj.testsuite.test.class=testsuite.simple.StringUtilsTest test
```

You can also run individual tests in a test case by specifying the names of the corresponding methods in the `com.mysql.cj.testsuite.test.methods` variable, separating multiple methods by commas; for example:

```
shell > ant -Dcom.mysql.cj.testsuite.test.class=testsuite.simple.StringUtilsTest \
-Dcom.mysql.cj.testsuite.test.methods=testIndexOfIgnoreCase,testGetBytes test
```

While the test results are partially reported by the console, complete reports in HTML and XML formats are provided. View the HTML report by opening `buildtest/junit/report/index.html`. XML version of the reports are located in the folder `buildtest/junit`.

Note

Going from Connector/J 5.1 to 8.0, a number of Ant properties for testing Connector/J have been renamed or removed; see [Section 4.3.1.5, “Change for Test Properties”](#) for details.

Chapter 5 Connector/J Examples

Examples of using Connector/J are located throughout this document. This section provides a summary and links to these examples.

- [Example 7.1, "Connector/J: Obtaining a connection from the `DriverManager`"](#)
- [Example 7.2, "Connector/J: Using `java.sql.Statement` to execute a `SELECT` query"](#)
- [Example 7.3, "Connector/J: Calling Stored Procedures"](#)
- [Example 7.4, "Connector/J: Using `Connection.prepareCall\(\)`"](#)
- [Example 7.5, "Connector/J: Registering output parameters"](#)
- [Example 7.6, "Connector/J: Setting `CallableStatement` input parameters"](#)
- [Example 7.7, "Connector/J: Retrieving results and output parameter values"](#)
- [Example 7.8, "Connector/J: Retrieving `AUTO_INCREMENT` column values using `Statement.getGeneratedKeys\(\)`"](#)
- [Example 7.9, "Connector/J: Retrieving `AUTO_INCREMENT` column values using `SELECT LAST_INSERT_ID\(\)`"](#)
- [Example 7.10, "Connector/J: Retrieving `AUTO_INCREMENT` column values in `Updatable ResultSets`"](#)
- [Example 8.1, "Connector/J: Using a connection pool with a J2EE application server"](#)
- [Example 14.1, "Connector/J: Example of transaction with retry logic"](#)

Chapter 6 Connector/J Reference

Table of Contents

6.1 Driver/Datasource Class Name	19
6.2 Connection URL Syntax	19
6.3 Configuration Properties	22
6.4 JDBC API Implementation Notes	53
6.5 Java, JDBC, and MySQL Types	55
6.6 Using Character Sets and Unicode	58
6.7 Connecting Securely Using SSL	59
6.8 Connecting Using Unix Domain Sockets	63
6.9 Connecting Using Named Pipes	63
6.10 Connecting Using PAM Authentication	64
6.11 Using Master/Slave Replication with ReplicationConnection	64
6.12 Mapping MySQL Error Numbers to JDBC SQLState Codes	64

This section of the manual contains reference material for MySQL Connector/J.

6.1 Driver/Datasource Class Name

The name of the class that implements `java.sql.Driver` in MySQL Connector/J is `com.mysql.cj.jdbc.Driver`.

6.2 Connection URL Syntax

This section explains the syntax of the URLs for connecting to MySQL.

This is the generic format of the connection URL:

```
protocol//[hosts][/database][?properties]
```

The URL consists of the following parts:

Important

Any reserved characters for URLs (for example, `/`, `:`, `@`, `(`, `)`, `[`, `]`, `&`, `#`, `=`, `?`, and space) that appear in any part of the connection URL must be percent encoded.

protocol

There are four possible protocols for a connection:

- `jdbc:mysql:` is for ordinary and basic failover connections.
- `jdbc:mysql:loadbalance:` is for configuring load balancing. See [Section 9.3, “Configuring Load Balancing with Connector/J”](#) for details.
- `jdbc:mysql:replication:` is for configuring a replication setup. See [Section 9.4, “Configuring Master/Slave Replication with Connector/J”](#) for details.
- `mysqlx:` is for connections using the X Protocol.

hosts

Depending on the situation, the *hosts* part may consist simply of a host name, or it can be a complex structure consisting of various elements like multiple host names, port numbers, host-specific properties, and user credentials.

- Single host:
 - Single-host connections without adding host-specific properties:
 - The *hosts* part is written in the format of *host:port*. This is an example of a simple single-host connection URL:

```
jdbc:mysql://host1:33060/sakila
```

- *host* can be an IPv4 or an IPv6 host name string, and in the latter case it must be put inside square brackets, for example “[1000:2000::abcd].” When *host* is not specified, the default value of *localhost* is used.
- *port* is a standard port number, i.e., an integer between 1 and 65535. The default port number for an ordinary MySQL connection is 3306, and it is 33060 for a connection using the X Protocol. If *port* is not specified, the corresponding default is used.
- Single-host connections adding host-specific properties:
 - In this case, the host is defined as a succession of *key=value* pairs. Keys are used to identify the host, the port, as well as any host-specific properties. There are two alternate formats for specifying keys:

- The “address-equals” form:

```
address=(host=host_or_ip)(port=port)(key1=value1)(key2=value2)...(keyN=valueN)
```

Here is a sample URL using the “address-equals” form :

```
jdbc:mysql://address=(host=myhost)(port=1111)(key1=value1)/db
```

- The “key-value” form:

```
(host=host,port=port,key1=value1,key2=value2,...,keyN=valueN)
```

Here is a sample URL using the “key-value” form :

```
jdbc:mysql://(host=myhost,port=1111,key1=value1)/db
```

- The host and the port are identified by the keys *host* and *port*. The descriptions of the format and default values of *host* and *port* in [Single host without host-specific properties \[20\]](#) above also apply here.
- Other keys that can be added include *user*, *password*, *protocol*, and so on. They override the global values set in the [properties part of the URL](#). Limit the overrides to user, password, network timeouts, and statement and metadata cache sizes; the effects of other per-host overrides are not defined.
- Different protocols may require different keys. For example, the *mysqlx:* scheme uses two special keys, *address* and *priority*. *address* is a *host:port* pair and *priority* an integer. For example:

```
mysqlx://(address=host:1111,priority=1,key1=value1)/db
```

- `key` is case-sensitive. Two keys differing in case only are considered conflicting, and there are no guarantees on which one will be used.

- Multiple hosts

There are two formats for specifying multiple hosts:

- List hosts in a comma-separated list:

```
host1,host2,...,hostN
```

Each host can be specified in any of the three ways described in [Single host \[20\]](#) above. Here are some examples:

```
jdbc:mysql://myhost1:1111,myhost2:2222/db
jdbc:mysql://address=(host=myhost1)(port=1111)(key1=value1),address=(host=myhost2)(port=2222)(key2=value2)/db
jdbc:mysql://(host=myhost1,port=1111,key1=value1),(host=myhost2,port=2222,key2=value2)/db
jdbc:mysql://myhost1:1111,(host=myhost2,port=2222,key2=value2)/db
mysqlx://(address=host1:1111,priority=1,key1=value1),(address=host2:2222,priority=2,key2=value2)/db
```

- List hosts in a comma-separated list, and then encloses the list by square brackets:

```
[host1,host2,...,hostN]
```

This is called the host sublist form, which allows sharing of the [user credentials](#) by all hosts in the list as if they are a single host. Each host in the list can be specified in any of the three ways described in [Single host \[20\]](#) above. Here are some examples:

```
jdbc:mysql://sandy:secret@[myhost1:1111,myhost2:2222]/db
jdbc:mysql://sandy:secret@[address=(host=myhost1)(port=1111)(key1=value1),address=(host=myhost2)(port=2222)(key2=value2)]/db
jdbc:mysql://sandy:secret@[myhost1:1111,address=(host=myhost2)(port=2222)(key2=value2)]/db
```

While it is not possible to write host sublists recursively, a host list may contain host sublists as its member hosts.

- User credentials

User credentials can be set outside of the connection URL—for example, as arguments when getting a connection from the `java.sql.DriverManager` (see [Section 6.3, “Configuration Properties”](#) for details). When set with the connection URL, there are several ways to specify them:

- Prefix the a single host, a host sublist (see [Multiple hosts \[21\]](#)), or any host in a list of hosts with the user credentials with an `@`:

```
user:password@host_or_host_sublist
```

For example:

```
mysqlx://sandy:secret@[ (address=host1:1111,priority=1,key1=value1) ,(address=host2:2222,priority=2,key2=value2) ]/db
```

- Use the keys `user` and `password` to specify credentials for each host:

```
(user=sandy)(password=mypass)
```

For example:

```
jdbc:mysql://[(host=myhost1,port=1111,user=sandy,password=secret),(host=myhost2,port=2222,user=finn,password=secret)]
jdbc:mysql://address=(host=myhost1)(port=1111)(user=sandy)(password=secret),address=(host=myhost2)(port=2222)
```

In both forms, when multiple user credentials are specified, the one to the left takes precedence—that is, going from left to right in the connection string, the first one found that is applicable to a host is the one that is used.

Inside a host sublist, no host can have user credentials in the @ format, but individual host can have user credentials specified in the key format.

database

The default database or catalog to open. If the database is not specified, the connection is made with no default database. In this case, either call the `setCatalog()` method on the `Connection` instance, or specify table names using the database name (that is, `SELECT dbname.tablename.colname FROM dbname.tablename...`) in your SQL statements. Opening a connection without specifying the database to use is, in general, only useful when building tools that work with multiple databases, such as GUI database managers.

Note

Always use the `Connection.setCatalog()` method to specify the desired database in JDBC applications, rather than the `USE database` statement.

properties

A succession of global properties applying to all hosts, preceded by ? and written as `key=value` pairs separated by the symbol "&." Here are some examples:

```
jdbc:mysql://(host=myhost1,port=1111),(host=myhost2,port=2222)/db?key1=value1&key2=value2&key3=value3
```

The following are true for the key-value pairs:

- `key` and `value` are just strings. Proper type conversion and validation are performed internally in Connector/J.
- `key` is case-sensitive. Two keys differing in case only are considered conflicting, and it is uncertain which one will be used.
- Any host-specific values specified with key-value pairs as explained in [Single host with host-specific properties \[20\]](#) and [Multiple hosts \[21\]](#) above override the global values set here.

See [Section 6.3, "Configuration Properties"](#) for details about the configuration properties.

6.3 Configuration Properties

Configuration properties define how Connector/J will make a connection to a MySQL server. Unless otherwise noted, properties can be set for a `DataSource` object or for a `Connection` object.

Configuration properties can be set in one of the following ways:

- Using the `set*()` methods on MySQL implementations of `java.sql.DataSource` (which is the preferred method when using implementations of `java.sql.DataSource`):
 - `com.mysql.cj.jdbc.MySQLDataSource`
 - `com.mysql.cj.jdbc.MySQLConnectionPoolDataSource`

- As a key/value pair in the `java.util.Properties` instance passed to `DriverManager.getConnection()` or `Driver.connect()`
- As a JDBC URL parameter in the URL given to `java.sql.DriverManager.getConnection()`, `java.sql.Driver.connect()` or the MySQL implementations of the `javax.sql.DataSource.setURL()` method. If you specify a configuration property in the URL without providing a value for it, nothing will be set; for example, adding `useServerPrepStmts` alone to the URL does not make Connector/J use server-side prepared statements; you need to add `useServerPrepStmts=true`.

Note

If the mechanism you use to configure a JDBC URL is XML-based, use the XML character literal `&` to separate configuration parameters, as the ampersand is a reserved character for XML.

The properties are listed in the following tables.

Authentication.

Properties and Descriptions
<p>user</p> <p>The user to connect as</p> <p>Since version: all versions</p>
<p>password</p> <p>The password to use when connecting</p> <p>Since version: all versions</p>

Connection.

Properties and Descriptions
<p>connectionAttributes</p> <p>A comma-delimited list of user-defined key:value pairs (in addition to standard MySQL-defined key:value pairs) to be passed to MySQL Server for display as connection attributes in the <code>PERFORMANCE_SCHEMA.SESSION_CONNECT_ATTRS</code> table. Example usage: <code>connectionAttributes=key1:value1,key2:value2</code> This functionality is available for use with MySQL Server version 5.6 or later only. Earlier versions of MySQL Server do not support connection attributes, causing this configuration option to be ignored. Setting <code>connectionAttributes=none</code> will cause connection attribute processing to be bypassed, for situations where Connection creation/initialization speed is critical.</p> <p>Since version: 5.1.25</p>
<p>connectionLifecycleInterceptors</p> <p>A comma-delimited list of classes that implement "com.mysql.cj.jdbc.interceptors.ConnectionLifecycleInterceptor" that should notified of connection lifecycle events (creation, destruction, commit, rollback, setCatalog and setAutoCommit) and potentially alter the execution of these commands. ConnectionLifecycleInterceptors are "stackable", more than one interceptor may be specified via the configuration property as a comma-delimited list, with the interceptors executed in order from left to right.</p> <p>Since version: 5.1.4</p>

Properties and Descriptions
<p>useConfigs</p> <p>Load the comma-delimited list of configuration properties before parsing the URL or applying user-specified properties. These configurations are explained in the 'Configurations' of the documentation.</p> <p>Since version: 3.1.5</p>
<p>authenticationPlugins</p> <p>Comma-delimited list of classes that implement com.mysql.cj.protocol.AuthenticationPlugin and which will be used for authentication unless disabled by "disabledAuthenticationPlugins" property.</p> <p>Since version: 5.1.19</p>
<p>createDatabaseIfNotExist</p> <p>Creates the database given in the URL if it doesn't yet exist. Assumes the configured user has permissions to create databases.</p> <p>Default: false</p> <p>Since version: 3.1.9</p>
<p>defaultAuthenticationPlugin</p> <p>Name of a class implementing com.mysql.cj.protocol.AuthenticationPlugin which will be used as the default authentication plugin (see below). It is an error to use a class which is not listed in "authenticationPlugins" nor it is one of the built-in plugins. It is an error to set as default a plugin which was disabled with "disabledAuthenticationPlugins" property. It is an error to set this value to null or the empty string (i.e. there must be at least a valid default authentication plugin specified for the connection, meeting all constraints listed above).</p> <p>Default: com.mysql.cj.protocol.a.authentication.MysqlNativePasswordPlugin</p> <p>Since version: 5.1.19</p>
<p>detectCustomCollations</p> <p>Should the driver detect custom charsets/collations installed on server (true/false, defaults to 'false'). If this option set to 'true' driver gets actual charsets/collations from server each time connection establishes. This could slow down connection initialization significantly.</p> <p>Default: false</p> <p>Since version: 5.1.29</p>
<p>disabledAuthenticationPlugins</p> <p>Comma-delimited list of classes implementing com.mysql.cj.protocol.AuthenticationPlugin or mechanisms, i.e. "mysql_native_password". The authentication plugins or mechanisms listed will not be used for authentication which will fail if it requires one of them. It is an error to disable the default authentication plugin (either the one named by "defaultAuthenticationPlugin" property or the hard-coded one if "defaultAuthenticationPlugin" property is not set).</p> <p>Since version: 5.1.19</p>
<p>disconnectOnExpiredPasswords</p>

<p>Properties and Descriptions</p> <p>If "disconnectOnExpiredPasswords" is set to "false" and password is expired then server enters "sandbox" mode and sends ERR(08001, ER_MUST_CHANGE_PASSWORD) for all commands that are not needed to set a new password until a new password is set.</p> <p>Default: true</p> <p>Since version: 5.1.23</p>
<p>interactiveClient</p> <p>Set the CLIENT_INTERACTIVE flag, which tells MySQL to timeout connections based on INTERACTIVE_TIMEOUT instead of WAIT_TIMEOUT</p> <p>Default: false</p> <p>Since version: 3.1.0</p>
<p>passwordCharacterEncoding</p> <p>What character encoding is used for passwords? Leaving this set to the default value (null), uses the value set in "characterEncoding" if there is one, otherwise uses UTF-8 as default encoding. If the password contains non-ASCII characters, the password encoding must match what server encoding was set to when the password was created. For passwords in other character encodings, the encoding will have to be specified with this property (or with "characterEncoding"), as it's not possible for the driver to auto-detect this.</p> <p>Since version: 5.1.7</p>
<p>propertiesTransform</p> <p>An implementation of com.mysql.cj.conf.ConnectionPropertiesTransform that the driver will use to modify URL properties passed to the driver before attempting a connection</p> <p>Since version: 3.1.4</p>
<p>rollbackOnPooledClose</p> <p>Should the driver issue a rollback() when the logical connection in a pool is closed?</p> <p>Default: true</p> <p>Since version: 3.0.15</p>
<p>useAffectedRows</p> <p>Don't set the CLIENT_FOUND_ROWS flag when connecting to the server (not JDBC-compliant, will break most applications that rely on "found" rows vs. "affected rows" for DML statements), but does cause "correct" update counts from "INSERT ... ON DUPLICATE KEY UPDATE" statements to be returned by the server.</p> <p>Default: false</p> <p>Since version: 5.1.7</p>

Session.

<p>Properties and Descriptions</p> <p>characterEncoding</p>

Properties and Descriptions
<p>What character encoding should the driver use when dealing with strings? (defaults is to 'autodetect')</p> <p>Since version: 1.1g</p>
<p>characterSetResults</p> <p>Character set to tell the server to return results as.</p> <p>Since version: 3.0.13</p>
<p>connectionCollation</p> <p>If set, tells the server to use this collation via 'set collation_connection'</p> <p>Since version: 3.0.13</p>
<p>sessionVariables</p> <p>A comma or semicolon separated list of name=value pairs to be sent as SET [SESSION] ... to the server when the driver connects.</p> <p>Since version: 3.1.8</p>
<p>useOldUTF8Behavior</p> <p>Use the UTF-8 behavior the driver did when communicating with 4.0 and older servers</p> <p>Default: false</p> <p>Since version: 3.1.6</p>

Networking.

Properties and Descriptions
<p>socksProxyHost</p> <p>Name or IP address of SOCKS host to connect through.</p> <p>Since version: 5.1.34</p>
<p>socksProxyPort</p> <p>Port of SOCKS server.</p> <p>Default: 1080</p> <p>Since version: 5.1.34</p>
<p>socketFactory</p> <p>The name of the class that the driver should use for creating socket connections to the server. This class must implement the interface 'com.mysql.cj.protocol.SocketFactory' and have public no-args constructor.</p> <p>Default: com.mysql.cj.protocol.StandardSocketFactory</p> <p>Since version: 3.0.3</p>
<p>connectTimeout</p>

<p>Properties and Descriptions</p> <p>Timeout for socket connect (in milliseconds), with 0 being no timeout. Only works on JDK-1.4 or newer. Defaults to '0'.</p> <p>Default: 0</p> <p>Since version: 3.0.1</p>
<p>socketTimeout</p> <p>Timeout (in milliseconds) on network socket operations (0, the default means no timeout).</p> <p>Default: 0</p> <p>Since version: 3.0.1</p>
<p>localSocketAddress</p> <p>Hostname or IP address given to explicitly configure the interface that the driver will bind the client side of the TCP/IP connection to when connecting.</p> <p>Since version: 5.0.5</p>
<p>maxAllowedPacket</p> <p>Maximum allowed packet size to send to server. If not set, the value of system variable 'max_allowed_packet' will be used to initialize this upon connecting. This value will not take effect if set larger than the value of 'max_allowed_packet'. Also, due to an internal dependency with the property "blobSendChunkSize", this setting has a minimum value of "8203" if "useServerPrepStmts" is set to "true".</p> <p>Default: 65535</p> <p>Since version: 5.1.8</p>
<p>tcpKeepAlive</p> <p>If connecting using TCP/IP, should the driver set SO_KEEPALIVE?</p> <p>Default: true</p> <p>Since version: 5.0.7</p>
<p>tcpNoDelay</p> <p>If connecting using TCP/IP, should the driver set SO_TCP_NODELAY (disabling the Nagle Algorithm)?</p> <p>Default: true</p> <p>Since version: 5.0.7</p>
<p>tcpRcvBuf</p> <p>If connecting using TCP/IP, should the driver set SO_RCV_BUF to the given value? The default value of '0', means use the platform default value for this property)</p> <p>Default: 0</p> <p>Since version: 5.0.7</p>
<p>tcpSndBuf</p>

<p>Properties and Descriptions</p> <p>If connecting using TCP/IP, should the driver set SO_SND_BUF to the given value? The default value of '0', means use the platform default value for this property)</p> <p>Default: 0</p> <p>Since version: 5.0.7</p>
<p>tcpTrafficClass</p> <p>If connecting using TCP/IP, should the driver set traffic class or type-of-service fields ?See the documentation for java.net.Socket.setTrafficClass() for more information.</p> <p>Default: 0</p> <p>Since version: 5.0.7</p>
<p>useCompression</p> <p>Use zlib compression when communicating with the server (true/false)? Defaults to 'false'.</p> <p>Default: false</p> <p>Since version: 3.0.17</p>
<p>useUnbufferedInput</p> <p>Don't use BufferedInputStream for reading data from the server</p> <p>Default: true</p> <p>Since version: 3.0.11</p>

Security.

<p>Properties and Descriptions</p>
<p>allowMultiQueries</p> <p>Allow the use of ';' to delimit multiple queries during one statement (true/false), defaults to 'false', and does not affect the addBatch() and executeBatch() methods, which instead rely on rewriteBatchStatements.</p> <p>Default: false</p> <p>Since version: 3.1.1</p>
<p>useSSL</p> <p>Use SSL when communicating with the server (true/false), default is 'true' when connecting to MySQL 5.5.45+, 5.6.26+ or 5.7.6+, otherwise default is 'false'</p> <p>Default: false</p> <p>Since version: 3.0.2</p>
<p>requireSSL</p> <p>Require server support of SSL connection if useSSL=true? (defaults to 'false').</p> <p>Default: false</p>

Properties and Descriptions
Since version: 3.1.0
verifyServerCertificate
If "useSSL" is set to "true", should the driver verify the server's certificate? When using this feature, the keystore parameters should be specified by the "clientCertificateKeyStore*" properties, rather than system properties. Default is 'false' when connecting to MySQL 5.5.45+, 5.6.26+ or 5.7.6+ and "useSSL" was not explicitly set to "true". Otherwise default is 'true'
Default: true
Since version: 5.1.6
clientCertificateKeyStoreUrl
URL to the client certificate KeyStore (if not specified, use defaults)
Since version: 5.1.0
clientCertificateKeyStoreType
KeyStore type for client certificates (NULL or empty means use the default, which is "JKS". Standard keystore types supported by the JVM are "JKS" and "PKCS12", your environment may have more available depending on what security products are installed and available to the JVM.
Default: JKS
Since version: 5.1.0
clientCertificateKeyStorePassword
Password for the client certificates KeyStore
Since version: 5.1.0
trustCertificateKeyStoreUrl
URL to the trusted root certificate KeyStore (if not specified, use defaults)
Since version: 5.1.0
trustCertificateKeyStoreType
KeyStore type for trusted root certificates (NULL or empty means use the default, which is "JKS". Standard keystore types supported by the JVM are "JKS" and "PKCS12", your environment may have more available depending on what security products are installed and available to the JVM.
Default: JKS
Since version: 5.1.0
trustCertificateKeyStorePassword
Password for the trusted root certificates KeyStore
Since version: 5.1.0
enabledSSLCipherSuites

<p>Properties and Descriptions</p> <p>If "useSSL" is set to "true", overrides the cipher suites enabled for use on the underlying SSL sockets. This may be required when using external JSSE providers or to specify cipher suites compatible with both MySQL server and used JVM.</p> <p>Since version: 5.1.35</p>
<p>enabledTLSProtocols</p> <p>If "useSSL" is set to "true", overrides the TLS protocols enabled for use on the underlying SSL sockets. This may be used to restrict connections to specific TLS versions.</p> <p>Since version: 8.0.8</p>
<p>allowLoadLocalInfile</p> <p>Should the driver allow use of 'LOAD DATA LOCAL INFILE...' (defaults to 'true').</p> <p>Default: true</p> <p>Since version: 3.0.3</p>
<p>allowUrlInLocalInfile</p> <p>Should the driver allow URLs in 'LOAD DATA LOCAL INFILE' statements?</p> <p>Default: false</p> <p>Since version: 3.1.4</p>
<p>allowPublicKeyRetrieval</p> <p>Allows special handshake roundtrip to get server RSA public key directly from server.</p> <p>Default: false</p> <p>Since version: 5.1.31</p>
<p>paranoid</p> <p>Take measures to prevent exposure sensitive information in error messages and clear data structures holding sensitive data when possible? (defaults to 'false')</p> <p>Default: false</p> <p>Since version: 3.0.1</p>
<p>serverRSAPublicKeyFile</p> <p>File path to the server RSA public key file for sha256_password authentication. If not specified, the public key will be retrieved from the server.</p> <p>Since version: 5.1.31</p>

Statements.

<p>Properties and Descriptions</p> <p>continueBatchOnError</p>
--

<p>Properties and Descriptions</p> <p>Should the driver continue processing batch commands if one statement fails. The JDBC spec allows either way (defaults to 'true').</p> <p>Default: true</p> <p>Since version: 3.0.3</p>
<p>dontTrackOpenResources</p> <p>The JDBC specification requires the driver to automatically track and close resources, however if your application doesn't do a good job of explicitly calling close() on statements or result sets, this can cause memory leakage. Setting this property to true relaxes this constraint, and can be more memory efficient for some applications. Also the automatic closing of the Statement and current ResultSet in Statement.closeOnCompletion() and Statement.getMoreResults ([Statement.CLOSE_CURRENT_RESULT Statement.CLOSE_ALL_RESULTS]), respectively, ceases to happen. This property automatically sets holdResultsOpenOverStatementClose=true.</p> <p>Default: false</p> <p>Since version: 3.1.7</p>
<p>queryInterceptors</p> <p>A comma-delimited list of classes that implement "com.mysql.cj.interceptors.QueryInterceptor" that should be placed "in between" query execution to influence the results. QueryInterceptors are "chainable", the results returned by the "current" interceptor will be passed on to the next in in the chain, from left-to-right order, as specified in this property.</p> <p>Since version: 8.0.7</p>
<p>queryTimeoutKillsConnection</p> <p>If the timeout given in Statement.setQueryTimeout() expires, should the driver forcibly abort the Connection instead of attempting to abort the query?</p> <p>Default: false</p> <p>Since version: 5.1.9</p>

Prepared Statements.

<p>Properties and Descriptions</p>
<p>allowNanAndInf</p> <p>Should the driver allow NaN or +/- INF values in PreparedStatement.setDouble()?</p> <p>Default: false</p> <p>Since version: 3.1.5</p>
<p>autoClosePStmtStreams</p> <p>Should the driver automatically call .close() on streams/readers passed as arguments via set*() methods?</p> <p>Default: false</p> <p>Since version: 3.1.12</p>

Properties and Descriptions
<p>compensateOnDuplicateKeyUpdateCounts</p> <p>Should the driver compensate for the update counts of "ON DUPLICATE KEY" INSERT statements (2 = 1, 0 = 1) when using prepared statements?</p> <p>Default: false</p> <p>Since version: 5.1.7</p>
<p>emulateUnsupportedPstmts</p> <p>Should the driver detect prepared statements that are not supported by the server, and replace them with client-side emulated versions?</p> <p>Default: true</p> <p>Since version: 3.1.7</p>
<p>generateSimpleParameterMetadata</p> <p>Should the driver generate simplified parameter metadata for PreparedStatements when no metadata is available either because the server couldn't support preparing the statement, or server-side prepared statements are disabled?</p> <p>Default: false</p> <p>Since version: 5.0.5</p>
<p>processEscapeCodesForPrepStmts</p> <p>Should the driver process escape codes in queries that are prepared? Default escape processing behavior in non-prepared statements must be defined with the property 'enableEscapeProcessing'.</p> <p>Default: true</p> <p>Since version: 3.1.12</p>
<p>useServerPrepStmts</p> <p>Use server-side prepared statements if the server supports them?</p> <p>Default: false</p> <p>Since version: 3.1.0</p>
<p>useStreamLengthsInPrepStmts</p> <p>Honor stream length parameter in PreparedStatement/ResultSet.setXXXStream() method calls (true/false, defaults to 'true')?</p> <p>Default: true</p> <p>Since version: 3.0.2</p>

Result Sets.

Properties and Descriptions
<p>clobberStreamingResults</p>

<p>Properties and Descriptions</p> <p>This will cause a 'streaming' ResultSet to be automatically closed, and any outstanding data still streaming from the server to be discarded if another query is executed before all the data has been read from the server.</p> <p>Default: false</p> <p>Since version: 3.0.9</p>
<p>emptyStringsConvertToZero</p> <p>Should the driver allow conversions from empty string fields to numeric values of '0'?</p> <p>Default: true</p> <p>Since version: 3.1.8</p>
<p>holdResultsOpenOverStatementClose</p> <p>Should the driver close result sets on Statement.close() as required by the JDBC specification?</p> <p>Default: false</p> <p>Since version: 3.1.7</p>
<p>jdbcCompliantTruncation</p> <p>Should the driver throw java.sql.DataTruncation exceptions when data is truncated as is required by the JDBC specification when connected to a server that supports warnings (MySQL 4.1.0 and newer)? This property has no effect if the server sql-mode includes STRICT_TRANS_TABLES.</p> <p>Default: true</p> <p>Since version: 3.1.2</p>
<p>maxRows</p> <p>The maximum number of rows to return (0, the default means return all rows).</p> <p>Default: -1</p> <p>Since version: all versions</p>
<p>netTimeoutForStreamingResults</p> <p>What value should the driver automatically set the server setting 'net_write_timeout' to when the streaming result sets feature is in use? (value has unit of seconds, the value '0' means the driver will not try and adjust this value)</p> <p>Default: 600</p> <p>Since version: 5.1.0</p>
<p>padCharsWithSpace</p> <p>If a result set column has the CHAR type and the value does not fill the amount of characters specified in the DDL for the column, should the driver pad the remaining characters with space (for ANSI compliance)?</p> <p>Default: false</p>

Properties and Descriptions
<p>Since version: 5.0.6</p>
<p>populateInsertRowWithDefaultValues</p> <p>When using ResultSets that are CONCUR_UPDATABLE, should the driver pre-populate the "insert" row with default values from the DDL for the table used in the query so those values are immediately available for ResultSet accessors? This functionality requires a call to the database for metadata each time a result set of this type is created. If disabled (the default), the default values will be populated by the an internal call to refreshRow() which pulls back default values and/or values changed by triggers.</p> <p>Default: false</p> <p>Since version: 5.0.5</p>
<p>strictUpdates</p> <p>Should the driver do strict checking (all primary keys selected) of updatable result sets (true, false, defaults to 'true')?</p> <p>Default: true</p> <p>Since version: 3.0.4</p>
<p>tinyInt1isBit</p> <p>Should the driver treat the datatype TINYINT(1) as the BIT type (because the server silently converts BIT -> TINYINT(1) when creating tables)?</p> <p>Default: true</p> <p>Since version: 3.0.16</p>
<p>transformedBitsBoolean</p> <p>If the driver converts TINYINT(1) to a different type, should it use BOOLEAN instead of BIT for future compatibility with MySQL-5.0, as MySQL-5.0 has a BIT type?</p> <p>Default: false</p> <p>Since version: 3.1.9</p>

Metadata.

Properties and Descriptions
<p>getProceduresReturnsFunctions</p> <p>Pre-JDBC4 DatabaseMetaData API has only the getProcedures() and getProcedureColumns() methods, so they return metadata info for both stored procedures and functions. JDBC4 was extended with the getFunctions() and getFunctionColumns() methods and the expected behaviours of previous methods are not well defined. For JDBC4 and higher, default 'true' value of the option means that calls of DatabaseMetaData.getProcedures() and DatabaseMetaData.getProcedureColumns() return metadata for both procedures and functions as before, keeping backward compatibility. Setting this property to 'false' decouples Connector/J from its pre-JDBC4 behaviours for DatabaseMetaData.getProcedures() and DatabaseMetaData.getProcedureColumns(), forcing them to return metadata for procedures only.</p> <p>Default: true</p>

Properties and Descriptions
Since version: 5.1.26
noAccessToProcedureBodies
When determining procedure parameter types for CallableStatements, and the connected user can't access procedure bodies through "SHOW CREATE PROCEDURE" or select on mysql.proc should the driver instead create basic metadata (all parameters reported as INOUT VARCHARs) instead of throwing an exception?
Default: false
Since version: 5.0.3
nullCatalogMeansCurrent
When DatabaseMetadataMethods ask for a 'catalog' parameter, does the value null mean use the current catalog?
Default: false
Since version: 3.1.8
useHostsInPrivileges
Add '@hostname' to users in DatabaseMetaData.getColumn/TablePrivileges() (true/false), defaults to 'true'.
Default: true
Since version: 3.0.2
useInformationSchema
Should the driver use the INFORMATION_SCHEMA to derive information used by DatabaseMetaData? Default is 'true' when connecting to MySQL 8.0.3+, otherwise default is 'false'.
Default: false
Since version: 5.0.0

BLOB/CLOB processing.

Properties and Descriptions
autoDeserialize
Should the driver automatically detect and de-serialize objects stored in BLOB fields?
Default: false
Since version: 3.1.5
blobSendChunkSize
Chunk size to use when sending BLOB/CLOBs via ServerPreparedStatements. Note that this value cannot exceed the value of "maxAllowedPacket" and, if that is the case, then this value will be corrected automatically.
Default: 1048576

Properties and Descriptions
Since version: 3.1.9
blobsAreStrings
Should the driver always treat BLOBs as Strings - specifically to work around dubious metadata returned by the server for GROUP BY clauses?
Default: false
Since version: 5.0.8
clobCharacterEncoding
The character encoding to use for sending and retrieving TEXT, MEDIUMTEXT and LONGTEXT values instead of the configured connection characterEncoding
Since version: 5.0.0
emulateLocators
Should the driver emulate java.sql.Blobs with locators? With this feature enabled, the driver will delay loading the actual Blob data until the one of the retrieval methods (getInputStream(), getBytes(), and so forth) on the blob data stream has been accessed. For this to work, you must use a column alias with the value of the column to the actual name of the Blob. The feature also has the following restrictions: The SELECT that created the result set must reference only one table, the table must have a primary key; the SELECT must alias the original blob column name, specified as a string, to an alternate name; the SELECT must cover all columns that make up the primary key.
Default: false
Since version: 3.1.0
functionsNeverReturnBlobs
Should the driver always treat data from functions returning BLOBs as Strings - specifically to work around dubious metadata returned by the server for GROUP BY clauses?
Default: false
Since version: 5.0.8
locatorFetchBufferSize
If 'emulateLocators' is configured to 'true', what size buffer should be used when fetching BLOB data for getBinaryInputStream?
Default: 1048576
Since version: 3.2.1

Datetime types processing.

Properties and Descriptions
noDatetimeStringSync
Don't ensure that ResultSet.getDatetimeType().toString().equals(ResultSet.getString())
Default: false

Properties and Descriptions
Since version: 3.1.7
sendFractionalSeconds
Send fractional part from TIMESTAMP seconds. If set to false, the nanoseconds value of TIMESTAMP values will be truncated before sending any data to the server. This option applies only to prepared statements, callable statements or updatable result sets.
Default: true
Since version: 5.1.37
serverTimezone
Override detection/mapping of time zone. Used when time zone from server doesn't map to Java time zone
Since version: 3.0.2
treatUtilDateAsTimestamp
Should the driver treat java.util.Date as a TIMESTAMP for the purposes of PreparedStatement.setObject()?
Default: true
Since version: 5.0.5
yearIsDateType
Should the JDBC driver treat the MySQL type "YEAR" as a java.sql.Date, or as a SHORT?
Default: true
Since version: 3.1.9
zeroDateTimeBehavior
What should happen when the driver encounters DATETIME values that are composed entirely of zeros (used by MySQL to represent invalid dates)? Valid values are "EXCEPTION", "ROUND" and "CONVERT_TO_NULL".
Default: EXCEPTION
Since version: 3.1.4

High Availability and Clustering.

Properties and Descriptions
autoReconnect
Should the driver try to re-establish stale and/or dead connections? If enabled the driver will throw an exception for a queries issued on a stale or dead connection, which belong to the current transaction, but will attempt reconnect before the next query issued on the connection in a new transaction. The use of this feature is not recommended, because it has side effects related to session state and data consistency when applications don't handle SQLExceptions properly, and is only designed to be used when you are unable to configure your application to handle SQLExceptions resulting from dead and stale connections

<p>Properties and Descriptions</p> <p>properly. Alternatively, as a last option, investigate setting the MySQL server variable "wait_timeout" to a high value, rather than the default of 8 hours.</p> <p>Default: false</p> <p>Since version: 1.1</p>
<p>autoReconnectForPools</p> <p>Use a reconnection strategy appropriate for connection pools (defaults to 'false')</p> <p>Default: false</p> <p>Since version: 3.1.3</p>
<p>failOverReadOnly</p> <p>When failing over in autoReconnect mode, should the connection be set to 'read-only'?</p> <p>Default: true</p> <p>Since version: 3.0.12</p>
<p>maxReconnects</p> <p>Maximum number of reconnects to attempt if autoReconnect is true, default is '3'.</p> <p>Default: 3</p> <p>Since version: 1.1</p>
<p>reconnectAtTxEnd</p> <p>If autoReconnect is set to true, should the driver attempt reconnections at the end of every transaction?</p> <p>Default: false</p> <p>Since version: 3.0.10</p>
<p>retriesAllDown</p> <p>When using loadbalancing or failover, the number of times the driver should cycle through available hosts, attempting to connect. Between cycles, the driver will pause for 250ms if no servers are available.</p> <p>Default: 120</p> <p>Since version: 5.1.6</p>
<p>initialTimeout</p> <p>If autoReconnect is enabled, the initial time to wait between re-connect attempts (in seconds, defaults to '2').</p> <p>Default: 2</p> <p>Since version: 1.1</p>
<p>queriesBeforeRetryMaster</p>

<p>Properties and Descriptions</p> <p>Number of queries to issue before falling back to the primary host when failed over (when using multi-host failover). Whichever condition is met first, 'queriesBeforeRetryMaster' or 'secondsBeforeRetryMaster' will cause an attempt to be made to reconnect to the primary host. Setting both properties to 0 disables the automatic fall back to the primary host at transaction boundaries. Defaults to 50.</p> <p>Default: 50</p> <p>Since version: 3.0.2</p>
<p>secondsBeforeRetryMaster</p> <p>How long should the driver wait, when failed over, before attempting to reconnect to the primary host? Whichever condition is met first, 'queriesBeforeRetryMaster' or 'secondsBeforeRetryMaster' will cause an attempt to be made to reconnect to the master. Setting both properties to 0 disables the automatic fall back to the primary host at transaction boundaries. Time in seconds, defaults to 30</p> <p>Default: 30</p> <p>Since version: 3.0.2</p>
<p>allowMasterDownConnections</p> <p>By default, a replication-aware connection will fail to connect when configured master hosts are all unavailable at initial connection. Setting this property to 'true' allows to establish the initial connection, by failing over to the slave servers, in read-only state. It won't prevent subsequent failures when switching back to the master hosts i.e. by setting the replication connection to read/write state.</p> <p>Default: false</p> <p>Since version: 5.1.27</p>
<p>allowSlaveDownConnections</p> <p>By default, a replication-aware connection will fail to connect when configured slave hosts are all unavailable at initial connection. Setting this property to 'true' allows to establish the initial connection. It won't prevent failures when switching to slaves i.e. by setting the replication connection to read-only state. The property 'readFromMasterWhenNoSlaves' should be used for this purpose.</p> <p>Default: false</p> <p>Since version: 6.0.2</p>
<p>ha.enableJMX</p> <p>Enables JMX-based management of load-balanced connection groups, including live addition/removal of hosts from load-balancing pool. Enables JMX-based management of replication connection groups, including live slave promotion, addition of new slaves and removal of master or slave hosts from load-balanced master and slave connection pools.</p> <p>Default: false</p> <p>Since version: 5.1.27</p>
<p>loadBalanceHostRemovalGracePeriod</p> <p>Sets the grace period to wait for a host being removed from a load-balanced connection, to be released when it is currently the active host.</p>

<p>Properties and Descriptions</p> <p>Default: 15000</p> <p>Since version: 6.0.3</p>
<p>readFromMasterWhenNoSlaves</p> <p>Replication-aware connections distribute load by using the master hosts when in read/write state and by using the slave hosts when in read-only state. If, when setting the connection to read-only state, none of the slave hosts are available, an SQLException is thrown back. Setting this property to 'true' allows to fail over to the master hosts, while setting the connection state to read-only, when no slave hosts are available at switch instant.</p> <p>Default: false</p> <p>Since version: 6.0.2</p>
<p>selfDestructOnPingMaxOperations</p> <p>If set to a non-zero value, the driver will report close the connection and report failure when Connection.ping() or Connection.isValid(int) is called if the connection's count of commands sent to the server exceeds this value.</p> <p>Default: 0</p> <p>Since version: 5.1.6</p>
<p>selfDestructOnPingSecondsLifetime</p> <p>If set to a non-zero value, the driver will close the connection and report failure when Connection.ping() or Connection.isValid(int) is called if the connection's lifetime exceeds this value (in milliseconds).</p> <p>Default: 0</p> <p>Since version: 5.1.6</p>
<p>ha.loadBalanceStrategy</p> <p>If using a load-balanced connection to connect to SQL nodes in a MySQL Cluster/NDB configuration (by using the URL prefix "jdbc:mysql:loadbalance://"), which load balancing algorithm should the driver use: (1) "random" - the driver will pick a random host for each request. This tends to work better than round-robin, as the randomness will somewhat account for spreading loads where requests vary in response time, while round-robin can sometimes lead to overloaded nodes if there are variations in response times across the workload. (2) "bestResponseTime" - the driver will route the request to the host that had the best response time for the previous transaction. (3) "serverAffinity" - the driver initially attempts to enforce server affinity while still respecting and benefiting from the fault tolerance aspects of the load-balancing implementation. The server affinity ordered list is provided using the property 'serverAffinityOrder'. If none of the servers listed in the affinity list is responsive, the driver then refers to the "random" strategy to proceed with choosing the next server.</p> <p>Default: random</p> <p>Since version: 5.0.6</p>
<p>loadBalanceAutoCommitStatementRegex</p> <p>When load-balancing is enabled for auto-commit statements (via loadBalanceAutoCommitStatementThreshold), the statement counter will only increment when the SQL matches the regular expression. By default, every statement issued matches.</p>

Properties and Descriptions
Since version: 5.1.15
<p>loadBalanceAutoCommitStatementThreshold</p> <p>When auto-commit is enabled, the number of statements which should be executed before triggering load-balancing to rebalance. Default value of 0 causes load-balanced connections to only rebalance when exceptions are encountered, or auto-commit is disabled and transactions are explicitly committed or rolled back.</p> <p>Default: 0</p> <p>Since version: 5.1.15</p>
<p>loadBalanceBlacklistTimeout</p> <p>Time in milliseconds between checks of servers which are unavailable, by controlling how long a server lives in the global blacklist.</p> <p>Default: 0</p> <p>Since version: 5.1.0</p>
<p>loadBalanceConnectionGroup</p> <p>Logical group of load-balanced connections within a classloader, used to manage different groups independently. If not specified, live management of load-balanced connections is disabled.</p> <p>Since version: 5.1.13</p>
<p>loadBalanceExceptionChecker</p> <p>Fully-qualified class name of custom exception checker. The class must implement <code>com.mysql.cj.jdbc.ha.LoadBalanceExceptionChecker</code> interface, and is used to inspect <code>SQLExceptions</code> and determine whether they should trigger fail-over to another host in a load-balanced deployment.</p> <p>Default: <code>com.mysql.cj.jdbc.ha.StandardLoadBalanceExceptionChecker</code></p> <p>Since version: 5.1.13</p>
<p>loadBalancePingTimeout</p> <p>Time in milliseconds to wait for ping response from each of load-balanced physical connections when using load-balanced Connection.</p> <p>Default: 0</p> <p>Since version: 5.1.13</p>
<p>loadBalanceSQLExceptionSubclassFailover</p> <p>Comma-delimited list of classes/interfaces used by default load-balanced exception checker to determine whether a given <code>SQLException</code> should trigger failover. The comparison is done using <code>Class.isInstance(SQLException)</code> using the thrown <code>SQLException</code>.</p> <p>Since version: 5.1.13</p>
<p>loadBalanceSQLStateFailover</p>

<p>Properties and Descriptions</p> <p>Comma-delimited list of SQLState codes used by default load-balanced exception checker to determine whether a given SQLException should trigger failover. The SQLState of a given SQLException is evaluated to determine whether it begins with any value in the comma-delimited list.</p> <p>Since version: 5.1.13</p>
<p>loadBalanceValidateConnectionOnSwapServer</p> <p>Should the load-balanced Connection explicitly check whether the connection is live when swapping to a new physical connection at commit/rollback?</p> <p>Default: false</p> <p>Since version: 5.1.13</p>
<p>pinGlobalTxToPhysicalConnection</p> <p>When using XAConnections, should the driver ensure that operations on a given XID are always routed to the same physical connection? This allows the XAConnection to support "XA START ... JOIN" after "XA END" has been called</p> <p>Default: false</p> <p>Since version: 5.0.1</p>
<p>replicationConnectionGroup</p> <p>Logical group of replication connections within a classloader, used to manage different groups independently. If not specified, live management of replication connections is disabled.</p> <p>Since version: 8.0.7</p>
<p>resourceId</p> <p>A globally unique name that identifies the resource that this datasource or connection is connected to, used for XAResource.isSameRM() when the driver can't determine this value based on hostnames used in the URL</p> <p>Since version: 5.0.1</p>
<p>serverAffinityOrder</p> <p>A comma separated list containing the host/port pairs that are to be used in load-balancing "serverAffinity" strategy. Only the sub-set of the hosts enumerated in the main hosts section in this URL will be used and they must be identical in case and type, i.e., can't use an IP address in one place and the corresponding host name in the other.</p> <p>Since version: 8.0.8</p>

Performance Extensions.

<p>Properties and Descriptions</p> <p>callableStmtCacheSize</p> <p>If 'cacheCallableStmts' is enabled, how many callable statements should be cached?</p> <p>Default: 100</p>

Properties and Descriptions
Since version: 3.1.2
metadataCacheSize
The number of queries to cache ResultSetMetadata for if cacheResultSetMetaData is set to 'true' (default 50)
Default: 50
Since version: 3.1.1
useLocalSessionState
Should the driver refer to the internal values of autocommit and transaction isolation that are set by Connection.setAutoCommit() and Connection.setTransactionIsolation() and transaction state as maintained by the protocol, rather than querying the database or blindly sending commands to the database for commit() or rollback() method calls?
Default: false
Since version: 3.1.7
useLocalTransactionState
Should the driver use the in-transaction state provided by the MySQL protocol to determine if a commit() or rollback() should actually be sent to the database?
Default: false
Since version: 5.1.7
prepStmtCacheSize
If prepared statement caching is enabled, how many prepared statements should be cached?
Default: 25
Since version: 3.0.10
prepStmtCacheSqlLimit
If prepared statement caching is enabled, what's the largest SQL the driver will cache the parsing for?
Default: 256
Since version: 3.0.10
parseInfoCacheFactory
Name of a class implementing com.mysql.cj.CacheAdapterFactory, which will be used to create caches for the parsed representation of client-side prepared statements.
Default: com.mysql.cj.PerConnectionLRUFactory
Since version: 5.1.1
serverConfigCacheFactory
Name of a class implementing com.mysql.cj.CacheAdapterFactory<String, Map<String, String>>, which will be used to create caches for MySQL server configuration values

<p>Properties and Descriptions</p> <p>Default: com.mysql.cj.util.PerVmServerConfigCacheFactory</p> <p>Since version: 5.1.1</p>
<p>alwaysSendSetIsolation</p> <p>Should the driver always communicate with the database when Connection.setTransactionIsolation() is called? If set to false, the driver will only communicate with the database when the requested transaction isolation is different than the whichever is newer, the last value that was set via Connection.setTransactionIsolation(), or the value that was read from the server when the connection was established. Note that useLocalSessionState=true will force the same behavior as alwaysSendSetIsolation=false, regardless of how alwaysSendSetIsolation is set.</p> <p>Default: true</p> <p>Since version: 3.1.7</p>
<p>maintainTimeStats</p> <p>Should the driver maintain various internal timers to enable idle time calculations as well as more verbose error messages when the connection to the server fails? Setting this property to false removes at least two calls to System.currentTimeMillis() per query.</p> <p>Default: true</p> <p>Since version: 3.1.9</p>
<p>useCursorFetch</p> <p>Should the driver use cursor-based fetching to retrieve rows? If set to "true" and "defaultFetchSize" > 0 (or setFetchSize() > 0 is called on a statement) then the cursor-based result set will be used. Please note that "useServerPrepStmts" is automatically set to "true" in this case because cursor functionality is available only for server-side prepared statements.</p> <p>Default: false</p> <p>Since version: 5.0.0</p>
<p>cacheCallableStmts</p> <p>Should the driver cache the parsing stage of CallableStatements</p> <p>Default: false</p> <p>Since version: 3.1.2</p>
<p>cachePrepStmts</p> <p>Should the driver cache the parsing stage of PreparedStatements of client-side prepared statements, the "check" for suitability of server-side prepared and server-side prepared statements themselves?</p> <p>Default: false</p> <p>Since version: 3.0.10</p>
<p>cacheResultSetMetadata</p> <p>Should the driver cache ResultSetMetaData for Statements and PreparedStatements? (Req. JDK-1.4+, true/false, default 'false')</p>

<p>Properties and Descriptions</p> <p>Default: false</p> <p>Since version: 3.1.1</p>
<p>cacheServerConfiguration</p> <p>Should the driver cache the results of 'SHOW VARIABLES' and 'SHOW COLLATION' on a per-URL basis?</p> <p>Default: false</p> <p>Since version: 3.1.5</p>
<p>defaultFetchSize</p> <p>The driver will call setFetchSize(n) with this value on all newly-created Statements</p> <p>Default: 0</p> <p>Since version: 3.1.9</p>
<p>dontCheckOnDuplicateKeyUpdateInSQL</p> <p>Stops checking if every INSERT statement contains the "ON DUPLICATE KEY UPDATE" clause. As a side effect, obtaining the statement's generated keys information will return a list where normally it wouldn't. Also be aware that, in this case, the list of generated keys returned may not be accurate. The effect of this property is canceled if set simultaneously with 'rewriteBatchedStatements=true'.</p> <p>Default: false</p> <p>Since version: 5.1.32</p>
<p>elideSetAutoCommits</p> <p>If using MySQL-4.1 or newer, should the driver only issue 'set autocommit=n' queries when the server's state doesn't match the requested state by Connection.setAutoCommit(boolean)?</p> <p>Default: false</p> <p>Since version: 3.1.3</p>
<p>enableEscapeProcessing</p> <p>Sets the default escape processing behavior for Statement objects. The method Statement.setEscapeProcessing() can be used to specify the escape processing behavior for an individual Statement object. Default escape processing behavior in prepared statements must be defined with the property 'processEscapeCodesForPrepStmts'.</p> <p>Default: true</p> <p>Since version: 6.0.1</p>
<p>enableQueryTimeouts</p> <p>When enabled, query timeouts set via Statement.setQueryTimeout() use a shared java.util.Timer instance for scheduling. Even if the timeout doesn't expire before the query is processed, there will be memory used by the TimerTask for the given timeout which won't be reclaimed until the time the timeout would have expired if it hadn't been cancelled by the driver. High-load environments might want to consider disabling this functionality.</p>

<p>Properties and Descriptions</p> <p>Default: true</p> <p>Since version: 5.0.6</p>
<p>largeRowSizeThreshold</p> <p>What size result set row should the JDBC driver consider "large", and thus use a more memory-efficient way of representing the row internally?</p> <p>Default: 2048</p> <p>Since version: 5.1.1</p>
<p>readOnlyPropagatesToServer</p> <p>Should the driver issue appropriate statements to implicitly set the transaction access mode on server side when Connection.setReadOnly() is called? Setting this property to 'true' enables InnoDB read-only potential optimizations but also requires an extra roundtrip to set the right transaction state. Even if this property is set to 'false', the driver will do its best effort to prevent the execution of database-state-changing queries. Requires minimum of MySQL 5.6.</p> <p>Default: true</p> <p>Since version: 5.1.35</p>
<p>rewriteBatchedStatements</p> <p>Should the driver use multiqueries (irregardless of the setting of "allowMultiQueries") as well as rewriting of prepared statements for INSERT into multi-value inserts when executeBatch() is called? Notice that this has the potential for SQL injection if using plain java.sql.Statements and your code doesn't sanitize input correctly. Notice that for prepared statements, server-side prepared statements can not currently take advantage of this rewrite option, and that if you don't specify stream lengths when using PreparedStatement.set*Stream(), the driver won't be able to determine the optimum number of parameters per batch and you might receive an error from the driver that the resultant packet is too large. Statement.getGeneratedKeys() for these rewritten statements only works when the entire batch includes INSERT statements. Please be aware using rewriteBatchedStatements=true with INSERT .. ON DUPLICATE KEY UPDATE that for rewritten statement server returns only one value as sum of all affected (or found) rows in batch and it isn't possible to map it correctly to initial statements; in this case driver returns 0 as a result of each batch statement if total count was 0, and the Statement.SUCCESS_NO_INFO as a result of each batch statement if total count was > 0.</p> <p>Default: false</p> <p>Since version: 3.1.13</p>
<p>useReadAheadInput</p> <p>Use newer, optimized non-blocking, buffered input stream when reading from the server?</p> <p>Default: true</p> <p>Since version: 3.1.5</p>

Debugging/Profiling.

<p>Properties and Descriptions</p> <p>logger</p>
--

<p>Properties and Descriptions</p> <p>The name of a class that implements "com.mysql.cj.log.Log" that will be used to log messages to. (default is "com.mysql.cj.log.StandardLogger", which logs to STDERR)</p> <p>Default: com.mysql.cj.log.StandardLogger</p> <p>Since version: 3.1.1</p>
<p>gatherPerfMetrics</p> <p>Should the driver gather performance metrics, and report them via the configured logger every 'reportMetricsIntervalMillis' milliseconds?</p> <p>Default: false</p> <p>Since version: 3.1.2</p>
<p>profileSQL</p> <p>Trace queries and their execution/fetch times to the configured logger (true/false) defaults to 'false'</p> <p>Default: false</p> <p>Since version: 3.1.0</p>
<p>reportMetricsIntervalMillis</p> <p>If 'gatherPerfMetrics' is enabled, how often should they be logged (in ms)?</p> <p>Default: 30000</p> <p>Since version: 3.1.2</p>
<p>maxQuerySizeToLog</p> <p>Controls the maximum length/size of a query that will get logged when profiling or tracing</p> <p>Default: 2048</p> <p>Since version: 3.1.3</p>
<p>packetDebugBufferSize</p> <p>The maximum number of packets to retain when 'enablePacketDebug' is true</p> <p>Default: 20</p> <p>Since version: 3.1.3</p>
<p>slowQueryThresholdMillis</p> <p>If 'logSlowQueries' is enabled, how long should a query (in ms) before it is logged as 'slow'?</p> <p>Default: 2000</p> <p>Since version: 3.1.2</p>
<p>slowQueryThresholdNanos</p> <p>If 'useNanosForElapsedTime' is set to true, and this property is set to a non-zero value, the driver will use this threshold (in nanosecond units) to determine if a query was slow.</p>

<p>Properties and Descriptions</p> <p>Default: 0</p> <p>Since version: 5.0.7</p>
<p>useUsageAdvisor</p> <p>Should the driver issue 'usage' warnings advising proper and efficient usage of JDBC and MySQL Connector/J to the log (true/false, defaults to 'false')?</p> <p>Default: false</p> <p>Since version: 3.1.1</p>
<p>autoGenerateTestcaseScript</p> <p>Should the driver dump the SQL it is executing, including server-side prepared statements to STDERR?</p> <p>Default: false</p> <p>Since version: 3.1.9</p>
<p>autoSlowLog</p> <p>Instead of using slowQueryThreshold* to determine if a query is slow enough to be logged, maintain statistics that allow the driver to determine queries that are outside the 99th percentile?</p> <p>Default: true</p> <p>Since version: 5.1.4</p>
<p>clientInfoProvider</p> <p>The name of a class that implements the com.mysql.cj.jdbc.ClientInfoProvider interface in order to support JDBC-4.0's Connection.get/setClientInfo() methods</p> <p>Default: com.mysql.cj.jdbc.CommentClientInfoProvider</p> <p>Since version: 5.1.0</p>
<p>enablePacketDebug</p> <p>When enabled, a ring-buffer of 'packetDebugBufferSize' packets will be kept, and dumped when exceptions are thrown in key areas in the driver's code</p> <p>Default: false</p> <p>Since version: 3.1.3</p>
<p>explainSlowQueries</p> <p>If 'logSlowQueries' is enabled, should the driver automatically issue an 'EXPLAIN' on the server and send the results to the configured log at a WARN level?</p> <p>Default: false</p> <p>Since version: 3.1.2</p>
<p>logSlowQueries</p> <p>Should queries that take longer than 'slowQueryThresholdMillis' be logged?</p>

<p>Properties and Descriptions</p> <p>Default: false</p> <p>Since version: 3.1.2</p>
<p>logXaCommands</p> <p>Should the driver log XA commands sent by MysqlXaConnection to the server, at the DEBUG level of logging?</p> <p>Default: false</p> <p>Since version: 5.0.5</p>
<p>profilerEventHandler</p> <p>Name of a class that implements the interface com.mysql.cj.log.ProfilerEventHandler that will be used to handle profiling/tracing events.</p> <p>Default: com.mysql.cj.log.LoggingProfilerEventHandler</p> <p>Since version: 5.1.6</p>
<p>resultSetSizeThreshold</p> <p>If the usage advisor is enabled, how many rows should a result set contain before the driver warns that it is suspiciously large?</p> <p>Default: 100</p> <p>Since version: 5.0.5</p>
<p>traceProtocol</p> <p>Should trace-level network protocol be logged?</p> <p>Default: false</p> <p>Since version: 3.1.2</p>
<p>useNanosForElapsedTime</p> <p>For profiling/debugging functionality that measures elapsed time, should the driver try to use nanoseconds resolution if available (JDK >= 1.5)?</p> <p>Default: false</p> <p>Since version: 5.0.7</p>
<p>useOnlyServerErrorMessages</p> <p>Don't prepend 'standard' SQLState error messages to error messages returned by the server.</p> <p>Default: true</p> <p>Since version: 3.0.15</p>

Exceptions/Warnings.

Properties and Descriptions
<p>dumpQueriesOnException</p> <p>Should the driver dump the contents of the query sent to the server in the message for SQLExceptions?</p> <p>Default: false</p> <p>Since version: 3.1.3</p>
<p>exceptionInterceptors</p> <p>Comma-delimited list of classes that implement com.mysql.cj.exceptions.ExceptionInterceptor. These classes will be instantiated one per Connection instance, and all SQLExceptions thrown by the driver will be allowed to be intercepted by these interceptors, in a chained fashion, with the first class listed as the head of the chain.</p> <p>Since version: 5.1.8</p>
<p>ignoreNonTxTables</p> <p>Ignore non-transactional table warning for rollback? (defaults to 'false').</p> <p>Default: false</p> <p>Since version: 3.0.9</p>
<p>includeInnodbStatusInDeadlockExceptions</p> <p>Include the output of "SHOW ENGINE INNODB STATUS" in exception messages when deadlock exceptions are detected?</p> <p>Default: false</p> <p>Since version: 5.0.7</p>
<p>includeThreadDumpInDeadlockExceptions</p> <p>Include a current Java thread dump in exception messages when deadlock exceptions are detected?</p> <p>Default: false</p> <p>Since version: 5.1.15</p>
<p>includeThreadNamesAsStatementComment</p> <p>Include the name of the current thread as a comment visible in "SHOW PROCESSLIST", or in Innodb deadlock dumps, useful in correlation with "includeInnodbStatusInDeadlockExceptions=true" and "includeThreadDumpInDeadlockExceptions=true".</p> <p>Default: false</p> <p>Since version: 5.1.15</p>

Tunes for integration with other products.

Properties and Descriptions
<p>overrideSupportsIntegrityEnhancementFacility</p>

<p>Properties and Descriptions</p> <p>Should the driver return "true" for DatabaseMetaData.supportsIntegrityEnhancementFacility() even if the database doesn't support it to workaround applications that require this method to return "true" to signal support of foreign keys, even though the SQL specification states that this facility contains much more than just foreign key support (one such application being OpenOffice)?</p> <p>Default: false</p> <p>Since version: 3.1.12</p>
<p>ultraDevHack</p> <p>Create PreparedStatements for prepareCall() when required, because UltraDev is broken and issues a prepareCall() for _all_ statements? (true/false, defaults to 'false')</p> <p>Default: false</p> <p>Since version: 2.0.3</p>

JDBC compliance.

<p>Properties and Descriptions</p>
<p>useColumnNamesInFindColumn</p> <p>Prior to JDBC-4.0, the JDBC specification had a bug related to what could be given as a "column name" to ResultSet methods like findColumn(), or getters that took a String property. JDBC-4.0 clarified "column name" to mean the label, as given in an "AS" clause and returned by ResultSetMetaData.getColumnLabel(), and if no AS clause, the column name. Setting this property to "true" will give behavior that is congruent to JDBC-3.0 and earlier versions of the JDBC specification, but which because of the specification bug could give unexpected results. This property is preferred over "useOldAliasMetadataBehavior" unless you need the specific behavior that it provides with respect to ResultSetMetadata.</p> <p>Default: false</p> <p>Since version: 5.1.7</p>
<p>pedantic</p> <p>Follow the JDBC spec to the letter.</p> <p>Default: false</p> <p>Since version: 3.0.0</p>
<p>useOldAliasMetadataBehavior</p> <p>Should the driver use the legacy behavior for "AS" clauses on columns and tables, and only return aliases (if any) for ResultSetMetaData.getColumnname() or ResultSetMetaData.getTableName() rather than the original column/table name? In 5.0.x, the default value was true.</p> <p>Default: false</p> <p>Since version: 5.0.4</p>

X Protocol and X DevAPI.

Properties and Descriptions
<p>xdevapi.asyncResponseTimeout</p> <p>Timeout (in seconds) for getting server response via X Protocol.</p> <p>Default: 300</p> <p>Since version: 8.0.7</p>
<p>xdevapi.auth</p> <p>Authentication mechanism to use with the X Protocol. Allowed values are "SHA256_MEMORY", "MYSQL41", "PLAIN", and "EXTERNAL". Value is case insensitive. If the property is not set, the mechanism is chosen depending on the connection type: "PLAIN" is used for TLS connections and "SHA256_MEMORY" or "MYSQL41" is used for unencrypted connections.</p> <p>Default: PLAIN</p> <p>Since version: 8.0.8</p>
<p>xdevapi.ssl-mode</p> <p>By default the network connections in X DevAPI sessions are SSL encrypted. This property permits to turn secure connections off or to choose different levels of security. The following values are allowed: DISABLED - Establish unencrypted connections; REQUIRED - (default) Establish secure connections if the server supports them, fails otherwise; VERIFY_CA - Like REQUIRED but additionally verify the server TLS certificate against the configured Certificate Authority (CA) certificates; VERIFY_IDENTITY - Like VERIFY_CA, but additionally verify that the server certificate matches the host to which the connection is attempted.</p> <p>Default: REQUIRED</p> <p>Since version: 8.0.7</p>
<p>xdevapi.ssl-truststore</p> <p>URL to the trusted CA certificates key store (if not specified, use defaults)</p> <p>Since version: 6.0.6</p>
<p>xdevapi.ssl-truststore-password</p> <p>Password for the trusted CA certificates key store</p> <p>Since version: 6.0.6</p>
<p>xdevapi.ssl-truststore-type</p> <p>Type of the trusted CA certificates key store (NULL or empty means use the default, which is "JKS")</p> <p>Default: JKS</p> <p>Since version: 6.0.6</p>
<p>xdevapi.useAsyncProtocol</p> <p>Use asynchronous variant of X Protocol</p> <p>Default: false</p>

Properties and Descriptions

Since version: 6.0.0

6.4 JDBC API Implementation Notes

MySQL Connector/J, as a rigorous implementation of the [JDBC API](#), passes all of the tests in the publicly available version of Oracle's JDBC compliance test suite. The JDBC specification is flexible on how certain functionality should be implemented. This section gives details on an interface-by-interface level about implementation decisions that might affect how you code applications with MySQL Connector/J.

- **BLOB**

You can emulate BLOBs with locators by adding the property `emulateLocators=true` to your JDBC URL. Using this method, the driver will delay loading the actual BLOB data until you retrieve the other data and then use retrieval methods (`getInputStream()`, `getBytes()`, and so forth) on the BLOB data stream.

You must use a column alias with the value of the column to the actual name of the BLOB, for example:

```
SELECT id, 'data' as blob_data from blobtable
```

You must also follow these rules:

- The `SELECT` must reference only one table. The table must have a [primary key](#).
- The `SELECT` must alias the original BLOB column name, specified as a string, to an alternate name.
- The `SELECT` must cover all columns that make up the primary key.

The BLOB implementation does not allow in-place modification (they are copies, as reported by the `DatabaseMetaData.locatorsUpdateCopies()` method). Because of this, use the corresponding `PreparedStatement.setBlob()` or `ResultSet.updateBlob()` (in the case of updatable result sets) methods to save changes back to the database.

- **Connection**

The `isClosed()` method does not ping the server to determine if it is available. In accordance with the JDBC specification, it only returns true if `closed()` has been called on the connection. If you need to determine if the connection is still valid, issue a simple query, such as `SELECT 1`. The driver will throw an exception if the connection is no longer valid.

- **DatabaseMetaData**

[Foreign key](#) information (`getImportedKeys()/getExportedKeys()` and `getCrossReference()`) is only available from `InnoDB` tables. The driver uses `SHOW CREATE TABLE` to retrieve this information, so if any other storage engines add support for foreign keys, the driver would transparently support them as well.

- **PreparedStatement**

Two variants of prepared statements are implemented by Connector/J, the client-side and the server-side prepared statements. Client-side prepared statements are used by default because early MySQL versions did not support the prepared statement feature or had problems with its implementation. Server-side prepared statements and binary-encoded result sets are used when the server supports them. To enable usage of server-side prepared statements, set `useServerPrepStmts=true`.

Be careful when using a server-side prepared statement with **large** parameters that are set using `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, `setCharacterStream()`, `setNCharacterStream()`, `setBlob()`, `setClob()`, or `setNClob()`. To re-execute the statement with any large parameter changed to a nonlarge parameter, call `clearParameters()` and set all parameters again. The reason for this is as follows:

- During both server-side prepared statements and client-side emulation, large data is exchanged only when `PreparedStatement.execute()` is called.
- Once that has been done, the stream used to read the data on the client side is closed (as per the JDBC spec), and cannot be read from again.
- If a parameter changes from large to nonlarge, the driver must reset the server-side state of the prepared statement to allow the parameter that is being changed to take the place of the prior large value. This removes all of the large data that has already been sent to the server, thus requiring the data to be re-sent, using the `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, `setCharacterStream()`, `setNCharacterStream()`, `setBlob()`, `setClob()`, or `setNClob()` method.

Consequently, to change the type of a parameter to a nonlarge one, you must call `clearParameters()` and set all parameters of the prepared statement again before it can be re-executed.

- **ResultSet**

By default, ResultSets are completely retrieved and stored in memory. In most cases this is the most efficient way to operate and, due to the design of the MySQL network protocol, is easier to implement. If you are working with ResultSets that have a large number of rows or large values and cannot allocate heap space in your JVM for the memory required, you can tell the driver to stream the results back one row at a time.

To enable this functionality, create a `Statement` instance in the following manner:

```
stmt = conn.createStatement( java.sql.ResultSet.TYPE_FORWARD_ONLY,  
                             java.sql.ResultSet.CONCUR_READ_ONLY );  
stmt.setFetchSize( Integer.MIN_VALUE );
```

The combination of a forward-only, read-only result set, with a fetch size of `Integer.MIN_VALUE` serves as a signal to the driver to stream result sets row-by-row. After this, any result sets created with the statement will be retrieved row-by-row.

There are some caveats with this approach. You must read all of the rows in the result set (or close it) before you can issue any other queries on the connection, or an exception will be thrown.

The earliest the locks these statements hold can be released (whether they be `MyISAM` table-level locks or row-level locks in some other storage engine such as `InnoDB`) is when the statement completes.

If the statement is within scope of a transaction, then locks are released when the transaction completes (which implies that the statement needs to complete first). As with most other databases, statements are not complete until all the results pending on the statement are read or the active result set for the statement is closed.

Therefore, if using streaming results, process them as quickly as possible if you want to maintain concurrent access to the tables referenced by the statement producing the result set.

Another alternative is to use cursor-based streaming to retrieve a set number of rows each time. This can be done by setting the connection property `useCursorFetch` to true, and then calling `setFetchSize(int)` with `int` being the desired number of rows to be fetched each time:

```
conn = DriverManager.getConnection("jdbc:mysql://localhost/?useCursorFetch=true", "user", "s3cr3t");
stmt = conn.createStatement();
stmt.setFetchSize(100);
rs = stmt.executeQuery("SELECT * FROM your_table_here");
```

- **Statement**

Connector/J includes support for both `Statement.cancel()` and `Statement.setQueryTimeout()`. Both require a separate connection to issue the `KILL QUERY` statement. In the case of `setQueryTimeout()`, the implementation creates an additional thread to handle the timeout functionality.

Note

Failures to cancel the statement for `setQueryTimeout()` may manifest themselves as `RuntimeException` rather than failing silently, as there is currently no way to unblock the thread that is executing the query being cancelled due to timeout expiration and have it throw the exception instead.

MySQL does not support SQL cursors, and the JDBC driver does not emulate them, so `setCursorName()` has no effect.

Connector/J also supplies two additional methods:

- `setLocalInfileInputStream()` sets an `InputStream` instance that will be used to send data to the MySQL server for a `LOAD DATA LOCAL INFILE` statement rather than a `FileInputStream` or `URLInputStream` that represents the path given as an argument to the statement.

This stream will be read to completion upon execution of a `LOAD DATA LOCAL INFILE` statement, and will automatically be closed by the driver, so it needs to be reset before each call to `execute*()` that would cause the MySQL server to request data to fulfill the request for `LOAD DATA LOCAL INFILE`.

If this value is set to `NULL`, the driver will revert to using a `FileInputStream` or `URLInputStream` as required.

- `getLocalInfileInputStream()` returns the `InputStream` instance that will be used to send data in response to a `LOAD DATA LOCAL INFILE` statement.

This method returns `NULL` if no such stream has been set using `setLocalInfileInputStream()`.

6.5 Java, JDBC, and MySQL Types

MySQL Connector/J is flexible in the way it handles conversions between MySQL data types and Java data types.

In general, any MySQL data type can be converted to a `java.lang.String`, and any numeric type can be converted to any of the Java numeric types, although round-off, overflow, or loss of precision may occur.

Note

All `TEXT` types return `Types.LONGVARCHAR` with different `getPrecision()` values (65535, 255, 16777215, and 2147483647 respectively) with `getColumnType()` returning `-1`. This behavior is intentional even though `TINYTEXT` does not fall, regarding to its size, within the `LONGVARCHAR` category. This is to avoid different handling inside the same base type. And `getColumnType()` returns `-1` because the internal server handling is of type `TEXT`, which is similar to `BLOB`.

Also note that `getColumnTypeName()` will return `VARCHAR` even though `getColumnType()` returns `Types.LONGVARCHAR`, because `VARCHAR` is the designated column database-specific name for this type.

Connector/J issues warnings or throws `DataTruncation` exceptions as is required by the JDBC specification, unless the connection was configured not to do so by using the property `jdbcCompliantTruncation` and setting it to `false`.

The conversions that are always guaranteed to work are listed in the following table. The first column lists one or more MySQL data types, and the second column lists one or more Java types to which the MySQL types can be converted.

Table 6.1 Possible Conversions Between MySQL and Java Data Types

These MySQL Data Types	Can always be converted to these Java types
<code>CHAR</code> , <code>VARCHAR</code> , <code>BLOB</code> , <code>TEXT</code> , <code>ENUM</code> , and <code>SET</code>	<code>java.lang.String</code> , <code>java.io.InputStream</code> , <code>java.io.Reader</code> , <code>java.sql.Blob</code> , <code>java.sql.Clob</code>
<code>FLOAT</code> , <code>REAL</code> , <code>DOUBLE PRECISION</code> , <code>NUMERIC</code> , <code>DECIMAL</code> , <code>TINYINT</code> , <code>SMALLINT</code> , <code>MEDIUMINT</code> , <code>INTEGER</code> , <code>BIGINT</code>	<code>java.lang.String</code> , <code>java.lang.Short</code> , <code>java.lang.Integer</code> , <code>java.lang.Long</code> , <code>java.lang.Double</code> , <code>java.math.BigDecimal</code>
<code>DATE</code> , <code>TIME</code> , <code>DATETIME</code> , <code>TIMESTAMP</code>	<code>java.lang.String</code> , <code>java.sql.Date</code> , <code>java.sql.Timestamp</code>

Note

Round-off, overflow or loss of precision may occur if you choose a Java numeric data type that has less precision or capacity than the MySQL data type you are converting to/from.

The `ResultSet.getObject()` method uses the type conversions between MySQL and Java types, following the JDBC specification where appropriate. The values returned by `ResultSetMetaData.getColumnTypeName()` and `ResultSetMetaData.getColumnClassName()` are shown in the table below. For more information on the JDBC types, see the reference on the `java.sql.Types` class.

Table 6.2 MySQL Types and Return Values for `ResultSetMetaData.getColumnTypeName()` and `ResultSetMetaData.getColumnClassName()`

MySQL Type Name	Return value of <code>getColumnTypeName</code>	Return value of <code>getColumnClassName</code>
<code>BIT(1)</code>	<code>BIT</code>	<code>java.lang.Boolean</code>
<code>BIT(> 1)</code>	<code>BIT</code>	<code>byte[]</code>
<code>TINYINT</code>	<code>TINYINT</code>	<code>java.lang.Boolean</code> if the configuration property <code>tinyIntIsBit</code> is set to <code>true</code> (the default) and

MySQL Type Name	Return value of <code>getColumnTypeName</code>	Return value of <code>getColumnClassName</code>
		the storage size is 1, or <code>java.lang.Integer</code> if not.
BOOL, BOOLEAN	TINYINT	See TINYINT, above as these are aliases for TINYINT(1), currently.
SMALLINT[(M)] [UNSIGNED]	SMALLINT [UNSIGNED]	<code>java.lang.Integer</code> (regardless of whether it is UNSIGNED or not)
MEDIUMINT[(M)] [UNSIGNED]	MEDIUMINT [UNSIGNED]	<code>java.lang.Integer</code> (regardless of whether it is UNSIGNED or not)
INT, INTEGER[(M)] [UNSIGNED]	INTEGER [UNSIGNED]	<code>java.lang.Integer</code> , if UNSIGNED <code>java.lang.Long</code>
BIGINT[(M)] [UNSIGNED]	BIGINT [UNSIGNED]	<code>java.lang.Long</code> , if UNSIGNED <code>java.math.BigInteger</code>
FLOAT[(M,D)]	FLOAT	<code>java.lang.Float</code>
DOUBLE[(M,B)]	DOUBLE	<code>java.lang.Double</code>
DECIMAL[(M[,D])]	DECIMAL	<code>java.math.BigDecimal</code>
DATE	DATE	<code>java.sql.Date</code>
DATETIME	DATETIME	<code>java.sql.Timestamp</code>
TIMESTAMP[(M)]	TIMESTAMP	<code>java.sql.Timestamp</code>
TIME	TIME	<code>java.sql.Time</code>
YEAR[(2 4)]	YEAR	If <code>yearIsDateType</code> configuration property is set to <code>false</code> , then the returned object type is <code>java.sql.Short</code> . If set to <code>true</code> (the default), then the returned object is of type <code>java.sql.Date</code> with the date set to January 1st, at midnight.
CHAR(M)	CHAR	<code>java.lang.String</code> (unless the character set for the column is BINARY, then <code>byte[]</code> is returned.
VARCHAR(M) [BINARY]	VARCHAR	<code>java.lang.String</code> (unless the character set for the column is BINARY, then <code>byte[]</code> is returned.
BINARY(M)	BINARY	<code>byte[]</code>
VARBINARY(M)	VARBINARY	<code>byte[]</code>
TINYBLOB	TINYBLOB	<code>byte[]</code>
TINYTEXT	VARCHAR	<code>java.lang.String</code>
BLOB	BLOB	<code>byte[]</code>
TEXT	VARCHAR	<code>java.lang.String</code>
MEDIUMBLOB	MEDIUMBLOB	<code>byte[]</code>
MEDIUMTEXT	VARCHAR	<code>java.lang.String</code>
LOBLOB	LOBLOB	<code>byte[]</code>
LONGTEXT	VARCHAR	<code>java.lang.String</code>
ENUM('value1', 'value2', ...)	CHAR..)	<code>java.lang.String</code>
SET('value1', 'value2', ...)	CHAR..)	<code>java.lang.String</code>

6.6 Using Character Sets and Unicode

All strings sent from the JDBC driver to the server are converted automatically from native Java Unicode form to the client character encoding, including all queries sent using `Statement.execute()`, `Statement.executeUpdate()`, `Statement.executeQuery()` as well as all `PreparedStatement` and `CallableStatement` parameters with the exclusion of parameters set using `setBytes()`, `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, and `setBlob()`.

Number of Encodings Per Connection

Connector/J supports a single character encoding between client and server, and any number of character encodings for data returned by the server to the client in the results.

Setting the Character Encoding

The character encoding between client and server is automatically detected upon connection. You specify the encoding on the server using the `character_set_server`. The driver automatically uses the encoding specified by the server. For more information, see [Server Character Set and Collation](#).

For example, to use the 4-byte UTF-8 character set with Connector/J, configure the MySQL server with `character_set_server=utf8mb4`, and leave `characterEncoding` out of the Connector/J connection string. Connector/J will then autodetect the UTF-8 setting.

Warning

In order to use the `utf8mb4` character set for the connection, the server MUST be configured with `character_set_server=utf8mb4`; if that is not the case, when UTF-8 is used for `characterEncoding` in the connection string, it will map to the MySQL character set name `utf8`, which is an alias for `utf8mb3`.

To override the automatically detected encoding on the client side, use the `characterEncoding` property in the connection URL to the server. Use Java-style names when specifying character encodings. The following table lists MySQL character set names and their corresponding Java-style names:

Table 6.3 MySQL to Java Encoding Name Translations

MySQL Character Set Name	Java-Style Character Encoding Name
<code>ascii</code>	US-ASCII
<code>big5</code>	Big5
<code>gbk</code>	GBK
<code>sjis</code>	SJIS or Cp932
<code>cp932</code>	Cp932 or MS932
<code>gb2312</code>	EUC_CN
<code>ujis</code>	EUC_JP
<code>euackr</code>	EUC_KR
<code>latin1</code>	Cp1252
<code>latin2</code>	ISO8859_2
<code>greek</code>	ISO8859_7
<code>hebrew</code>	ISO8859_8

MySQL Character Set Name	Java-Style Character Encoding Name
<code>cp866</code>	<code>Cp866</code>
<code>tis620</code>	<code>TIS620</code>
<code>cp1250</code>	<code>Cp1250</code>
<code>cp1251</code>	<code>Cp1251</code>
<code>cp1257</code>	<code>Cp1257</code>
<code>macroman</code>	<code>MacRoman</code>
<code>macce</code>	<code>MacCentralEurope</code>
<code>utf8</code>	<code>UTF-8</code>
<code>ucs2</code>	<code>UnicodeBig</code>

Warning

Do not issue the query `SET NAMES` with Connector/J, as the driver will not detect that the character set has been changed by the query, and will continue to use the character set configured when the connection was first set up.

6.7 Connecting Securely Using SSL

Connector/J can encrypt all data communicated between the JDBC driver and the server (except for the initial handshake) using SSL. There is a performance penalty for enabling connection encryption, the severity of which depends on multiple factors including (but not limited to) the size of the query, the amount of data returned, the server hardware, the SSL library used, the network bandwidth, and so on.

The system works through two Java keystore files: one file contains the certificate information for the server (`truststore` in the examples below), and another contains the keys and certificate for the client (`keystore` in the examples below). All Java keystore files are protected by the password supplied to the `keytool` when you created the files. You need the file names and the associated passwords to create an SSL connection.

For SSL support to work, you must have the following:

- A MySQL server that supports SSL, and compiled and configured to do so. For more information, see [Using Encrypted Connections](#) and [Building MySQL with Support for Encrypted Connections](#).
- A signed client certificate, if using [mutual \(two-way\) authentication](#).

By default, Connector/J establishes secure connections with the MySQL servers. Note that MySQL servers 5.7 and 8.0, when compiled with OpenSSL, can automatically generate missing SSL files at startup and configure the SSL connection accordingly.

As long as the server is correctly configured to use SSL, there is no need to configure anything on the Connector/J client to use encrypted connections (the exception is when Connector/J is connecting to very old server versions like 5.5.44 and earlier, 5.6.25 and earlier, or 5.7.5 and earlier, in which case the client must set the connection property `useSSL=true` in order to use encrypted connections). The client can demand SSL to be used by setting the connection property `requireSSL=true`; the connection then fails if the server is not configured to use SSL. Without `requireSSL=true`, the connection just falls back to non-encrypted mode if the server is not configured to use SSL.

For additional security, you can setup the client for a one-way (server or client) or two-way (server and client) SSL authentication, allowing the client or the server to authenticate each other's identity.

Setting up Server Authentication

Server authentication via server certificate verification is enabled when the Connector/J connection property `verifyServerCertificate` is `true` (which is the default setting when `useSSL=true`).

Note

Standard Java SSL sockets do not support hostname verification, thus Connector/J does not support it. Host authentication is by certificates only.

Certificates signed by a trusted CA. When `verifyServerCertificate=true` (which is the default when `useSSL` is true), if no additional configurations are made regarding server authentication, Java verifies the server certificate using its default trusted CA certificates, usually from `$JAVA_HOME/lib/security/cacerts`.

Using self-signed certificates. It is pretty common though for MySQL server certificates to be self signed or signed by a self-signed CA certificate; the auto-generated certificates and keys created by the MySQL server are based on the latter—that is, the server generates all required keys and a self-signed CA certificate that is used to sign a server and a client certificate. The server then configures itself to use the CA certificate and the server certificate. Although the client certificate file is placed in the same directory, it is not used by the server.

To verify the server certificate, Connector/J needs to be able to read the certificate that signed it, that is, the server certificate that signed itself or the self-signed CA certificate. This can be accomplished by either importing the certificate (`ca.pem` or any other certificate) into the Java default truststore (although tampering the default truststore is not recommended) or by importing it into a custom Java truststore file and configuring the Connector/J driver accordingly. Use Java's `keytool` (typically located in the `bin` subdirectory of your JDK or JRE installation) to import the server certificates:

```
shell> keytool -importcert -alias MySQLCACert -file ca.pem \
  -keystore truststore -storepass mypassword
```

Supply the proper arguments for the command options. If the truststore file does not already exist, a new one will be created; otherwise the certificate will be added to the existing file. Interaction with `keytool` looks like this:

```
Owner: CN=MySQL_Server_5.7.17_Auto_Generated_CA_Certificate
Issuer: CN=MySQL_Server_5.7.17_Auto_Generated_CA_Certificate
Serial number: 1
Valid from: Thu Feb 16 11:42:43 EST 2017 until: Sun Feb 14 11:42:43 EST 2027
Certificate fingerprints:
  MD5:  18:87:97:37:EA:CB:0B:5A:24:AB:27:76:45:A4:78:C1
  SHA1: 2B:0D:D9:69:2C:99:BF:1E:2A:25:4E:8D:2D:38:B8:70:66:47:FA:ED
  SHA256: C3:29:67:1B:E5:37:06:F7:A9:93:DF:C7:B3:27:5E:09:C7:FD:EE:2D:18:86:F4:9C:40:D8:26:CB:DA:95:A0:24
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 1
Trust this certificate? [no]: yes
Certificate was added to keystore
```

The output of the command shows all details about the imported certificate. Make sure you remember the password you have supplied. Also, be mindful that the password will have to be written as plain text in your Connector/J configuration file or application source code.

The next step is to configure Java or Connector/J to read the truststore you just created or modified. This can be done by using one of the following three methods:

- Using the Java command line arguments:

```
-Djavax.net.ssl.trustStore=path_to_truststore_file
-Djavax.net.ssl.trustStorePassword=mypassword
```

- Setting the system properties directly in the client code:

```
System.setProperty("javax.net.ssl.trustStore", "path_to_truststore_file");
System.setProperty("javax.net.ssl.trustStorePassword", "mypassword");
```

- Setting the Connector/J connection properties:

```
clientCertificateKeyStoreUrl=file:path_to_truststore_file
clientCertificateKeyStorePassword=mypassword
```

Notice that when used together, the connection properties override the values set by the other two methods. Also, whatever values set with connection properties are used in that connection only, while values set using the system-wide values are used for all connections (unless overridden by the connection properties).

With the above setup and the connection property `verifyServerCertificate=true`, all connections established are going to be SSL-encrypted, with the server being authenticated in the SSL handshake process, and the client can now safely trust the server it is connecting to.

Setting up Client Authentication

The server may want to authenticate a client and require the client to provide an SSL certificate to it, which it verifies against its known certificate authorities or performs additional checks on the client identity if needed (see [CREATE USER SSL/TLS Options](#) for details). In that case, Connector/J needs to have access to the client certificate, so it can be sent to the server while establishing new database connections. This is done using the Java keystore files.

To allow client authentication, the client connecting to the server must have its own set of keys and an SSL certificate. The client certificate must be signed so that the server can verify it. While you can have the client certificates signed by official certificate authorities, it is more common to use an intermediate, private, CA certificate to sign client certificates. Such an intermediate CA certificate may be self-signed or signed by a trusted root CA. The requirement is that the server knows a CA certificate that is capable of validating the client certificate.

Some MySQL server builds are able to generate SSL keys and certificates for communication encryption, including a certificate and a private key (contained in the `client-cert.pem` and `client-key.pem` files), which can be used by any client. This SSL certificate is already signed by the self-signed CA certificate `ca.pem`, which the server may have already been configured to use.

If you do not want to use the client keys and certificate files generated by the server, you can also generate new ones using the procedures described in [Creating SSL and RSA Certificates and Keys](#). Notice that, according to the setup of the server, you may have to reuse the already existing CA certificate the server is configured to work with to sign the new client certificate, instead of creating a new one.

Once you have the client private key and certificate files you want to use, you need to import them into a Java keystore so that they can be used by the Java SSL library and Connector/J. The following instructions explain how to create the keystore file:

- Convert the client key and certificate files to a PKCS #12 archive:

```
shell> openssl pkcs12 -export -in client-cert.pem -inkey client-key.pem \
  -name "mysqlclient" -passout pass:mypassword -out client-keystore.p12
```

- Import the client key and certificate into a Java keystore:

```
shell> keytool -importkeystore -srckeystore client-keystore.p12 -srcstoretype pkcs12 \  
-srcstorepass mypassword -destkeystore keystore -deststoretype JKS -deststorepass mypassword
```

Supply the proper arguments for the command options. If the keystore file does not already exist, a new one will be created; otherwise the certificate will be added to the existing file. Output by `keytool` looks like this:

```
Entry for alias mysqlclient successfully imported.  
Import command completed: 1 entries successfully imported, 0 entries failed or cancelled
```

Make sure you remember the password you have chosen. Also, be mindful that the password will have to be written as plain text in your Connector/J configuration file or application source code.

After the step, you can delete the PKCS #12 archive (`client-keystore.p12` in the example).

The next step is to configure Java or Connector/J so that it reads the truststore you just created or modified. This can be done by using one of the following three methods:

- Using the Java command line arguments:

```
-Djavax.net.ssl.keyStore=path_to_keystore_file  
-Djavax.net.ssl.keyStorePassword=mypassword
```

- Setting the system properties directly in the client code:

```
System.setProperty("javax.net.ssl.keyStore", "path_to_keystore_file");  
System.setProperty("javax.net.ssl.keyStorePassword", "mypassword");
```

- Through Connector/J connection properties:

```
clientCertificateKeyStoreUrl=file:path_to_truststore_file  
clientCertificateKeyStorePassword=mypassword
```

Notice that when used together, the connection properties override the values set by the other two methods. Also, whatever values set with connection properties are used in that connection only, while values set using the system-wide values are used for all connections (unless overridden by the connection properties).

With the above setups, all connections established are going to be SSL-encrypted with the client being authenticated in the SSL handshake process, and the server can now safely trust the client that is requesting a connection to it.

Setting up 2-Way Authentication

Apply the steps outlined in both [Setting up Server Authentication](#) and [Setting up Client Authentication](#) to set up a mutual, two-way authentication process in which the server and the client authenticate each other before establishing a connection.

Although the typical setup described above uses the same CA certificate in both ends for mutual authentication, it does not have to be the case. The only requirements are that the CA certificate configured in the server must be able to validate the client certificate and the CA certificate imported into the client truststore must be able to validate the server certificate; the two CA certificates used on the two ends can be distinct.

Debugging an SSL Connection

JSSE provides debugging information to `stdout` when you set the system property `-Djavax.net.debug=all`. Java then tells you what keystores and truststores are being used, as well as what is going on during the SSL handshake and certificate exchange. That will be helpful when you are trying to debug a failed SSL connection.

6.8 Connecting Using Unix Domain Sockets

Connector/J does not natively support connections to MySQL Servers with Unix domain sockets. However, there is provision for using 3rd-party libraries that supply the function via a pluggable socket factory. Such a custom factory should implement the `com.mysql.cj.protocol.SocketFactory` interface or the legacy `com.mysql.jdbc.SocketFactory` interface of Connector/J. Follow these requirements when you use such a custom socket factory for Unix sockets :

- The MySQL Server must be configured with the system variable `--socket` (for native protocol connections using the JDBC API) or `--mysqlx-socket` (for X Protocol connections using the X DevAPI), which must contain the file path of the Unix socket file.
- The fully-qualified class name of the custom factory should be passed to Connector/J via the connection property `socketFactory`. For example, with the `unixsocket` library, set:

```
socketFactory=org.newsclub.net.mysql.AFUNIXDatabaseSocketFactory
```

You might also need to pass other parameters to the custom factory as connection properties. For example, for the `unixsocket` library, provide the file path of the socket file with the property `unixsocket.file`:

```
unixsocket.file=path_to_socket_file
```

- When using the X Protocol, set the connection property `xdevapi.useAsyncProtocol=false` (that is the default setting for Connector/J 8.0.12 and later). Unix socket is not supported for asynchronous socket channels. When `xdevapi.useAsyncProtocol=true`, the `socketFactory` property is ignored.

Note

For X Protocol connections, the provision to use custom socket factory for Unix socket connections is only available for Connector/J 8.0.12 and later.

6.9 Connecting Using Named Pipes

Note

Support for named pipes is not available for X Protocol connections.

Connector/J also supports access to MySQL using named pipes on Windows platforms with the `NamedPipeSocketFactory` as a plugin-sockets factory. If you do not use a `namedPipePath` property, the default of `'\\.\pipe\MySQL'` is used. If you use the `NamedPipeSocketFactory`, the host name and port number values in the JDBC URL are ignored. To enable this feature, set the `socketFactory` property:

```
socketFactory=com.mysql.cj.protocol.NamedPipeSocketFactory
```

Set this property, as well as the path of the named pipe, with the following connection URL:

```
jdbc:mysql:///test?socketFactory=com.mysql.cj.protocol.NamedPipeSocketFactory&namedPipePath=\\.\pipe\MySQL
```

To create your own socket factories, follow the sample code in `com.mysql.cj.protocol.NamedPipeSocketFactory` or `com.mysql.cj.protocol.StandardSocketFactory`.

An alternate approach is to use the following two properties in connection URLs for establishing named pipe connections on Windows platforms:

- `(protocol=pipe)` for named pipes (default value for the property is `tcp`).
- `(path=path_to_pipe)` for path of named pipes. Default value for the path is `\\.pipe\MySQL`.

The “address-equals” or “key-value” form of host specification (see [Single host \[20\]](#) for details) greatly simplifies the URL for a named pipe connection on Windows. For example, to use the default named pipe of `\\.pipe\MySQL`, just specify:

```
jdbc:mysql://address=(protocol=pipe)/test
```

To use the custom named pipe of `\\.pipe\MySQL80` :

```
jdbc:mysql://address=(protocol=pipe)(path=\\.pipe\MySQL80)/test
```

With `(protocol=pipe)`, the `NamedPipeSocketFactory` is automatically selected.

Named pipes only work when connecting to a MySQL server on the same physical machine where the JDBC driver is running. In simple performance tests, named pipe access is between 30%-50% faster than the standard TCP/IP access. However, this varies per system, and named pipes are slower than TCP/IP in many Windows configurations.

6.10 Connecting Using PAM Authentication

Java applications using Connector/J can connect to MySQL servers that use the pluggable authentication module (PAM) authentication scheme.

For PAM authentication to work, you must have the following:

- A MySQL server that supports PAM authentication. See [PAM Pluggable Authentication](#) for more information. Connector/J implements the same cleartext authentication method as in [Client-Side Cleartext Pluggable Authentication](#).
- SSL capability, as explained in [Section 6.7, “Connecting Securely Using SSL”](#). Because the PAM authentication scheme sends the original password to the server, the connection to the server must be encrypted.

PAM authentication support is enabled by default in Connector/J 8.0, so no extra configuration is needed.

To disable the PAM authentication feature, specify `mysql_clear_password` (the method) or `com.mysql.cj.protocol.a.authentication.MysqlClearPasswordPlugin` (the class name) in the comma-separated list of arguments for the `disabledAuthenticationPlugins` connection option. See [Section 6.3, “Configuration Properties”](#) for details about that connection option.

6.11 Using Master/Slave Replication with ReplicationConnection

See [Section 9.4, “Configuring Master/Slave Replication with Connector/J”](#) for details on the topic.

6.12 Mapping MySQL Error Numbers to JDBC SQLState Codes

The table below provides a mapping of the MySQL error numbers to JDBC `SQLState` values.

Table 6.4 Mapping of MySQL Error Numbers to SQLStates

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1022	ER_DUP_KEY	23000
1037	ER_OUTOFMEMORY	HY001
1038	ER_OUT_OF_SORTMEMORY	HY001
1040	ER_CON_COUNT_ERROR	08004
1042	ER_BAD_HOST_ERROR	08S01
1043	ER_HANDSHAKE_ERROR	08S01
1044	ER_DBACCESS_DENIED_ERROR	42000
1045	ER_ACCESS_DENIED_ERROR	28000
1046	ER_NO_DB_ERROR	3D000
1047	ER_UNKNOWN_COM_ERROR	08S01
1048	ER_BAD_NULL_ERROR	23000
1049	ER_BAD_DB_ERROR	42000
1050	ER_TABLE_EXISTS_ERROR	42S01
1051	ER_BAD_TABLE_ERROR	42S02
1052	ER_NON_UNIQ_ERROR	23000
1053	ER_SERVER_SHUTDOWN	08S01
1054	ER_BAD_FIELD_ERROR	42S22
1055	ER_WRONG_FIELD_WITH_GROUP	42000
1056	ER_WRONG_GROUP_FIELD	42000
1057	ER_WRONG_SUM_SELECT	42000
1058	ER_WRONG_VALUE_COUNT	21S01
1059	ER_TOO_LONG_IDENT	42000
1060	ER_DUP_FIELDNAME	42S21
1061	ER_DUP_KEYNAME	42000
1062	ER_DUP_ENTRY	23000
1063	ER_WRONG_FIELD_SPEC	42000
1064	ER_PARSE_ERROR	42000
1065	ER_EMPTY_QUERY	42000
1066	ER_NONUNIQ_TABLE	42000
1067	ER_INVALID_DEFAULT	42000
1068	ER_MULTIPLE_PRI_KEY	42000
1069	ER_TOO_MANY_KEYS	42000
1070	ER_TOO_MANY_KEY_PARTS	42000
1071	ER_TOO_LONG_KEY	42000
1072	ER_KEY_COLUMN_DOES_NOT_EXISTS	42000

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1073	ER_BLOB_USED_AS_KEY	42000
1074	ER_TOO_BIG_FIELDLENGTH	42000
1075	ER_WRONG_AUTO_KEY	42000
1080	ER_FORCING_CLOSE	08S01
1081	ER_IPSOCK_ERROR	08S01
1082	ER_NO_SUCH_INDEX	42S12
1083	ER_WRONG_FIELD_TERMINATORS	42000
1084	ER_BLOBS_AND_NO_TERMINATED	42000
1090	ER_CANT_REMOVE_ALL_FIELDS	42000
1091	ER_CANT_DROP_FIELD_OR_KEY	42000
1101	ER_BLOB_CANT_HAVE_DEFAULT	42000
1102	ER_WRONG_DB_NAME	42000
1103	ER_WRONG_TABLE_NAME	42000
1104	ER_TOO_BIG_SELECT	42000
1106	ER_UNKNOWN_PROCEDURE	42000
1107	ER_WRONG_PARAMCOUNT_TO_PROCEDURE	42000
1109	ER_UNKNOWN_TABLE	42S02
1110	ER_FIELD_SPECIFIED_TWICE	42000
1112	ER_UNSUPPORTED_EXTENSION	42000
1113	ER_TABLE_MUST_HAVE_COLUMNS	42000
1115	ER_UNKNOWN_CHARACTER_SET	42000
1118	ER_TOO_BIG_ROWSIZE	42000
1120	ER_WRONG_OUTER_JOIN	42000
1121	ER_NULL_COLUMN_IN_INDEX	42000
1131	ER_PASSWORD_ANONYMOUS_USER	42000
1132	ER_PASSWORD_NOT_ALLOWED	42000
1133	ER_PASSWORD_NO_MATCH	42000
1136	ER_WRONG_VALUE_COUNT_ON_ROW	21S01
1138	ER_INVALID_USE_OF_NULL	22004
1139	ER_REGEXP_ERROR	42000
1140	ER_MIX_OF_GROUP_FUNC_AND_FIELDS	42000
1141	ER_NONEXISTING_GRANT	42000
1142	ER_TABLEACCESS_DENIED_ERROR	42000
1143	ER_COLUMNACCESS_DENIED_ERROR	42000
1144	ER_ILLEGAL_GRANT_FOR_TABLE	42000
1145	ER_GRANT_WRONG_HOST_OR_USER	42000

Mapping MySQL Error Numbers to JDBC SQLState Codes

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1146	ER_NO_SUCH_TABLE	42S02
1147	ER_NONEXISTING_TABLE_GRANT	42000
1148	ER_NOT_ALLOWED_COMMAND	42000
1149	ER_SYNTAX_ERROR	42000
1152	ER_ABORTING_CONNECTION	08S01
1153	ER_NET_PACKET_TOO_LARGE	08S01
1154	ER_NET_READ_ERROR_FROM_PIPE	08S01
1155	ER_NET_FCNTL_ERROR	08S01
1156	ER_NET_PACKETS_OUT_OF_ORDER	08S01
1157	ER_NET_UNCOMPRESS_ERROR	08S01
1158	ER_NET_READ_ERROR	08S01
1159	ER_NET_READ_INTERRUPTED	08S01
1160	ER_NET_ERROR_ON_WRITE	08S01
1161	ER_NET_WRITE_INTERRUPTED	08S01
1162	ER_TOO_LONG_STRING	42000
1163	ER_TABLE_CANT_HANDLE_BLOB	42000
1164	ER_TABLE_CANT_HANDLE_AUTO_INCREMENT	42000
1166	ER_WRONG_COLUMN_NAME	42000
1167	ER_WRONG_KEY_COLUMN	42000
1169	ER_DUP_UNIQUE	23000
1170	ER_BLOB_KEY_WITHOUT_LENGTH	42000
1171	ER_PRIMARY_CANT_HAVE_NULL	42000
1172	ER_TOO_MANY_ROWS	42000
1173	ER_REQUIRES_PRIMARY_KEY	42000
1176	ER_KEY_DOES_NOT_EXISTS	42000
1177	ER_CHECK_NO_SUCH_TABLE	42000
1178	ER_CHECK_NOT_IMPLEMENTED	42000
1179	ER_CANT_DO_THIS_DURING_AN_TRANSACTION	25000
1184	ER_NEW_ABORTING_CONNECTION	08S01
1189	ER_MASTER_NET_READ	08S01
1190	ER_MASTER_NET_WRITE	08S01
1203	ER_TOO_MANY_USER_CONNECTIONS	42000
1205	ER_LOCK_WAIT_TIMEOUT	40001
1207	ER_READ_ONLY_TRANSACTION	25000
1211	ER_NO_PERMISSION_TO_CREATE_USER	42000
1213	ER_LOCK_DEADLOCK	40001

Mapping MySQL Error Numbers to JDBC SQLState Codes

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1216	ER_NO_REFERENCED_ROW	23000
1217	ER_ROW_IS_REFERENCED	23000
1218	ER_CONNECT_TO_MASTER	08S01
1222	ER_WRONG_NUMBER_OF_COLUMNS_IN_SELECT	21000
1226	ER_USER_LIMIT_REACHED	42000
1227	ER_SPECIFIC_ACCESS_DENIED_ERROR	42000
1230	ER_NO_DEFAULT	42000
1231	ER_WRONG_VALUE_FOR_VAR	42000
1232	ER_WRONG_TYPE_FOR_VAR	42000
1234	ER_CANT_USE_OPTION_HERE	42000
1235	ER_NOT_SUPPORTED_YET	42000
1239	ER_WRONG_FK_DEF	42000
1241	ER_OPERAND_COLUMNS	21000
1242	ER_SUBQUERY_NO_1_ROW	21000
1247	ER_ILLEGAL_REFERENCE	42S22
1248	ER_DERIVED_MUST_HAVE_ALIAS	42000
1249	ER_SELECT_REDUCED	01000
1250	ER_TABLENAME_NOT_ALLOWED_HERE	42000
1251	ER_NOT_SUPPORTED_AUTH_MODE	08004
1252	ER_SPATIAL_CANT_HAVE_NULL	42000
1253	ER_COLLATION_CHARSET_MISMATCH	42000
1261	ER_WARN_TOO_FEW_RECORDS	01000
1262	ER_WARN_TOO_MANY_RECORDS	01000
1263	ER_WARN_NULL_TO_NOTNULL	22004
1264	ER_WARN_DATA_OUT_OF_RANGE	22003
1265	ER_WARN_DATA_TRUNCATED	01000
1280	ER_WRONG_NAME_FOR_INDEX	42000
1281	ER_WRONG_NAME_FOR_CATALOG	42000
1286	ER_UNKNOWN_STORAGE_ENGINE	42000
1292	ER_TRUNCATED_WRONG_VALUE	22007
1303	ER_SP_NO_RECURSIVE_CREATE	2F003
1304	ER_SP_ALREADY_EXISTS	42000
1305	ER_SP_DOES_NOT_EXIST	42000
1308	ER_SP_LILABEL_MISMATCH	42000
1309	ER_SP_LABEL_REDEFINE	42000
1310	ER_SP_LABEL_MISMATCH	42000

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1311	ER_SP_UNINIT_VAR	01000
1312	ER_SP_BADSELECT	0A000
1313	ER_SP_BADRETURN	42000
1314	ER_SP_BADSTATEMENT	0A000
1315	ER_UPDATE_LOG_DEPRECATED_IGNORED	42000
1316	ER_UPDATE_LOG_DEPRECATED_TRANSLATED	42000
1317	ER_QUERY_INTERRUPTED	70100
1318	ER_SP_WRONG_NO_OF_ARGS	42000
1319	ER_SP_COND_MISMATCH	42000
1320	ER_SP_NORETURN	42000
1321	ER_SP_NORETURNEND	2F005
1322	ER_SP_BAD_CURSOR_QUERY	42000
1323	ER_SP_BAD_CURSOR_SELECT	42000
1324	ER_SP_CURSOR_MISMATCH	42000
1325	ER_SP_CURSOR_ALREADY_OPEN	24000
1326	ER_SP_CURSOR_NOT_OPEN	24000
1327	ER_SP_UNDECLARED_VAR	42000
1329	ER_SP_FETCH_NO_DATA	02000
1330	ER_SP_DUP_PARAM	42000
1331	ER_SP_DUP_VAR	42000
1332	ER_SP_DUP_COND	42000
1333	ER_SP_DUP_CURS	42000
1335	ER_SP_SUBSELECT_NYI	0A000
1336	ER_STMT_NOT_ALLOWED_IN_SF_OR_TRG	0A000
1337	ER_SP_VARCOND_AFTER_CURSHNDLR	42000
1338	ER_SP_CURSOR_AFTER_HANDLER	42000
1339	ER_SP_CASE_NOT_FOUND	20000
1365	ER_DIVISION_BY_ZERO	22012
1367	ER_ILLEGAL_VALUE_FOR_TYPE	22007
1370	ER_PROCACCESS_DENIED_ERROR	42000
1397	ER_XAER_NOTA	XAE04
1398	ER_XAER_INVAL	XAE05
1399	ER_XAER_RMFAIL	XAE07
1400	ER_XAER_OUTSIDE	XAE09
1401	ER_XA_RMERR	XAE03
1402	ER_XA_RBROLLBACK	XA100

Mapping MySQL Error Numbers to JDBC SQLState Codes

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1403	ER_NONEXISTING_PROC_GRANT	42000
1406	ER_DATA_TOO_LONG	22001
1407	ER_SP_BAD_SQLSTATE	42000
1410	ER_CANT_CREATE_USER_WITH_GRANT	42000
1413	ER_SP_DUP_HANDLER	42000
1414	ER_SP_NOT_VAR_ARG	42000
1415	ER_SP_NO_RETSET	0A000
1416	ER_CANT_CREATE_GEOMETRY_OBJECT	22003
1425	ER_TOO_BIG_SCALE	42000
1426	ER_TOO_BIG_PRECISION	42000
1427	ER_M_BIGGER_THAN_D	42000
1437	ER_TOO_LONG_BODY	42000
1439	ER_TOO_BIG_DISPLAYWIDTH	42000
1440	ER_XAER_DUPID	XAE08
1441	ER_DATETIME_FUNCTION_OVERFLOW	22008
1451	ER_ROW_IS_REFERENCED_2	23000
1452	ER_NO_REFERENCED_ROW_2	23000
1453	ER_SP_BAD_VAR_SHADOW	42000
1458	ER_SP_WRONG_NAME	42000
1460	ER_SP_NO_AGGREGATE	42000
1461	ER_MAX_PREPARED_STMT_COUNT_REACHED	42000
1463	ER_NON_GROUPING_FIELD_USED	42000
1557	ER_FOREIGN_DUPLICATE_KEY	23000
1568	ER_CANT_CHANGE_TX_ISOLATION	25001
1582	ER_WRONG_PARAMCOUNT_TO_NATIVE_FCT	42000
1583	ER_WRONG_PARAMETERS_TO_NATIVE_FCT	42000
1584	ER_WRONG_PARAMETERS_TO_STORED_FCT	42000
1586	ER_DUP_ENTRY_WITH_KEY_NAME	23000
1613	ER_XA_RBTIMEOUT	XA106
1614	ER_XA_RBDEADLOCK	XA102
1630	ER_FUNC_INEXISTENT_NAME_COLLISION	42000
1641	ER_DUP_SIGNAL_SET	42000
1642	ER_SIGNAL_WARN	01000
1643	ER_SIGNAL_NOT_FOUND	02000
1645	ER_RESIGNAL_WITHOUT_ACTIVE_HANDLER	0K000
1687	ER_SPATIAL_MUST_HAVE_GEOM_COL	42000

MySQL Error Number	MySQL Error Name	SQL Standard SQLState
1690	ER_DATA_OUT_OF_RANGE	22003
1698	ER_ACCESS_DENIED_NO_PASSWORD_ERROR	28000
1701	ER_TRUNCATE_ILLEGAL_FK	42000
1758	ER_DA_INVALID_CONDITION_NUMBER	35000
1761	ER_FOREIGN_DUPLICATE_KEY_WITH_CHILD_INFO	23000
1762	ER_FOREIGN_DUPLICATE_KEY_WITHOUT_CHILD_INFO	23000
1792	ER_CANT_EXECUTE_IN_READ_ONLY_TRANSACTION	25006
1845	ER_ALTER_OPERATION_NOT_SUPPORTED	0A000
1846	ER_ALTER_OPERATION_NOT_SUPPORTED_REASON	0A000
1859	ER_DUP_UNKNOWN_IN_INDEX	23000
1873	ER_ACCESS_DENIED_CHANGE_USER_ERROR	28000
1887	ER_GET_STACKED_DA_WITHOUT_ACTIVE_HANDLER	0Z002
1903	ER_INVALID_ARGUMENT_FOR_LOGARITHM	2201E

Chapter 7 JDBC Concepts

Table of Contents

7.1 Connecting to MySQL Using the JDBC <code>DriverManager</code> Interface	73
7.2 Using JDBC <code>Statement</code> Objects to Execute SQL	74
7.3 Using JDBC <code>CallableStatements</code> to Execute Stored Procedures	75
7.4 Retrieving <code>AUTO_INCREMENT</code> Column Values through JDBC	78

This section provides some general JDBC background.

7.1 Connecting to MySQL Using the JDBC `DriverManager` Interface

When you are using JDBC outside of an application server, the `DriverManager` class manages the establishment of connections.

Specify to the `DriverManager` which JDBC drivers to try to make Connections with. The easiest way to do this is to use `Class.forName()` on the class that implements the `java.sql.Driver` interface. With MySQL Connector/J, the name of this class is `com.mysql.cj.jdbc.Driver`. With this method, you could use an external configuration file to supply the driver class name and driver parameters to use when connecting to a database.

The following section of Java code shows how you might register MySQL Connector/J from the `main()` method of your application. If testing this code, first read the installation section at [Chapter 4, Connector/J Installation](#), to make sure you have connector installed correctly and the `CLASSPATH` set up. Also, ensure that MySQL is configured to accept external TCP/IP connections.

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;

// Notice, do not import com.mysql.cj.jdbc.*
// or you will have problems!

public class LoadDriver {
    public static void main(String[] args) {
        try {
            // The newInstance() call is a work around for some
            // broken Java implementations

            Class.forName("com.mysql.cj.jdbc.Driver").newInstance();
        } catch (Exception ex) {
            // handle the error
        }
    }
}
```

After the driver has been registered with the `DriverManager`, you can obtain a `Connection` instance that is connected to a particular database by calling `DriverManager.getConnection()`:

Example 7.1 Connector/J: Obtaining a connection from the `DriverManager`

If you have not already done so, please review the portion of [Section 7.1, “Connecting to MySQL Using the JDBC `DriverManager` Interface”](#) above before working with the example below.

This example shows how you can obtain a `Connection` instance from the `DriverManager`. There are a few different signatures for the `getConnection()` method. Consult the API documentation that comes with your JDK for more specific information on how to use them.

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;

Connection conn = null;
...
try {
    conn =
        DriverManager.getConnection("jdbc:mysql://localhost/test?" +
                                   "user=minty&password=greatsqlldb");

    // Do something with the Connection

    ...
} catch (SQLException ex) {
    // handle any errors
    System.out.println("SQLException: " + ex.getMessage());
    System.out.println("SQLState: " + ex.getSQLState());
    System.out.println("VendorError: " + ex.getErrorCode());
}
```

Once a `Connection` is established, it can be used to create `Statement` and `PreparedStatement` objects, as well as retrieve metadata about the database. This is explained in the following sections.

7.2 Using JDBC `Statement` Objects to Execute SQL

`Statement` objects allow you to execute basic SQL queries and retrieve the results through the `ResultSet` class, which is described later.

To create a `Statement` instance, you call the `createStatement()` method on the `Connection` object you have retrieved using one of the `DriverManager.getConnection()` or `DataSource.getConnection()` methods described earlier.

Once you have a `Statement` instance, you can execute a `SELECT` query by calling the `executeQuery(String)` method with the SQL you want to use.

To update data in the database, use the `executeUpdate(String SQL)` method. This method returns the number of rows matched by the update statement, not the number of rows that were modified.

If you do not know ahead of time whether the SQL statement will be a `SELECT` or an `UPDATE/INSERT`, then you can use the `execute(String SQL)` method. This method will return true if the SQL query was a `SELECT`, or false if it was an `UPDATE`, `INSERT`, or `DELETE` statement. If the statement was a `SELECT` query, you can retrieve the results by calling the `getResultSet()` method. If the statement was an `UPDATE`, `INSERT`, or `DELETE` statement, you can retrieve the affected rows count by calling `getUpdateCount()` on the `Statement` instance.

Example 7.2 Connector/J: Using `java.sql.Statement` to execute a `SELECT` query

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;
import java.sql.Statement;
import java.sql.ResultSet;

// assume that conn is an already created JDBC connection (see previous examples)
```



```

Statement stmt = null;
ResultSet rs = null;

try {
    stmt = conn.createStatement();
    rs = stmt.executeQuery("SELECT foo FROM bar");

    // or alternatively, if you don't know ahead of time that
    // the query will be a SELECT...

    if (stmt.execute("SELECT foo FROM bar")) {
        rs = stmt.getResultSet();
    }

    // Now do something with the ResultSet ....
}
catch (SQLException ex){
    // handle any errors
    System.out.println("SQLException: " + ex.getMessage());
    System.out.println("SQLState: " + ex.getSQLState());
    System.out.println("VendorError: " + ex.getErrorCode());
}
finally {
    // it is a good idea to release
    // resources in a finally{} block
    // in reverse-order of their creation
    // if they are no-longer needed

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException sqlEx) { } // ignore

        rs = null;
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlEx) { } // ignore

        stmt = null;
    }
}

```

7.3 Using JDBC `CallableStatements` to Execute Stored Procedures

Connector/J fully implements the `java.sql.CallableStatement` interface.

For more information on MySQL stored procedures, please refer to [Using Stored Routines \(Procedures and Functions\)](#).

Connector/J exposes stored procedure functionality through JDBC's `CallableStatement` interface.

The following example shows a stored procedure that returns the value of `inOutParam` incremented by 1, and the string passed in using `inputParam` as a `ResultSet`:

Example 7.3 Connector/J: Calling Stored Procedures

```
CREATE PROCEDURE demoSp(IN inputParam VARCHAR(255), \
```

```

                                INOUT inOutParam INT)
BEGIN
  DECLARE z INT;
  SET z = inOutParam + 1;
  SET inOutParam = z;

  SELECT inputParam;

  SELECT CONCAT('zyxw', inputParam);
END

```

To use the `demoSp` procedure with Connector/J, follow these steps:

1. Prepare the callable statement by using `Connection.prepareCall()`.

Notice that you have to use JDBC escape syntax, and that the parentheses surrounding the parameter placeholders are not optional:

Example 7.4 Connector/J: Using `Connection.prepareCall()`

```

import java.sql.CallableStatement;

...

//
// Prepare a call to the stored procedure 'demoSp'
// with two parameters
//
// Notice the use of JDBC-escape syntax ({call ...})
//

CallableStatement cStmt = conn.prepareCall("{call demoSp(?, ?)}");

cStmt.setString(1, "abcdefg");

```

Note

`Connection.prepareCall()` is an expensive method, due to the metadata retrieval that the driver performs to support output parameters. For performance reasons, minimize unnecessary calls to `Connection.prepareCall()` by reusing `CallableStatement` instances in your code.

2. Register the output parameters (if any exist)

To retrieve the values of output parameters (parameters specified as `OUT` or `INOUT` when you created the stored procedure), JDBC requires that they be specified before statement execution using the various `registerOutputParameter()` methods in the `CallableStatement` interface:

Example 7.5 Connector/J: Registering output parameters

```

import java.sql.Types;

...
//
// Connector/J supports both named and indexed
// output parameters. You can register output
// parameters using either method, as well
// as retrieve output parameters using either
// method, regardless of what method was
// used to register them.

```

```

//
// The following examples show how to use
// the various methods of registering
// output parameters (you should of course
// use only one registration per parameter).
//
//
// Registers the second parameter as output, and
// uses the type 'INTEGER' for values returned from
// getObject()
//
cStmt.registerOutParameter(2, Types.INTEGER);
//
// Registers the named parameter 'inOutParam', and
// uses the type 'INTEGER' for values returned from
// getObject()
//
cStmt.registerOutParameter("inOutParam", Types.INTEGER);
...

```

3. Set the input parameters (if any exist)

Input and in/out parameters are set as for `PreparedStatement` objects. However, `CallableStatement` also supports setting parameters by name:

Example 7.6 Connector/J: Setting `CallableStatement` input parameters

```

...
//
// Set a parameter by index
//
cStmt.setString(1, "abcdefg");
//
// Alternatively, set a parameter using
// the parameter name
//
cStmt.setString("inputParam", "abcdefg");
//
// Set the 'in/out' parameter using an index
//
cStmt.setInt(2, 1);
//
// Alternatively, set the 'in/out' parameter
// by name
//
cStmt.setInt("inOutParam", 1);
...

```

4. Execute the `CallableStatement`, and retrieve any result sets or output parameters.

Although `CallableStatement` supports calling any of the `Statement` execute methods (`executeUpdate()`, `executeQuery()` or `execute()`), the most flexible method to call is `execute()`, as you do not need to know ahead of time if the stored procedure returns result sets:

Example 7.7 Connector/J: Retrieving results and output parameter values

```
...
boolean hadResults = cStmt.execute();

//
// Process all returned result sets
//

while (hadResults) {
    ResultSet rs = cStmt.getResultSet();

    // process result set
    ...

    hadResults = cStmt.getMoreResults();
}

//
// Retrieve output parameters
//
// Connector/J supports both index-based and
// name-based retrieval
//

int outputValue = cStmt.getInt(2); // index-based

outputValue = cStmt.getInt("inOutParam"); // name-based

...
```

7.4 Retrieving `AUTO_INCREMENT` Column Values through JDBC

`getGeneratedKeys()` is the preferred method to use if you need to retrieve `AUTO_INCREMENT` keys and through JDBC; this is illustrated in the first example below. The second example shows how you can retrieve the same value using a standard `SELECT LAST_INSERT_ID()` query. The final example shows how updatable result sets can retrieve the `AUTO_INCREMENT` value when using the `insertRow()` method.

Example 7.8 Connector/J: Retrieving `AUTO_INCREMENT` column values using `Statement.getGeneratedKeys()`

```
Statement stmt = null;
ResultSet rs = null;

try {

    //
    // Create a Statement instance that we can use for
    // 'normal' result sets assuming you have a
    // Connection 'conn' to a MySQL database already
    // available

    stmt = conn.createStatement();

    //
```

```

// Issue the DDL queries for the table for this example
//
stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
stmt.executeUpdate(
    "CREATE TABLE autoIncTutorial ("
    + "priKey INT NOT NULL AUTO_INCREMENT, "
    + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

//
// Insert one row that will generate an AUTO INCREMENT
// key in the 'priKey' field
//
stmt.executeUpdate(
    "INSERT INTO autoIncTutorial (dataField) "
    + "values ('Can I Get the Auto Increment Field?')",
    Statement.RETURN_GENERATED_KEYS);

//
// Example of using Statement.getGeneratedKeys()
// to retrieve the value of an auto-increment
// value
//

int autoIncKeyFromApi = -1;

rs = stmt.getGeneratedKeys();

if (rs.next()) {
    autoIncKeyFromApi = rs.getInt(1);
} else {

    // throw an exception from here
}

System.out.println("Key returned from getGeneratedKeys():"
    + autoIncKeyFromApi);
} finally {

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException ex) {
            // ignore
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException ex) {
            // ignore
        }
    }
}
}

```

Example 7.9 Connector/J: Retrieving `AUTO_INCREMENT` column values using `SELECT LAST_INSERT_ID()`

```

Statement stmt = null;
ResultSet rs = null;

try {

    //

```

```

// Create a Statement instance that we can use for
// 'normal' result sets.

stmt = conn.createStatement();

//
// Issue the DDL queries for the table for this example
//

stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
stmt.executeUpdate(
    "CREATE TABLE autoIncTutorial ( "
    + "priKey INT NOT NULL AUTO_INCREMENT, "
    + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

//
// Insert one row that will generate an AUTO INCREMENT
// key in the 'priKey' field
//

stmt.executeUpdate(
    "INSERT INTO autoIncTutorial (dataField) "
    + "values ('Can I Get the Auto Increment Field?')");

//
// Use the MySQL LAST_INSERT_ID()
// function to do the same thing as getGeneratedKeys()
//

int autoIncKeyFromFunc = -1;
rs = stmt.executeQuery("SELECT LAST_INSERT_ID()");

if (rs.next()) {
    autoIncKeyFromFunc = rs.getInt(1);
} else {
    // throw an exception from here
}

System.out.println("Key returned from " +
    "'SELECT LAST_INSERT_ID()': " +
    autoIncKeyFromFunc);
} finally {

    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException ex) {
            // ignore
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException ex) {
            // ignore
        }
    }
}
}

```

Example 7.10 Connector/J: Retrieving `AUTO_INCREMENT` column values in `Updatable ResultSets`

```

Statement stmt = null;
ResultSet rs = null;

```

```
try {  
  
    //  
    // Create a Statement instance that we can use for  
    // 'normal' result sets as well as an 'updatable'  
    // one, assuming you have a Connection 'conn' to  
    // a MySQL database already available  
    //  
    stmt = conn.createStatement(java.sql.ResultSet.TYPE_FORWARD_ONLY,  
                                java.sql.ResultSet.CONCUR_UPDATABLE);  
  
    //  
    // Issue the DDL queries for the table for this example  
    //  
    stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");  
    stmt.executeUpdate(  
        "CREATE TABLE autoIncTutorial ("  
        + "priKey INT NOT NULL AUTO_INCREMENT, "  
        + "dataField VARCHAR(64), PRIMARY KEY (priKey))");  
  
    //  
    // Example of retrieving an AUTO INCREMENT key  
    // from an updatable result set  
    //  
    rs = stmt.executeQuery("SELECT priKey, dataField "  
        + "FROM autoIncTutorial");  
  
    rs.moveToInsertRow();  
  
    rs.updateString("dataField", "AUTO INCREMENT here?");  
    rs.insertRow();  
  
    //  
    // the driver adds rows at the end  
    //  
    rs.last();  
  
    //  
    // We should now be on the row we just inserted  
    //  
    int autoIncKeyFromRS = rs.getInt("priKey");  
  
    System.out.println("Key returned for inserted row: "  
        + autoIncKeyFromRS);  
} finally {  
  
    if (rs != null) {  
        try {  
            rs.close();  
        } catch (SQLException ex) {  
            // ignore  
        }  
    }  
  
    if (stmt != null) {  
        try {  
            stmt.close();  
        } catch (SQLException ex) {  
            // ignore  
        }  
    }  
}
```

```
}
```

Running the preceding example code should produce the following output:

```
Key returned from getGeneratedKeys(): 1  
Key returned from SELECT LAST_INSERT_ID(): 1  
Key returned for inserted row: 1
```

At times, it can be tricky to use the `SELECT LAST_INSERT_ID()` query, as that function's value is scoped to a connection. So, if some other query happens on the same connection, the value is overwritten. On the other hand, the `getGeneratedKeys()` method is scoped by the `Statement` instance, so it can be used even if other queries happen on the same connection, but not on the same `Statement` instance.

Chapter 8 Connection Pooling with Connector/J

Connection pooling is a technique of creating and managing a pool of connections that are ready for use by any [thread](#) that needs them. Connection pooling can greatly increase the performance of your Java application, while reducing overall resource usage.

How Connection Pooling Works

Most applications only need a thread to have access to a JDBC connection when they are actively processing a [transaction](#), which often takes only milliseconds to complete. When not processing a transaction, the connection sits idle. Connection pooling enables the idle connection to be used by some other thread to do useful work.

In practice, when a thread needs to do work against a MySQL or other database with JDBC, it requests a connection from the pool. When the thread is finished using the connection, it returns it to the pool, so that it can be used by any other threads.

When the connection is loaned out from the pool, it is used exclusively by the thread that requested it. From a programming point of view, it is the same as if your thread called `DriverManager.getConnection()` every time it needed a JDBC connection. With connection pooling, your thread may end up using either a new connection or an already-existing connection.

Benefits of Connection Pooling

The main benefits to connection pooling are:

- Reduced connection creation time.

Although this is not usually an issue with the quick connection setup that MySQL offers compared to other databases, creating new JDBC connections still incurs networking and JDBC driver overhead that will be avoided if connections are recycled.

- Simplified programming model.

When using connection pooling, each individual thread can act as though it has created its own JDBC connection, allowing you to use straightforward JDBC programming techniques.

- Controlled resource usage.

If you create a new connection every time a thread needs one rather than using connection pooling, your application's resource usage can be wasteful, and it could lead to unpredictable behaviors for your application when it is under a heavy load.

Using Connection Pooling with Connector/J

The concept of connection pooling in JDBC has been standardized through the JDBC 2.0 Optional interfaces, and all major application servers have implementations of these APIs that work with MySQL Connector/J.

Generally, you configure a connection pool in your application server configuration files, and access it through the Java Naming and Directory Interface (JNDI). The following code shows how you might use a connection pool from an application deployed in a J2EE application server:

Example 8.1 Connector/J: Using a connection pool with a J2EE application server

```
import java.sql.Connection;  
import java.sql.SQLException;
```

```
import java.sql.Statement;

import javax.naming.InitialContext;
import javax.sql.DataSource;

public class MyServletJspOrEjb {

    public void doSomething() throws Exception {
        /*
         * Create a JNDI Initial context to be able to
         * lookup the DataSource
         *
         * In production-level code, this should be cached as
         * an instance or static variable, as it can
         * be quite expensive to create a JNDI context.
         *
         * Note: This code only works when you are using servlets
         * or EJBs in a J2EE application server. If you are
         * using connection pooling in standalone Java code, you
         * will have to create/configure datasources using whatever
         * mechanisms your particular connection pooling library
         * provides.
         */

        InitialContext ctx = new InitialContext();

        /*
         * Lookup the DataSource, which will be backed by a pool
         * that the application server provides. DataSource instances
         * are also a good candidate for caching as an instance
         * variable, as JNDI lookups can be expensive as well.
         */

        DataSource ds =
            (DataSource)ctx.lookup("java:comp/env/jdbc/MySQLDB");

        /*
         * The following code is what would actually be in your
         * Servlet, JSP or EJB 'service' method...where you need
         * to work with a JDBC connection.
         */

        Connection conn = null;
        Statement stmt = null;

        try {
            conn = ds.getConnection();

            /*
             * Now, use normal JDBC programming to work with
             * MySQL, making sure to close each resource when you're
             * finished with it, which permits the connection pool
             * resources to be recovered as quickly as possible
             */

            stmt = conn.createStatement();
            stmt.execute("SOME SQL QUERY");

            stmt.close();
            stmt = null;

            conn.close();
            conn = null;
        } finally {
            /*
             * close any jdbc instances here that weren't

```

```
    * explicitly closed during normal code path, so
    * that we don't 'leak' resources...
    */

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlex) {
            // ignore, as we can't do anything about it here
        }

        stmt = null;
    }

    if (conn != null) {
        try {
            conn.close();
        } catch (SQLException sqlex) {
            // ignore, as we can't do anything about it here
        }

        conn = null;
    }
}
}
```

As shown in the example above, after obtaining the JNDI [InitialContext](#), and looking up the [DataSource](#), the rest of the code follows familiar JDBC conventions.

When using connection pooling, always make sure that connections, and anything created by them (such as statements or result sets) are closed. This rule applies no matter what happens in your code (exceptions, flow-of-control, and so forth). When these objects are closed, they can be re-used; otherwise, they will be stranded, which means that the MySQL server resources they represent (such as buffers, locks, or sockets) are tied up for some time, or in the worst case can be tied up forever.

Sizing the Connection Pool

Each connection to MySQL has overhead (memory, CPU, context switches, and so forth) on both the client and server side. Every connection limits how many resources there are available to your application as well as the MySQL server. Many of these resources will be used whether or not the connection is actually doing any useful work! Connection pools can be tuned to maximize performance, while keeping resource utilization below the point where your application will start to fail rather than just run slower.

The optimal size for the connection pool depends on anticipated load and average database transaction time. In practice, the optimal connection pool size can be smaller than you might expect. If you take Oracle's Java Petstore blueprint application for example, a connection pool of 15-20 connections can serve a relatively moderate load (600 concurrent users) using MySQL and Tomcat with acceptable response times.

To correctly size a connection pool for your application, create load test scripts with tools such as Apache JMeter or The Grinder, and load test your application.

An easy way to determine a starting point is to configure your connection pool's maximum number of connections to be unbounded, run a load test, and measure the largest amount of concurrently used connections. You can then work backward from there to determine what values of minimum and maximum pooled connections give the best performance for your particular application.

Validating Connections

MySQL Connector/J can validate the connection by executing a lightweight ping against a server. In the case of load-balanced connections, this is performed against all active pooled internal connections that are retained. This is beneficial to Java applications using connection pools, as the pool can use this feature to validate connections. Depending on your connection pool and configuration, this validation can be carried out at different times:

1. Before the pool returns a connection to the application.
2. When the application returns a connection to the pool.
3. During periodic checks of idle connections.

To use this feature, specify a validation query in your connection pool that starts with `/* ping */`. Note that the syntax must be exactly as specified. This will cause the driver send a ping to the server and return a dummy lightweight result set. When using a [ReplicationConnection](#) or [LoadBalancedConnection](#), the ping will be sent across all active connections.

It is critical that the syntax be specified correctly. The syntax needs to be exact for reasons of efficiency, as this test is done for every statement that is executed:

```
protected static final String PING_MARKER = "/* ping */";
...
if (sql.charAt(0) == '/') {
if (sql.startsWith(PING_MARKER)) {
doPingInstead();
...
}
```

None of the following snippets will work, because the ping syntax is sensitive to whitespace, capitalization, and placement:

```
sql = "/* PING */ SELECT 1";
sql = "SELECT 1 /* ping*/";
sql = "/*ping*/ SELECT 1";
sql = " /* ping */ SELECT 1";
sql = "/*to ping or not to ping*/ SELECT 1";
```

All of the previous statements will issue a normal `SELECT` statement and will **not** be transformed into the lightweight ping. Further, for load-balanced connections, the statement will be executed against one connection in the internal pool, rather than validating each underlying physical connection. This results in the non-active physical connections assuming a stale state, and they may die. If Connector/J then re-balances, it might select a dead connection, resulting in an exception being passed to the application. To help prevent this, you can use [loadBalanceValidateConnectionOnSwapServer](#) to validate the connection before use.

If your Connector/J deployment uses a connection pool that allows you to specify a validation query, take advantage of it, but ensure that the query starts *exactly* with `/* ping */`. This is particularly important if you are using the load-balancing or replication-aware features of Connector/J, as it will help keep alive connections which otherwise will go stale and die, causing problems later.

Chapter 9 Multi-Host Connections

Table of Contents

9.1 Configuring Server Failover	87
9.2 Configuring Client-Side Failover when using the X Protocol	90
9.3 Configuring Load Balancing with Connector/J	90
9.4 Configuring Master/Slave Replication with Connector/J	93
9.5 Advanced Load-balancing and Failover Configuration	96

The following sections discuss a number of topics that involve multi-host connections, namely, server load-balancing, failover, and replication.

Developers should know the following things about multi-host connections that are managed through Connector/J:

- Each multi-host connection is a wrapper of the underlying physical connections.
- Each of the underlying physical connections has its own session. Sessions cannot be tracked, shared, or copied, given the MySQL architecture.
- Every switch between physical connections means a switch between sessions.
- Within a transaction boundary, there are no switches between physical connections. Beyond a transaction boundary, there is no guarantee that a switch does not occur.

Note

If an application reuses session-scope data (for example, variables, SSPs) beyond a transaction boundary, failures are possible, as a switch between the physical connections (which is also a switch between sessions) might occur. Therefore, the application should re-prepare the session data and also restart the last transaction in case of an exception, or it should re-prepare session data for each new transaction if it does not want to deal with exception handling.

9.1 Configuring Server Failover

MySQL Connector/J supports server failover. A failover happens when connection-related errors occur for an underlying, active connection. The connection errors are, by default, propagated to the client, which has to handle them by, for example, recreating the working objects (`Statement`, `ResultSet`, etc.) and restarting the processes. Sometimes, the driver might eventually fall back to the original host automatically before the client application continues to run, in which case the host switch is transparent and the client application will not even notice it.

A connection using failover support works just like a standard connection: the client does not experience any disruptions in the failover process. This means the client can rely on the same connection instance even if two successive statements might be executed on two different physical hosts. However, this does not mean the client does not have to deal with the exception that triggered the server switch.

The failover is configured at the initial setup stage of the server connection by the connection URL (see explanations for its format [here](#)):

```
jdbc:mysql://[primary host][:port],[secondary host 1][:port],[secondary host 2][:port]]...[[database]]»  
[?propertyName1=propertyValue1[&propertyName2=propertyValue2]...]
```

The host list in the connection URL comprises of two types of hosts, the primary and the secondary. When starting a new connection, the driver always tries to connect to the primary host first and, if required, fails over to the secondary hosts on the list sequentially when communication problems are experienced. Even if the initial connection to the primary host fails and the driver gets connected to a secondary host, the primary host never loses its special status: for example, it can be configured with an access mode distinct from those of the secondary hosts, and it can be put on a higher priority when a host is to be picked during a failover process.

The failover support is configured by the following connection properties (their functions are explained in the paragraphs below):

- `failOverReadOnly`
- `secondsBeforeRetryMaster`
- `queriesBeforeRetryMaster`
- `retriesAllDown`
- `autoReconnect`
- `autoReconnectForPools`

Configuring Connection Access Mode

As with any standard connection, the initial connection to the primary host is in read/write mode. However, if the driver fails to establish the initial connection to the primary host and it automatically switches to the next host on the list, the access mode now depends on the value of the property `failOverReadOnly`, which is “true” by default. The same happens if the driver is initially connected to the primary host and, because of some connection failure, it fails over to a secondary host. Every time the connection falls back to the primary host, its access mode will be read/write, irrespective of whether or not the primary host has been connected to before. The connection access mode can be changed any time at runtime by calling the method `Connection.setReadOnly(boolean)`, which partially overrides the property `failOverReadOnly`. When `failOverReadOnly=false` and the access mode is explicitly set to either true or false, it becomes the mode for every connection after a host switch, no matter what host type are being connected to; but, if `failOverReadOnly=true`, changing the access mode to read/write is only possible if the driver is connecting to the primary host; however, even if the access mode cannot be changed for the current connection, the driver remembers the client's last intention and, when falling back to the primary host, that is the mode that will be used. For an illustration, see the following successions of events with a two-host connection.

- Sequence A, with `failOverReadOnly=true`:
 1. Connects to primary host in read/write mode
 2. Sets `Connection.setReadOnly(true)`; primary host now in read-only mode
 3. Failover event; connects to secondary host in read-only mode
 4. Sets `Connection.setReadOnly(false)`; secondary host remains in read-only mode
 5. Falls back to primary host; connection now in read/write mode
- Sequence B, with `failOverReadOnly=false`

1. Connects to primary host in read/write mode
2. Sets `Connection.setReadOnly(true)`; primary host now in read-only mode
3. Failover event; connects to secondary host in read-only mode
4. Set `Connection.setReadOnly(false)`; connection to secondary host switches to read/write mode
5. Falls back to primary host; connection now in read/write mode

The difference between the two scenarios is in step 4: the access mode for the secondary host in sequence A does not change at that step, but the driver remembers and uses the set mode when falling back to the primary host, which would be read-only otherwise; but in sequence B, the access mode for the secondary host changes immediately.

Configuring Fallback to Primary Host

As already mentioned, the primary host is special in the failover arrangement when it comes to the host's access mode. Additionally, the driver tries to fall back to the primary host as soon as possible by default, even if no communication exception occurs. Two properties, `secondsBeforeRetryMaster` and `queriesBeforeRetryMaster`, determine when the driver is ready to retry a reconnection to the primary host (the `Master` in the property names stands for the primary host of our connection URL, which is not necessarily a master host in a replication setup):

- `secondsBeforeRetryMaster` determines how much time the driver waits before trying to fall back to the primary host
- `queriesBeforeRetryMaster` determines the number of queries that are executed before the driver tries to fall back to the primary host. Note that for the driver, each call to a `Statement.execute*()` method increments the query execution counter; therefore, when calls are made to `Statement.executeBatch()` or if `allowMultiQueries` or `rewriteBatchStatements` are enabled, the driver may not have an accurate count of the actual number of queries executed on the server. Also, the driver calls the `Statement.execute*()` methods internally in several occasions. All these mean you can only use `queriesBeforeRetryMaster` only as a coarse specification for when to fall back to the primary host.

In general, an attempt to fallback to the primary host is made when at least one of the conditions specified by the two properties is met, and the attempt always takes place at transaction boundaries. However, if auto-commit is turned off, the check happens only when the method `Connection.commit()` or `Connection.rollback()` is called. The automatic fallback to the primary host can be turned off by setting simultaneously `secondsBeforeRetryMaster` and `queriesBeforeRetryMaster` to "0". Setting only one of the properties to "0" only disables one part of the check.

Configuring Reconnection Attempts

When establishing a new connection or when a failover event occurs, the driver tries to connect successively to the next candidate on the host list. When the end of the list has been reached, it restarts all over again from the beginning of the list; however, the primary host is skipped over, if (a) NOT all the secondary hosts have already been tested at least once, AND (b) the fallback conditions defined by `secondsBeforeRetryMaster` and `queriesBeforeRetryMaster` are not yet fulfilled. Each run-through of the whole host list, (which is not necessarily completed at the end of the host list) counts as a single connection attempt. The driver tries as many connection attempts as specified by the value of the property `retriesAllDown`.

Seamless Reconnection

Although not recommended, you can make the driver perform failovers without invalidating the active `Statement` or `ResultSet` instances by setting either the parameter `autoReconnect` or `autoReconnectForPools` to `true`. This allows the client to continue using the same object instances after a failover event, without taking any exceptional measures. This, however, may lead to unexpected results: for example, if the driver is connected to the primary host with read/write access mode and it fails over to a secondary host in read-only mode, further attempts to issue data-changing queries will result in errors, and the client will not be aware of that. This limitation is particularly relevant when using data streaming: after the failover, the `ResultSet` looks to be alright, but the underlying connection may have changed already, and no backing cursor is available anymore.

9.2 Configuring Client-Side Failover when using the X Protocol

When using the X Protocol, Connector/J supports a client-side failover feature for establishing a Session. If multiple hosts are specified in the connection URL, when Connector/J fails to connect to a listed host, it tries to connect to another one. This is a sample X DevAPI URL for configuring client-side failover:

```
mysqlx://sandy:mypassword@[host1:33060,host2:33061]/test
```

An alternate format can also be used, with which the priority for connection can be set explicitly for each individual host:

```
mysqlx://sandy:mypassword@[ (address=host1:33060,priority=2) ,(address=host2:33061,priority=1) ]/test
```

With the client-side failover configured, when there is a failure to establish a connection, Connector/J keeps attempting to connect to a host on the host list in the order of the set priorities for the hosts, which are specified by any numbers between 0 to 100, with a larger number indicating a higher priority for connection. Priorities should either be set for all or no hosts. When no priorities are specified, the priorities for connection are established according to the order the hosts appear in the list, with a host appearing earlier in the list receiving a higher priority.

Notice that this feature only allows for a failover when Connector/J is trying to establish a connection, but not during operations after a connection has already been made.

9.3 Configuring Load Balancing with Connector/J

Connector/J has long provided an effective means to distribute read/write load across multiple MySQL server instances for Cluster or master-master replication deployments. You can dynamically configure load-balanced connections, with no service outage. In-process transactions are not lost, and no application exceptions are generated if any application is trying to use that particular server instance.

The load balancing is configured at the initial setup stage of the server connection by the following connection URL, which has a similar format as [the general JDBC URL for MySQL connection](#), but a specialized scheme:

```
jdbc:mysql:loadbalance://[host1][:port],[host2][:port][,[host3][:port]]...[/[database]] »
[?propertyName1=propertyValue1[&propertyName2=propertyValue2]...]
```

There are two configuration properties associated with this functionality:

- `loadBalanceConnectionGroup` – This provides the ability to group connections from different sources. This allows you to manage these JDBC sources within a single class loader in any combination

you choose. If they use the same configuration, and you want to manage them as a logical single group, give them the same name. This is the key property for management: if you do not define a name (string) for `loadBalanceConnectionGroup`, you cannot manage the connections. All load-balanced connections sharing the same `loadBalanceConnectionGroup` value, regardless of how the application creates them, will be managed together.

- `ha.enableJMX` – The ability to manage the connections is exposed when you define a `loadBalanceConnectionGroup`; but if you want to manage this externally, enable JMX by setting this property to `true`. This enables a JMX implementation, which exposes the management and monitoring operations of a connection group. Further, start your application with the `-Dcom.sun.management.jmxremote` JVM flag. You can then perform connect and perform operations using a JMX client such as `jconsole`.

Once a connection has been made using the correct connection properties, a number of monitoring properties are available:

- Current active host count.
- Current active physical connection count.
- Current active logical connection count.
- Total logical connections created.
- Total transaction count.

The following management operations can also be performed:

- Add host.
- Remove host.

The JMX interface, `com.mysql.cj.jdbc.jmx.LoadBalanceConnectionGroupManagerMBean`, has the following methods:

- `int getActiveHostCount(String group);`
- `int getTotalHostCount(String group);`
- `long getTotalLogicalConnectionCount(String group);`
- `long getActiveLogicalConnectionCount(String group);`
- `long getActivePhysicalConnectionCount(String group);`
- `long getTotalPhysicalConnectionCount(String group);`
- `long getTotalTransactionCount(String group);`
- `void removeHost(String group, String host) throws SQLException;`
- `void stopNewConnectionsToHost(String group, String host) throws SQLException;`
- `void addHost(String group, String host, boolean forExisting);`
- `String getActiveHostsList(String group);`
- `String getRegisteredConnectionGroups();`

The `getRegisteredConnectionGroups()` method returns the names of all connection groups defined in that class loader.

You can test this setup with the following code:

```
public class Test {

    private static String URL = "jdbc:mysql:loadbalance://" +
        "localhost:3306,localhost:3310/test?" +
        "loadBalanceConnectionGroup=first&ha.enableJMX=true";

    public static void main(String[] args) throws Exception {
        new Thread(new Repeater()).start();
        new Thread(new Repeater()).start();
        new Thread(new Repeater()).start();
    }

    static Connection getNewConnection() throws SQLException, ClassNotFoundException {
        Class.forName("com.mysql.cj.jdbc.Driver");
        return DriverManager.getConnection(URL, "root", "");
    }

    static void executeSimpleTransaction(Connection c, int conn, int trans){
        try {
            c.setAutoCommit(false);
            Statement s = c.createStatement();
            s.executeQuery("SELECT SLEEP(1) /* Connection: " + conn + ", transaction: " + trans + " */");
            c.commit();
        } catch (SQLException e) {
            e.printStackTrace();
        }
    }

    public static class Repeater implements Runnable {
        public void run() {
            for(int i=0; i < 100; i++){
                try {
                    Connection c = getNewConnection();
                    for(int j=0; j < 10; j++){
                        executeSimpleTransaction(c, i, j);
                        Thread.sleep(Math.round(100 * Math.random()));
                    }
                    c.close();
                    Thread.sleep(100);
                } catch (Exception e) {
                    e.printStackTrace();
                }
            }
        }
    }
}
```

After compiling, the application can be started with the `-Dcom.sun.management.jmxremote` flag, to enable remote management. `jconsole` can then be started. The `Test` main class will be listed by `jconsole`. Select this and click **Connect**. You can then navigate to the `com.mysql.cj.jdbc.jmx.LoadBalanceConnectionGroupManager` bean. At this point, you can click on various operations and examine the returned result.

If you now had an additional instance of MySQL running on port 3309, you could ensure that Connector/J starts using it by using the `addHost()`, which is exposed in `jconsole`. Note that these operations can be performed dynamically without having to stop the application running.

For further information on the combination of load balancing and failover, see [Section 9.5, “Advanced Load-balancing and Failover Configuration”](#).

9.4 Configuring Master/Slave Replication with Connector/J

This section describe a number of features of Connector/J's support for replication-aware deployments.

The replication is configured at the initial setup stage of the server connection by the connection URL, which has a similar format as [the general JDBC URL for MySQL connection](#), but a specialized scheme:

```
jdbc:mysql:replication://[master host][:port],[slave host 1][:port][,[slave host 2][:port]].../[database]
[?propertyName1=propertyValue1[&propertyName2=propertyValue2]...]
```

Users may specify the property `allowMasterDownConnections=true` to allow `Connection` objects to be created even though no master hosts are reachable. Such `Connection` objects report they are read-only, and `isMasterConnection()` returns false for them. The `Connection` tests for available master hosts when `Connection.setReadOnly(false)` is called, throwing an `SQLException` if it cannot establish a connection to a master, or switching to a master connection if the host is available.

Users may specify the property `allowSlavesDownConnections=true` to allow `Connection` objects to be created even though no slave hosts are reachable. A `Connection` then, at runtime, tests for available slave hosts when `Connection.setReadOnly(true)` is called (see explanation for the method below), throwing an `SQLException` if it cannot establish a connection to a slave, unless the property `readFromMasterWhenNoSlaves` is set to be “true” (see below for a description of the property).

Scaling out Read Load by Distributing Read Traffic to Slaves

Connector/J supports replication-aware connections. It can automatically send queries to a read/write master, or a failover or round-robin loadbalanced set of slaves based on the state of `Connection.getReadOnly()`.

An application signals that it wants a transaction to be read-only by calling `Connection.setReadOnly(true)`. The replication-aware connection will use one of the slave connections, which are load-balanced per slave host using a round-robin scheme. A given connection is sticky to a slave until a transaction boundary command (a commit or rollback) is issued, or until the slave is removed from service. After calling `Connection.setReadOnly(true)`, if you want to allow connection to a master when no slaves are available, set the property `readFromMasterWhenNoSlaves` to “true.” Notice that the master host will be used in read-only state in those cases, as if it is a slave host. Also notice that setting `readFromMasterWhenNoSlaves=true` might result in an extra load for the master host in a transparent manner.

If you have a write transaction, or if you have a read that is time-sensitive (remember, replication in MySQL is asynchronous), set the connection to be not read-only, by calling `Connection.setReadOnly(false)` and the driver will ensure that further calls are sent to the master MySQL server. The driver takes care of propagating the current state of autocommit, isolation level, and catalog between all of the connections that it uses to accomplish this load balancing functionality.

To enable this functionality, use the specialized replication scheme (`jdbc:mysql:replication://`) when connecting to the server.

Here is a short example of how a replication-aware connection might be used in a standalone application:

```
import java.sql.Connection;
import java.sql.ResultSet;
import java.util.Properties;
```

```
import java.sql.DriverManager;

public class ReplicationDemo {

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

        Properties props = new Properties();

        // We want this for failover on the slaves
        props.put("autoReconnect", "true");

        // We want to load balance between the slaves
        props.put("roundRobinLoadBalance", "true");

        props.put("user", "foo");
        props.put("password", "password");

        //
        // Looks like a normal MySQL JDBC url, with a
        // comma-separated list of hosts, the first
        // being the 'master', the rest being any number
        // of slaves that the driver will load balance against
        //

        Connection conn =
            DriverManager.getConnection("jdbc:mysql:replication://master,slave1,slave2,slave3/test",
                props);

        //
        // Perform read/write work on the master
        // by setting the read-only flag to "false"
        //

        conn.setReadOnly(false);
        conn.setAutoCommit(false);
        conn.createStatement().executeUpdate("UPDATE some_table ...");
        conn.commit();

        //
        // Now, do a query from a slave, the driver automatically picks one
        // from the list
        //

        conn.setReadOnly(true);

        ResultSet rs =
            conn.createStatement().executeQuery("SELECT a,b FROM alt_table");

        .....
    }
}
```

Consider using the Load Balancing JDBC Pool ([lbpool](#)) tool, which provides a wrapper around the standard JDBC driver and enables you to use DB connection pools that includes checks for system failures and uneven load distribution. For more information, see [Load Balancing JDBC Driver for MySQL \(mysql-lbpool\)](#).

Support for Multiple-Master Replication Topographies

Connector/J supports multi-master replication topographies.

The connection URL for replication discussed earlier (i.e., in the format of `jdbc:mysql:replication://master,slave1,slave2,slave3/test`) assumes that the first (and only the first) host is the master. Supporting deployments with an arbitrary number of masters and

slaves requires the "address-equals" URL syntax for multiple host connection discussed in [Section 6.2, "Connection URL Syntax"](#), with the property `type=[master|slave]`; for example:

```
jdbc:mysql://address=(type=master)(host=master1host),address=(type=master)(host=master2host),address=(type=
```

Connector/J uses a load-balanced connection internally for management of the master connections, which means that `ReplicationConnection`, when configured to use multiple masters, exposes the same options to balance load across master hosts as described in [Section 9.3, "Configuring Load Balancing with Connector/J"](#).

Live Reconfiguration of Replication Topography

Connector/J also supports live management of replication host (single or multi-master) topographies. This enables users to promote slaves for Java applications without requiring an application restart.

The replication hosts are most effectively managed in the context of a replication connection group. A `ReplicationConnectionGroup` class represents a logical grouping of connections which can be managed together. There may be one or more such replication connection groups in a given Java class loader (there can be an application with two different JDBC resources needing to be managed independently). This key class exposes host management methods for replication connections, and `ReplicationConnection` objects register themselves with the appropriate `ReplicationConnectionGroup` if a value for the new `replicationConnectionGroup` property is specified. The `ReplicationConnectionGroup` object tracks these connections until they are closed, and it is used to manipulate the hosts associated with these connections.

Some important methods related to host management include:

- `getMasterHosts()`: Returns a collection of strings representing the hosts configured as masters
- `getSlaveHosts()`: Returns a collection of strings representing the hosts configured as slaves
- `addSlaveHost(String host)`: Adds new host to pool of possible slave hosts for selection at start of new read-only workload
- `promoteSlaveToMaster(String host)`: Removes the host from the pool of potential slaves for future read-only processes (existing read-only process is allowed to continue to completion) and adds the host to the pool of potential master hosts
- `removeSlaveHost(String host, boolean closeGently)`: Removes the host (host name match must be exact) from the list of configured slaves; if `closeGently` is false, existing connections which have this host as currently active will be closed hardly (application should expect exceptions)
- `removeMasterHost(String host, boolean closeGently)`: Same as `removeSlaveHost()`, but removes the host from the list of configured masters

Some useful management metrics include:

- `getConnectionCountWithHostAsSlave(String host)`: Returns the number of `ReplicationConnection` objects that have the given host configured as a possible slave
- `getConnectionCountWithHostAsMaster(String host)`: Returns the number of `ReplicationConnection` objects that have the given host configured as a possible master
- `getNumberOfSlavesAdded()`: Returns the number of times a slave host has been dynamically added to the group pool
- `getNumberOfSlavesRemoved()`: Returns the number of times a slave host has been dynamically removed from the group pool

- `getNumberOfSlavePromotions()`: Returns the number of times a slave host has been promoted to a master
- `getTotalConnectionCount()`: Returns the number of `ReplicationConnection` objects which have been registered with this group
- `getActiveConnectionCount()`: Returns the number of `ReplicationConnection` objects currently being managed by this group

ReplicationConnectionGroupManager

`com.mysql.cj.jdbc.ha.ReplicationConnectionGroupManager` provides access to the replication connection groups, together with some utility methods.

- `getConnectionGroup(String groupName)`: Returns the `ReplicationConnectionGroup` object matching the `groupName` provided

The other methods in `ReplicationConnectionGroupManager` mirror those of `ReplicationConnectionGroup`, except that the first argument is a `String` group name. These methods will operate on all matching `ReplicationConnectionGroups`, which are helpful for removing a server from service and have it decommissioned across all possible `ReplicationConnectionGroups`.

These methods might be useful for in-JVM management of replication hosts if an application triggers topography changes. For managing host configurations from outside the JVM, JMX can be used.

Using JMX for Managing Replication Hosts

When Connector/J is started with `ha.enableJMX=true` and a value set for the property `replicationConnectionGroup`, a JMX MBean will be registered, allowing manipulation of replication hosts by a JMX client. The MBean interface is defined in `com.mysql.cj.jdbc.jmx.ReplicationGroupManagerMBean`, and leverages the `ReplicationConnectionGroupManager` static methods:

```
public abstract void addSlaveHost(String groupFilter, String host) throws SQLException;
public abstract void removeSlaveHost(String groupFilter, String host) throws SQLException;
public abstract void promoteSlaveToMaster(String groupFilter, String host) throws SQLException;
public abstract void removeMasterHost(String groupFilter, String host) throws SQLException;
public abstract String getMasterHostsList(String group);
public abstract String getSlaveHostsList(String group);
public abstract String getRegisteredConnectionGroups();
public abstract int getActiveMasterHostCount(String group);
public abstract int getActiveSlaveHostCount(String group);
public abstract int getSlavePromotionCount(String group);
public abstract long getTotalLogicalConnectionCount(String group);
public abstract long getActiveLogicalConnectionCount(String group);
```

9.5 Advanced Load-balancing and Failover Configuration

Connector/J provides a useful load-balancing implementation for MySQL Cluster or multi-master deployments, as explained in [Section 9.3, “Configuring Load Balancing with Connector/J”](#) and [Support for Multiple-Master Replication Topographies](#). This same implementation is used for balancing load between read-only slaves for replication-aware connections.

When trying to balance workload between multiple servers, the driver has to determine when it is safe to swap servers, doing so in the middle of a transaction, for example, could cause problems. It is important not to lose state information. For this reason, Connector/J will only try to pick a new server when one of the following happens:

1. At transaction boundaries (transactions are explicitly committed or rolled back).
2. A communication exception (SQL State starting with "08") is encountered.
3. When a `SQLException` matches conditions defined by user, using the extension points defined by the `loadBalanceSQLStateFailover`, `loadBalanceSQLExceptionSubclassFailover` or `loadBalanceExceptionChecker` properties.

The third condition revolves around three properties, which allow you to control which `SQLExceptions` trigger failover:

- `loadBalanceExceptionChecker` - The `loadBalanceExceptionChecker` property is really the key. This takes a fully-qualified class name which implements the new `com.mysql.cj.jdbc.ha.LoadBalanceExceptionChecker` interface. This interface is very simple, and you only need to implement the following method:

```
public boolean shouldExceptionTriggerFailover(SQLException ex)
```

A `SQLException` is passed in, and a boolean returned. A value of `true` triggers a failover, `false` does not.

You can use this to implement your own custom logic. An example where this might be useful is when dealing with transient errors when using MySQL Cluster, where certain buffers may become overloaded. The following code snippet illustrates this:

```
public class NdbLoadBalanceExceptionChecker
    extends StandardLoadBalanceExceptionChecker {

    public boolean shouldExceptionTriggerFailover(SQLException ex) {
        return super.shouldExceptionTriggerFailover(ex)
            || checkNdbException(ex);
    }

    private boolean checkNdbException(SQLException ex){
        // Have to parse the message since most NDB errors
        // are mapped to the same DEMC.
        return (ex.getMessage().startsWith("Lock wait timeout exceeded") ||
            (ex.getMessage().startsWith("Got temporary error")
            && ex.getMessage().endsWith("from NDB")));
    }
}
```

The code above extends `com.mysql.cj.jdbc.ha.StandardLoadBalanceExceptionChecker`, which is the default implementation. There are a few convenient shortcuts built into this, for those who want to have some level of control using properties, without writing Java code. This default implementation uses the two remaining properties: `loadBalanceSQLStateFailover` and `loadBalanceSQLExceptionSubclassFailover`.

- `loadBalanceSQLStateFailover` - allows you to define a comma-delimited list of `SQLState` code prefixes, against which a `SQLException` is compared. If the prefix matches, failover is triggered. So, for example, the following would trigger a failover if a given `SQLException` starts with "00", or is "12345":

```
loadBalanceSQLStateFailover=00,12345
```

- `loadBalanceSQLExceptionSubclassFailover` - can be used in conjunction with `loadBalanceSQLStateFailover` or on its own. If you want certain subclasses of `SQLException` to

trigger failover, simply provide a comma-delimited list of fully-qualified class or interface names to check against. For example, if you want all `SQLTransientConnectionExceptions` to trigger failover, you would specify:

```
loadBalanceSQLExceptionSubclassFailover=java.sql.SQLTransientConnectionException
```

While the three failover conditions enumerated earlier suit most situations, if `autocommit` is enabled, Connector/J never re-balances, and continues using the same physical connection. This can be problematic, particularly when load-balancing is being used to distribute read-only load across multiple slaves. However, Connector/J can be configured to re-balance after a certain number of statements are executed, when `autocommit` is enabled. This functionality is dependent upon the following properties:

- `loadBalanceAutoCommitStatementThreshold` – defines the number of matching statements which will trigger the driver to potentially swap physical server connections. The default value, 0, retains the behavior that connections with `autocommit` enabled are never balanced.
- `loadBalanceAutoCommitStatementRegex` – the regular expression against which statements must match. The default value, blank, matches all statements. So, for example, using the following properties will cause Connector/J to re-balance after every third statement that contains the string “test”:

```
loadBalanceAutoCommitStatementThreshold=3  
loadBalanceAutoCommitStatementRegex=.*test.*
```

`loadBalanceAutoCommitStatementRegex` can prove useful in a number of situations. Your application may use temporary tables, server-side session state variables, or connection state, where letting the driver arbitrarily swap physical connections before processing is complete could cause data loss or other problems. This allows you to identify a trigger statement that is only executed when it is safe to swap physical connections.

Chapter 10 Using the Connector/J Interceptor Classes

An interceptor is a software design pattern that provides a transparent way to extend or modify some aspect of a program, similar to a user exit. No recompiling is required. With Connector/J, the interceptors are enabled and disabled by updating the connection string to refer to different sets of interceptor classes that you instantiate.

The connection properties that control the interceptors are explained in [Section 6.3, "Configuration Properties"](#):

- `connectionLifecycleInterceptors`, where you specify the fully qualified names of classes that implement the `com.mysql.cj.jdbc.interceptors.ConnectionLifecycleInterceptor` interface. In these kinds of interceptor classes, you might log events such as rollbacks, measure the time between transaction start and end, or count events such as calls to `setAutoCommit()`.
- `exceptionInterceptors`, where you specify the fully qualified names of classes that implement the `com.mysql.cj.exceptions.ExceptionInterceptor` interface. In these kinds of interceptor classes, you might add extra diagnostic information to exceptions that can have multiple causes or indicate a problem with server settings. `exceptionInterceptors` classes are called when handling an `Exception` thrown from Connector/J code.
- `queryInterceptors`, where you specify the fully qualified names of classes that implement the `com.mysql.cj.interceptors.QueryInterceptor` interface. In these kinds of interceptor classes, you might change or augment the processing done by certain kinds of statements, such as automatically checking for queried data in a `memcached` server, rewriting slow queries, logging information about statement execution, or route requests to remote servers.

Chapter 11 Using Connector/J with Tomcat

The following instructions are based on the instructions for Tomcat-5.x, available at <http://tomcat.apache.org/tomcat-5.5-doc/jndi-datasource-examples-howto.html> which is current at the time this document was written.

First, install the `.jar` file that comes with Connector/J in `$CATALINA_HOME/common/lib` so that it is available to all applications installed in the container.

Next, configure the JNDI DataSource by adding a declaration resource to `$CATALINA_HOME/conf/server.xml` in the context that defines your web application:

```
<Context ....>
...
<Resource name="jdbc/MySQLDB"
          auth="Container"
          type="javax.sql.DataSource" />
<ResourceParams name="jdbc/MySQLDB">
  <parameter>
    <name>factory</name>
    <value>org.apache.commons.dbcp.BasicDataSourceFactory</value>
  </parameter>
  <parameter>
    <name>maxActive</name>
    <value>10</value>
  </parameter>
  <parameter>
    <name>maxIdle</name>
    <value>5</value>
  </parameter>
  <parameter>
    <name>validationQuery</name>
    <value>SELECT 1</value>
  </parameter>
  <parameter>
    <name>testOnBorrow</name>
    <value>true</value>
  </parameter>
  <parameter>
    <name>testWhileIdle</name>
    <value>true</value>
  </parameter>
  <parameter>
    <name>timeBetweenEvictionRunsMillis</name>
    <value>10000</value>
  </parameter>
  <parameter>
    <name>minEvictableIdleTimeMillis</name>
    <value>60000</value>
  </parameter>
  <parameter>
    <name>username</name>
```

```

    <value>someuser</value>
  </parameter>

  <parameter>
    <name>password</name>
    <value>somepass</value>
  </parameter>

  <parameter>
    <name>driverClassName</name>
    <value>com.mysql.cj.jdbc.Driver</value>
  </parameter>

  <parameter>
    <name>url</name>
    <value>jdbc:mysql://localhost:3306/test</value>
  </parameter>

</ResourceParams>
</Context>

```

Connector/J introduces a facility whereby, rather than use a `validationQuery` value of `SELECT 1`, it is possible to use `validationQuery` with a value set to `/* ping */`. This sends a ping to the server which then returns a fake result set. This is a lighter weight solution. It also has the advantage that if using `ReplicationConnection` or `LoadBalancedConnection` type connections, the ping will be sent across all active connections. The following XML snippet illustrates how to select this option:

```

<parameter>
  <name>validationQuery</name>
  <value>/* ping */</value>
</parameter>

```

Note that `/* ping */` has to be specified exactly.

In general, follow the installation instructions that come with your version of Tomcat, as the way you configure datasources in Tomcat changes from time to time, and if you use the wrong syntax in your XML file, you will most likely end up with an exception similar to the following:

```

Error: java.sql.SQLException: Cannot load JDBC driver class 'null ' SQL
state: null

```

Note that the auto-loading of drivers having the `META-INF/service/java.sql.Driver` class in JDBC 4.0 and later causes an improper undeployment of the Connector/J driver in Tomcat on Windows. Namely, the Connector/J jar remains locked. This is an initialization problem that is not related to the driver. The possible workarounds, if viable, are as follows: use `antiResourceLocking=true` as a Tomcat Context attribute, or remove the `META-INF/` directory.

Chapter 12 Using Connector/J with JBoss

These instructions cover JBoss-4.x. To make the JDBC driver classes available to the application server, copy the `.jar` file that comes with Connector/J to the `lib` directory for your server configuration (which is usually called `default`). Then, in the same configuration directory, in the subdirectory named `deploy`, create a datasource configuration file that ends with `-ds.xml`, which tells JBoss to deploy this file as a JDBC Datasource. The file should have the following contents:

```
<datasources>
  <local-tx-datasource>

    <jndi-name>MySQLDB</jndi-name>
    <connection-url>jdbc:mysql://localhost:3306/dbname</connection-url>
    <driver-class>com.mysql.cj.jdbc.Driver</driver-class>
    <user-name>user</user-name>
    <password>pass</password>

    <min-pool-size>5</min-pool-size>

    <max-pool-size>20</max-pool-size>

    <idle-timeout-minutes>5</idle-timeout-minutes>

    <exception-sorter-class-name>
com.mysql.cj.jdbc.integration.jboss.ExtendedMySQLExceptionSorter
    </exception-sorter-class-name>
    <valid-connection-checker-class-name>
com.mysql.cj.jdbc.integration.jboss.MySqlValidConnectionChecker
    </valid-connection-checker-class-name>

  </local-tx-datasource>
</datasources>
```

Chapter 13 Using Connector/J with Spring

Table of Contents

13.1 Using <code>JdbcTemplate</code>	106
13.2 Transactional JDBC Access	108
13.3 Connection Pooling with Spring	109

The Spring Framework is a Java-based application framework designed for assisting in application design by providing a way to configure components. The technique used by Spring is a well known design pattern called Dependency Injection (see [Inversion of Control Containers and the Dependency Injection pattern](#)). This article will focus on Java-oriented access to MySQL databases with Spring 2.0. For those wondering, there is a .NET port of Spring appropriately named Spring.NET.

Spring is not only a system for configuring components, but also includes support for aspect oriented programming (AOP). This is one of the main benefits and the foundation for Spring's resource and transaction management. Spring also provides utilities for integrating resource management with JDBC and Hibernate.

For the examples in this section the MySQL world sample database will be used. The first task is to set up a MySQL data source through Spring. Components within Spring use the "bean" terminology. For example, to configure a connection to a MySQL server supporting the world sample database, you might use:

```
<util:map id="dbProps">
  <entry key="db.driver" value="com.mysql.cj.jdbc.Driver"/>
  <entry key="db.jdbcurl" value="jdbc:mysql://localhost/world"/>
  <entry key="db.username" value="myuser"/>
  <entry key="db.password" value="mypass"/>
</util:map>
```

In the above example, we are assigning values to properties that will be used in the configuration. For the datasource configuration:

```
<bean id="dataSource"
  class="org.springframework.jdbc.datasource.DriverManagerDataSource">
  <property name="driverClassName" value="${db.driver}"/>
  <property name="url" value="${db.jdbcurl}"/>
  <property name="username" value="${db.username}"/>
  <property name="password" value="${db.password}"/>
</bean>
```

The placeholders are used to provide values for properties of this bean. This means that we can specify all the properties of the configuration in one place instead of entering the values for each property on each bean. We do, however, need one more bean to pull this all together. The last bean is responsible for actually replacing the placeholders with the property values.

```
<bean
  class="org.springframework.beans.factory.config.PropertyPlaceholderConfigurer">
```

```
<property name="properties" ref="dbProps"/>
</bean>
```

Now that we have our MySQL data source configured and ready to go, we write some Java code to access it. The example below will retrieve three random cities and their corresponding country using the data source we configured with Spring.

```
// Create a new application context. this processes the Spring config
ApplicationContext ctx =
    new ClassPathXmlApplicationContext("exlappContext.xml");
// Retrieve the data source from the application context
DataSource ds = (DataSource) ctx.getBean("dataSource");
// Open a database connection using Spring's DataSourceUtils
Connection c = DataSourceUtils.getConnection(ds);
try {
    // retrieve a list of three random cities
    PreparedStatement ps = c.prepareStatement(
        "select City.Name as 'City', Country.Name as 'Country' " +
        "from City inner join Country on City.CountryCode = Country.Code " +
        "order by rand() limit 3");
    ResultSet rs = ps.executeQuery();
    while(rs.next()) {
        String city = rs.getString("City");
        String country = rs.getString("Country");
        System.out.printf("The city %s is in %s\n", city, country);
    }
} catch (SQLException ex) {
    // something has failed and we print a stack trace to analyse the error
    ex.printStackTrace();
    // ignore failure closing connection
    try { c.close(); } catch (SQLException e) { }
} finally {
    // properly release our connection
    DataSourceUtils.releaseConnection(c, ds);
}
```

This is very similar to normal JDBC access to MySQL with the main difference being that we are using `DataSourceUtils` instead of the `DriverManager` to create the connection.

While it may seem like a small difference, the implications are somewhat far reaching. Spring manages this resource in a way similar to a container managed data source in a J2EE application server. When a connection is opened, it can be subsequently accessed in other parts of the code if it is synchronized with a transaction. This makes it possible to treat different parts of your application as transactional instead of passing around a database connection.

13.1 Using JdbcTemplate

Spring makes extensive use of the Template method design pattern (see [Template Method Pattern](#)). Our immediate focus will be on the `JdbcTemplate` and related classes, specifically `NamedParameterJdbcTemplate`. The template classes handle obtaining and releasing a connection for data access when one is needed.

The next example shows how to use `NamedParameterJdbcTemplate` inside of a DAO (Data Access Object) class to retrieve a random city given a country code.

```
public class Ex2JdbcDao {
    /**
     * Data source reference which will be provided by Spring.
     */
}
```



```

private DataSource dataSource;

/**
 * Our query to find a random city given a country code. Notice
 * the ":country" parameter toward the end. This is called a
 * named parameter.
 */
private String queryString = "select Name from City " +
    "where CountryCode = :country order by rand() limit 1";

/**
 * Retrieve a random city using Spring JDBC access classes.
 */
public String getRandomCityByCountryCode(String cuntryCode) {
    // A template that permits using queries with named parameters
    NamedParameterJdbcTemplate template =
        new NamedParameterJdbcTemplate(dataSource);
    // A java.util.Map is used to provide values for the parameters
    Map params = new HashMap();
    params.put("country", cuntryCode);
    // We query for an Object and specify what class we are expecting
    return (String)template.queryForObject(queryString, params, String.class);
}

/**
 * A JavaBean setter-style method to allow Spring to inject the data source.
 * @param dataSource
 */
public void setDataSource(DataSource dataSource) {
    this.dataSource = dataSource;
}
}

```

The focus in the above code is on the `getRandomCityByCountryCode()` method. We pass a country code and use the `NamedParameterJdbcTemplate` to query for a city. The country code is placed in a Map with the key "country", which is the parameter is named in the SQL query.

To access this code, you need to configure it with Spring by providing a reference to the data source.

```

<bean id="dao" class="code.Ex2JdbcDao">
  <property name="dataSource" ref="dataSource"/>
</bean>

```

At this point, we can just grab a reference to the DAO from Spring and call `getRandomCityByCountryCode()`.

```

// Create the application context
ApplicationContext ctx =
    new ClassPathXmlApplicationContext("ex2appContext.xml");
// Obtain a reference to our DAO
Ex2JdbcDao dao = (Ex2JdbcDao) ctx.getBean("dao");

String countryCode = "USA";

// Find a few random cities in the US
for(int i = 0; i < 4; ++i)
    System.out.printf("A random city in %s is %s\n", countryCode,
        dao.getRandomCityByCountryCode(countryCode));

```

This example shows how to use Spring's JDBC classes to completely abstract away the use of traditional JDBC classes including `Connection` and `PreparedStatement`.

13.2 Transactional JDBC Access

Spring allows us to add transactions into our code without having to deal directly with the JDBC classes. For that purpose, Spring provides a transaction management package that not only replaces JDBC transaction management, but also enables declarative transaction management (configuration instead of code).

To use transactional database access, we will need to change the storage engine of the tables in the world database. The downloaded script explicitly creates MyISAM tables, which do not support transactional semantics. The InnoDB storage engine does support transactions and this is what we will be using. We can change the storage engine with the following statements.

```
ALTER TABLE City ENGINE=InnoDB;
ALTER TABLE Country ENGINE=InnoDB;
ALTER TABLE CountryLanguage ENGINE=InnoDB;
```

A good programming practice emphasized by Spring is separating interfaces and implementations. What this means is that we can create a Java interface and only use the operations on this interface without any internal knowledge of what the actual implementation is. We will let Spring manage the implementation and with this it will manage the transactions for our implementation.

First you create a simple interface:

```
public interface Ex3Dao {
    Integer createCity(String name, String countryCode,
        String district, Integer population);
}
```

This interface contains one method that will create a new city record in the database and return the id of the new record. Next you need to create an implementation of this interface.

```
public class Ex3DaoImpl implements Ex3Dao {
    protected DataSource dataSource;
    protected SqlUpdate updateQuery;
    protected SqlFunction idQuery;

    public Integer createCity(String name, String countryCode,
        String district, Integer population) {
        updateQuery.update(new Object[] { name, countryCode,
            district, population });
        return getLastId();
    }

    protected Integer getLastId() {
        return idQuery.run();
    }
}
```

You can see that we only operate on abstract query objects here and do not deal directly with the JDBC API. Also, this is the complete implementation. All of our transaction management will be dealt with in the configuration. To get the configuration started, we need to create the DAO.

```
<bean id="dao" class="code.Ex3DaoImpl">
    <property name="dataSource" ref="dataSource"/>
    <property name="updateQuery">...</property>
    <property name="idQuery">...</property>
</bean>
```

Now we need to set up the transaction configuration. The first thing we must do is create transaction manager to manage the data source and a specification of what transaction properties are required for the `dao` methods.

```
<bean id="transactionManager"
  class="org.springframework.jdbc.datasource.DataSourceTransactionManager">
  <property name="dataSource" ref="dataSource"/>
</bean>

<tx:advice id="txAdvice" transaction-manager="transactionManager">
  <tx:attributes>
    <tx:method name="*" />
  </tx:attributes>
</tx:advice>
```

The preceding code creates a transaction manager that handles transactions for the data source provided to it. The `txAdvice` uses this transaction manager and the attributes specify to create a transaction for all methods. Finally we need to apply this advice with an AOP pointcut.

```
<aop:config>
  <aop:pointcut id="daoMethods"
    expression="execution(* code.Ex3Dao.*(..))"/>
  <aop:advisor advice-ref="txAdvice" pointcut-ref="daoMethods"/>
</aop:config>
```

This basically says that all methods called on the `Ex3Dao` interface will be wrapped in a transaction. To make use of this, we only have to retrieve the `dao` from the application context and call a method on the `dao` instance.

```
Ex3Dao dao = (Ex3Dao) ctx.getBean("dao");
Integer id = dao.createCity(name, countryCode, district, pop);
```

We can verify from this that there is no transaction management happening in our Java code and it is all configured with Spring. This is a very powerful notion and regarded as one of the most beneficial features of Spring.

13.3 Connection Pooling with Spring

In many situations, such as web applications, there will be a large number of small database transactions. When this is the case, it usually makes sense to create a pool of database connections available for web requests as needed. Although MySQL does not spawn an extra process when a connection is made, there is still a small amount of overhead to create and set up the connection. Pooling of connections also alleviates problems such as collecting large amounts of sockets in the `TIME_WAIT` state.

Setting up pooling of MySQL connections with Spring is as simple as changing the data source configuration in the application context. There are a number of configurations that we can use. The first example is based on the [Jakarta Commons DBCP library](#). The example below replaces the source configuration that was based on `DriverManagerDataSource` with DBCP's `BasicDataSource`.

```
<bean id="dataSource" destroy-method="close"
```

```
class="org.apache.commons.dbcp.BasicDataSource">
  <property name="driverClassName" value="{db.driver}"/>
  <property name="url" value="{db.jdbcurl}"/>
  <property name="username" value="{db.username}"/>
  <property name="password" value="{db.password}"/>
  <property name="initialSize" value="3"/>
</bean>
```

The configuration of the two solutions is very similar. The difference is that DBCP will pool connections to the database instead of creating a new connection every time one is requested. We have also set a parameter here called `initialSize`. This tells DBCP that we want three connections in the pool when it is created.

Another way to configure connection pooling is to configure a data source in our J2EE application server. Using JBoss as an example, you can set up the MySQL connection pool by creating a file called `mysql-local-ds.xml` and placing it in the `server/default/deploy` directory in JBoss. Once we have this setup, we can use JNDI to look it up. With Spring, this lookup is very simple. The data source configuration looks like this.

```
<jee:jndi-lookup id="dataSource" jndi-name="java:MySQL_DS"/>
```

Chapter 14 Troubleshooting Connector/J Applications

This section explains the symptoms and resolutions for the most commonly encountered issues with applications using MySQL Connector/J.

Questions

- **14.1:** When I try to connect to the database with MySQL Connector/J, I get the following exception:

```
SQLException: Server configuration denies access to data source
SQLState: 08001
VendorError: 0
```

What is going on? I can connect just fine with the MySQL command-line client.

- **14.2:** My application throws an SQLException 'No Suitable Driver'. Why is this happening?
- **14.3:** I'm trying to use MySQL Connector/J in an applet or application and I get an exception similar to:

```
SQLException: Cannot connect to MySQL server on host:3306.
Is there a MySQL server running on the machine/port you
are trying to connect to?

(java.security.AccessControlException)
SQLState: 08S01
VendorError: 0
```

- **14.4:** I have a servlet/application that works fine for a day, and then stops working overnight
- **14.5:** I cannot connect to the MySQL server using Connector/J, and I'm sure the connection parameters are correct.
- **14.6:** My application is deployed through JBoss and I am using transactions to handle the statements on the MySQL database. Under heavy loads, I am getting an error and stack trace, but these only occur after a fixed period of heavy activity.
- **14.7:** Updating a table that contains a [primary key](#) that is either [FLOAT](#) or compound primary key that uses [FLOAT](#) fails to update the table and raises an exception.
- **14.8:** I get an [ER_NET_PACKET_TOO_LARGE](#) exception, even though the binary blob size I want to insert using JDBC is safely below the [max_allowed_packet](#) size.
- **14.9:** What should I do if I receive error messages similar to the following: "Communications link failure – Last packet sent to the server was X ms ago"?
- **14.10:** Why does Connector/J not reconnect to MySQL and re-issue the statement after a communication failure instead of throwing an Exception, even though I use the [autoReconnect](#) connection string option?
- **14.11:** How can I use 3-byte UTF8 with Connector/J?
- **14.12:** How can I use 4-byte UTF8, [utf8mb4](#) with Connector/J?
- **14.13:** Using [useServerPrepStmts=false](#) and certain character encodings can lead to corruption when inserting BLOBs. How can this be avoided?

Questions and Answers

- **14.1:** When I try to connect to the database with MySQL Connector/J, I get the following exception:

```
SQLException: Server configuration denies access to data source
SQLState: 08001
VendorError: 0
```

What is going on? I can connect just fine with the MySQL command-line client.

Connector/J normally uses TCP/IP sockets to connect to MySQL (see [Section 6.8, “Connecting Using Unix Domain Sockets”](#) and [Section 6.9, “Connecting Using Named Pipes”](#) for exceptions). The security manager on the MySQL server uses its grant tables to determine whether a TCP/IP connection is permitted. You must therefore add the necessary security credentials to the MySQL server for the connection by issuing a [GRANT](#) statement to your MySQL Server. See [GRANT Syntax](#), for more information.

Warning

Changing privileges and permissions improperly on MySQL can potentially cause your server installation to have non-optimal security properties.

Note

Testing your connectivity with the `mysql` command-line client will not work unless you add the `--host` flag, and use something other than `localhost` for the host. The `mysql` command-line client will try to use Unix domain sockets if you use the special host name `localhost`. If you are testing TCP/IP connectivity to `localhost`, use `127.0.0.1` as the host name instead.

14.2: My application throws an SQLException 'No Suitable Driver'. Why is this happening?

There are three possible causes for this error:

- The Connector/J driver is not in your `CLASSPATH`, see [Chapter 4, Connector/J Installation](#).
- The format of your connection URL is incorrect, or you are referencing the wrong JDBC driver.
- When using `DriverManager`, the `jdbc.drivers` system property has not been populated with the location of the Connector/J driver.

14.3: I'm trying to use MySQL Connector/J in an applet or application and I get an exception similar to:

```
SQLException: Cannot connect to MySQL server on host:3306.
Is there a MySQL server running on the machine/port you
are trying to connect to?

(java.security.AccessControlException)
SQLState: 08S01
VendorError: 0
```

Either you're running an Applet, your MySQL server has been installed with the "skip-networking" option set, or your MySQL server has a firewall sitting in front of it.

Applets can only make network connections back to the machine that runs the web server that served the `.class` files for the applet. This means that MySQL must run on the same machine (or you must have some sort of port re-direction) for this to work. This also means that you will not be able to test applets from your local file system, but must always deploy them to a web server.

Connector/J normally uses TCP/IP sockets to connect to MySQL (see [Section 6.8, “Connecting Using Unix Domain Sockets”](#) and [Section 6.9, “Connecting Using Named Pipes”](#) for exceptions). TCP/IP communication with MySQL can be affected by the server option `--skip-networking` or the server firewall. If MySQL has been started with the `--skip-networking` option, you need to comment it out

in the file `/etc/mysql/my.cnf` or `/etc/my.cnf` for TCP/IP connections to work. (Note that your server configuration file might also exist in the `data` directory of your MySQL server, or somewhere else, depending on how MySQL was compiled; binaries created by Oracle always look for `/etc/my.cnf` and `datadir/my.cnf`; see [Using Option Files](#) for details.) If your MySQL server has been firewalled, you will need to have the firewall configured to allow TCP/IP connections from the host where your Java code is running to the MySQL server on the port that MySQL is listening to (by default, 3306).

14.4: I have a servlet/application that works fine for a day, and then stops working overnight

MySQL closes connections after 8 hours of inactivity. You either need to use a connection pool that handles stale connections or use the `autoReconnect` parameter (see [Section 6.3, "Configuration Properties"](#)).

Also, catch `SQLExceptions` in your application and deal with them, rather than propagating them all the way until your application exits. This is just good programming practice. MySQL Connector/J will set the `SQLState` (see `java.sql.SQLException.getSQLState()` in your API docs) to `08S01` when it encounters network-connectivity issues during the processing of a query. Attempt to reconnect to MySQL at this point.

The following (simplistic) example shows what code that can handle these exceptions might look like:

Example 14.1 Connector/J: Example of transaction with retry logic

```
public void doBusinessOp() throws SQLException {
    Connection conn = null;
    Statement stmt = null;
    ResultSet rs = null;

    //
    // How many times do you want to retry the transaction
    // (or at least _getting_ a connection)?
    //
    int retryCount = 5;

    boolean transactionCompleted = false;

    do {
        try {
            conn = getConnection(); // assume getting this from a
                                   // javax.sql.DataSource, or the
                                   // java.sql.DriverManager

            conn.setAutoCommit(false);

            //
            // Okay, at this point, the 'retry-ability' of the
            // transaction really depends on your application logic,
            // whether or not you're using autocommit (in this case
            // not), and whether you're using transactional storage
            // engines
            //
            // For this example, we'll assume that it's _not_ safe
            // to retry the entire transaction, so we set retry
            // count to 0 at this point
            //
            // If you were using exclusively transaction-safe tables,
            // or your application could recover from a connection going
            // bad in the middle of an operation, then you would not
            // touch 'retryCount' here, and just let the loop repeat
            // until retryCount == 0.
            //
            retryCount = 0;
        }
    }
}
```

```

    stmt = conn.createStatement();

    String query = "SELECT foo FROM bar ORDER BY baz";

    rs = stmt.executeQuery(query);

    while (rs.next()) {
    }

    rs.close();
    rs = null;

    stmt.close();
    stmt = null;

    conn.commit();
    conn.close();
    conn = null;

    transactionCompleted = true;
} catch (SQLException sqlEx) {

    //
    // The two SQL states that are 'retry-able' are 08S01
    // for a communications error, and 40001 for deadlock.
    //
    // Only retry if the error was due to a stale connection,
    // communications problem or deadlock
    //

    String sqlState = sqlEx.getSQLState();

    if ("08S01".equals(sqlState) || "40001".equals(sqlState)) {
        retryCount -= 1;
    } else {
        retryCount = 0;
    }
} finally {
    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException sqlEx) {
            // You'd probably want to log this...
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlEx) {
            // You'd probably want to log this as well...
        }
    }

    if (conn != null) {
        try {
            //
            // If we got here, and conn is not null, the
            // transaction should be rolled back, as not
            // all work has been done

            try {
                conn.rollback();
            } finally {
                conn.close();
            }
        } catch (SQLException sqlEx) {

```



```

//
// If we got an exception here, something
// pretty serious is going on, so we better
// pass it up the stack, rather than just
// logging it...

        throw sqlEx;
    }
}
} while (!transactionCompleted && (retryCount > 0));
}

```

Note

Use of the `autoReconnect` option is not recommended because there is no safe method of reconnecting to the MySQL server without risking some corruption of the connection state or database state information. Instead, use a connection pool, which will enable your application to connect to the MySQL server using an available connection from the pool. The `autoReconnect` facility is deprecated, and may be removed in a future release.

14.5: I cannot connect to the MySQL server using Connector/J, and I'm sure the connection parameters are correct.

Make sure that the `skip-networking` option has not been enabled on your server. Connector/J must be able to communicate with your server over TCP/IP; named sockets are not supported. Also ensure that you are not filtering connections through a firewall or other network security system. For more information, see [Can't connect to \[local\] MySQL server](#).

14.6: My application is deployed through JBoss and I am using transactions to handle the statements on the MySQL database. Under heavy loads, I am getting an error and stack trace, but these only occur after a fixed period of heavy activity.

This is a JBoss, not Connector/J, issue and is connected to the use of transactions. Under heavy loads the time taken for transactions to complete can increase, and the error is caused because you have exceeded the predefined timeout.

You can increase the timeout value by setting the `TransactionTimeout` attribute to the `TransactionManagerService` within the `/conf/jboss-service.xml` file (pre-4.0.3) or `/deploy/jta-service.xml` for JBoss 4.0.3 or later. See `TransactionTimeout` within the JBoss wiki for more information.

14.7: Updating a table that contains a primary key that is either `FLOAT` or compound primary key that uses `FLOAT` fails to update the table and raises an exception.

Connector/J adds conditions to the `WHERE` clause during an `UPDATE` to check the old values of the primary key. If there is no match, then Connector/J considers this a failure condition and raises an exception.

The problem is that rounding differences between supplied values and the values stored in the database may mean that the values never match, and hence the update fails. The issue will affect all queries, not just those from Connector/J.

To prevent this issue, use a primary key that does not use `FLOAT`. If you have to use a floating point column in your primary key, use `DOUBLE` or `DECIMAL` types in place of `FLOAT`.

14.8: I get an `ER_NET_PACKET_TOO_LARGE` exception, even though the binary blob size I want to insert using JDBC is safely below the `max_allowed_packet` size.

This is because the `hexEscapeBlock()` method in `com.mysql.cj.AbstractPreparedQuery.streamToBytes()` may almost double the size of your data.

14.9: What should I do if I receive error messages similar to the following: “Communications link failure – Last packet sent to the server was X ms ago”?

Generally speaking, this error suggests that the network connection has been closed. There can be several root causes:

- Firewalls or routers may clamp down on idle connections (the MySQL client/server protocol does not ping).
- The MySQL Server may be closing idle connections that exceed the `wait_timeout` or `interactive_timeout` threshold.

Although network connections can be volatile, the following can be helpful in avoiding problems:

- Ensure connections are valid when used from the connection pool. Use a query that starts with `/* ping */` to execute a lightweight ping instead of full query. Note, the syntax of the ping needs to be exactly as specified here.
- Minimize the duration a connection object is left idle while other application logic is executed.
- Explicitly validate the connection before using it if the connection has been left idle for an extended period of time.
- Ensure that `wait_timeout` and `interactive_timeout` are set sufficiently high.
- Ensure that `tcpKeepalive` is enabled.
- Ensure that any configurable firewall or router timeout settings allow for the maximum expected connection idle time.

Note

Do not expect to be able to reuse a connection without problems if it has been lying idle for a period. If a connection is to be reused after being idle for any length of time, ensure that you explicitly test it before reusing it.

14.10: Why does Connector/J not reconnect to MySQL and re-issue the statement after a communication failure instead of throwing an Exception, even though I use the `autoReconnect` connection string option?

There are several reasons for this. The first is transactional integrity. The MySQL Reference Manual states that “there is no safe method of reconnecting to the MySQL server without risking some corruption of the connection state or database state information”. Consider the following series of statements for example:

```
conn.createStatement().execute(
    "UPDATE checking_account SET balance = balance - 1000.00 WHERE customer='Smith'");
conn.createStatement().execute(
    "UPDATE savings_account SET balance = balance + 1000.00 WHERE customer='Smith'");
conn.commit();
```

Consider the case where the connection to the server fails after the `UPDATE` to `checking_account`. If no exception is thrown, and the application never learns about the problem, it will continue executing.

However, the server did not commit the first transaction in this case, so that will get rolled back. But execution continues with the next transaction, and increases the `savings_account` balance by 1000. The application did not receive an exception, so it continued regardless, eventually committing the second transaction, as the commit only applies to the changes made in the new connection. Rather than a transfer taking place, a deposit was made in this example.

Note that running with `autocommit` enabled does not solve this problem. When Connector/J encounters a communication problem, there is no means to determine whether the server processed the currently executing statement or not. The following theoretical states are equally possible:

- The server never received the statement, and therefore no related processing occurred on the server.
- The server received the statement, executed it in full, but the response was not received by the client.

If you are running with `autocommit` enabled, it is not possible to guarantee the state of data on the server when a communication exception is encountered. The statement may have reached the server, or it may not. All you know is that communication failed at some point, before the client received confirmation (or data) from the server. This does not only affect `autocommit` statements though. If the communication problem occurred during `Connection.commit()`, the question arises of whether the transaction was committed on the server before the communication failed, or whether the server received the commit request at all.

The second reason for the generation of exceptions is that transaction-scoped contextual data may be vulnerable, for example:

- Temporary tables.
- User-defined variables.
- Server-side prepared statements.

These items are lost when a connection fails, and if the connection silently reconnects without generating an exception, this could be detrimental to the correct execution of your application.

In summary, communication errors generate conditions that may well be unsafe for Connector/J to simply ignore by silently reconnecting. It is necessary for the application to be notified. It is then for the application developer to decide how to proceed in the event of connection errors and failures.

14.11: How can I use 3-byte UTF8 with Connector/J?

To use 3-byte UTF8 with Connector/J set `characterEncoding=utf8` and set `useUnicode=true` in the connection string.

14.12: How can I use 4-byte UTF8, `utf8mb4` with Connector/J?

To use 4-byte UTF8 with Connector/J configure the MySQL server with `character_set_server=utf8mb4`. Connector/J will then use that setting as long as `characterEncoding` has not been set in the connection string. This is equivalent to autodetection of the character set.

14.13: Using `useServerPrepStmts=false` and certain character encodings can lead to corruption when inserting BLOBs. How can this be avoided?

When using certain character encodings, such as SJIS, CP932, and BIG5, it is possible that BLOB data contains characters that can be interpreted as control characters, for example, backslash, `\`. This can lead to corrupted data when inserting BLOBs into the database. There are two things that need to be done to avoid this:

-
1. Set the connection string option `useServerPrepStmts` to `true`.
 2. Set `SQL_MODE` to `NO_BACKSLASH_ESCAPES`.

Chapter 15 Known Issues and Limitations

The following are some known issues and limitations for MySQL Connector/J:

- When Connector/J retrieves timestamps for a daylight saving time (DST) switch day using the `getTimestamp()` method on the result set, some of the returned values might be wrong. The errors can be avoided by using the following connection options when connecting to a database:

```
serverTimezone=UTC
```

- The functionality of the property `elideSetAutoCommits` has been disabled due to Bug# 66884. Any value given for the property is ignored by Connector/J.

Chapter 16 Connector/J Support

Table of Contents

16.1 Connector/J Community Support	121
16.2 How to Report Connector/J Bugs or Problems	121

16.1 Connector/J Community Support

Oracle provides assistance to the user community by means of its mailing lists. For Connector/J related issues, you can get help from experienced users by using the MySQL and Java mailing list. Archives and subscription information is available online at <http://lists.mysql.com/java>.

For information about subscribing to MySQL mailing lists or to browse list archives, visit <http://lists.mysql.com/>. See [MySQL Mailing Lists](#).

Community support from experienced users is also available through the [JDBC Forum](#). You may also find help from other users in the other MySQL Forums, located at <http://forums.mysql.com>. See [MySQL Community Support at the MySQL Forums](#).

You can also join the [#connectors](#) channel in the [MySQL Community Slack workspace](#), where you can get help directly from MySQL developers and other users.

16.2 How to Report Connector/J Bugs or Problems

The normal place to report bugs is <http://bugs.mysql.com/>, which is the address for our bugs database. This database is public, and can be browsed and searched by anyone. If you log in to the system, you will also be able to enter new reports.

If you find a sensitive security bug in MySQL Server, please let us know immediately by sending an email message to [<secalert_us@oracle.com>](mailto:secalert_us@oracle.com). Exception: Support customers should report all problems, including security bugs, to Oracle Support at <http://support.oracle.com/>.

Writing a good bug report takes patience, but doing it right the first time saves time both for us and for yourself. A good bug report, containing a full test case for the bug, makes it very likely that we will fix sooner rather than later.

This section will help you write your report correctly so that you do not waste your time doing things that may not help us much or at all.

If you have a repeatable bug report, please report it to the bugs database at <http://bugs.mysql.com/>. Any bug that we are able to repeat has a high chance of being fixed sooner rather than later.

To report other problems, you can use one of the MySQL mailing lists.

Remember that it is possible for us to respond to a message containing too much information, but not to one containing too little. People often omit facts because they think they know the cause of a problem and assume that some details do not matter.

A good principle is this: If you are in doubt about stating something, state it. It is faster and less troublesome to write a couple more lines in your report than to wait longer for the answer if we must ask you to provide information that was missing from the initial report.

The most common errors made in bug reports are (a) not including the version number of Connector/J or MySQL used, and (b) not fully describing the platform on which Connector/J is installed (including the JVM version, and the platform type and version number that MySQL itself is installed on).

This is highly relevant information, and in 99 cases out of 100, the bug report is useless without it. Very often we get questions like, “Why doesn't this work for me?” Then we find that the feature requested was not implemented in that MySQL version, or that a bug described in a report has already been fixed in newer MySQL versions.

Sometimes the error is platform-dependent; in such cases, it is next to impossible for us to fix anything without knowing the operating system and the version number of the platform.

If at all possible, create a repeatable, standalone testcase that doesn't involve any third-party classes.

To streamline this process, we ship a base class for testcases with Connector/J, named `'com.mysql.cj.jdbc.util.BaseBugReport'`. To create a testcase for Connector/J using this class, create your own class that inherits from `com.mysql.cj.jdbc.util.BaseBugReport` and override the methods `setUp()`, `tearDown()` and `runTest()`.

In the `setUp()` method, create code that creates your tables, and populates them with any data needed to demonstrate the bug.

In the `runTest()` method, create code that demonstrates the bug using the tables and data you created in the `setUp()` method.

In the `tearDown()` method, drop any tables you created in the `setUp()` method.

In any of the above three methods, use one of the variants of the `getConnection()` method to create a JDBC connection to MySQL:

- `getConnection()` - Provides a connection to the JDBC URL specified in `getUrl()`. If a connection already exists, that connection is returned, otherwise a new connection is created.
- `getNewConnection()` - Use this if you need to get a new connection for your bug report (that is, there is more than one connection involved).
- `getConnection(String url)` - Returns a connection using the given URL.
- `getConnection(String url, Properties props)` - Returns a connection using the given URL and properties.

If you need to use a JDBC URL that is different from `'jdbc:mysql://test'`, override the method `getUrl()` as well.

Use the `assertTrue(boolean expression)` and `assertTrue(String failureMessage, boolean expression)` methods to create conditions that must be met in your testcase demonstrating the behavior you are expecting (vs. the behavior you are observing, which is why you are most likely filing a bug report).

Finally, create a `main()` method that creates a new instance of your testcase, and calls the `run` method:

```
public static void main(String[] args) throws Exception {
    new MyBugReport().run();
}
```

Once you have finished your testcase, and have verified that it demonstrates the bug you are reporting, upload it with your bug report to <http://bugs.mysql.com/>.

Index

C

character sets
 with Connector/J, 58
client-side failover, 90
connecting
 through JDBC and Connector/J, 19
 with Unix domain socket, 63
 with Windows named pipes, 63
connection pooling, 83, 109
connection properties, 22
connection URL, 19
Connector/J
 known issues, 119
 limitations, 119
 reporting problems, 121
 troubleshooting, 111

E

error codes, 64
ER_ABORTING_CONNECTION, 64
ER_ACCESS_DENIED_ERROR, 64
ER_BAD_FIELD_ERROR, 64
ER_BAD_HOST_ERROR, 64
ER_BAD_TABLE_ERROR, 64
ER_BLOBS_AND_NO_TERMINATED, 64
ER_BLOB_CANT_HAVE_DEFAULT, 64
ER_BLOB_KEY_WITHOUT_LENGTH, 64
ER_BLOB_USED_AS_KEY, 64
ER_CANT_DO_THIS_DURING_AN_TRANSACTION, 64
ER_CANT_DROP_FIELD_OR_KEY, 64
ER_CANT_REMOVE_ALL_FIELDS, 64
ER_CANT_USE_OPTION_HERE, 64
ER_CHECK_NOT_IMPLEMENTED, 64
ER_CHECK_NO_SUCH_TABLE, 64
ER_COLLATION_CHARSET_MISMATCH, 64
ER_COLUMNACCESS_DENIED_ERROR, 64
ER_CONNECT_TO_MASTER, 64
ER_CON_COUNT_ERROR, 64
ER_DBACCESS_DENIED_ERROR, 64
ER_DERIVED_MUST_HAVE_ALIAS, 64
ER_DUP_ENTRY, 64
ER_DUP_FIELDNAME, 64
ER_DUP_KEY, 64
ER_DUP_KEYNAME, 64
ER_DUP_UNIQUE, 64
ER_EMPTY_QUERY, 64
ER_FIELD_SPECIFIED_TWICE, 64
ER_FORCING_CLOSE, 64
ER_GRANT_WRONG_HOST_OR_USER, 64
ER_HANDSHAKE_ERROR, 64

ER_HOST_IS_BLOCKED, 64
ER_HOST_NOT_PRIVILEGED, 64
ER_ILLEGAL_GRANT_FOR_TABLE, 64
ER_ILLEGAL_REFERENCE, 64
ER_INVALID_DEFAULT, 64
ER_INVALID_USE_OF_NULL, 64
ER_IPSOCK_ERROR, 64
ER_KEY_COLUMN_DOES_NOT_EXITS, 64
ER_LOCK_DEADLOCK, 64
ER_LOCK_WAIT_TIMEOUT, 64
ER_MASTER_NET_READ, 64
ER_MASTER_NET_WRITE, 64
ER_MIX_OF_GROUP_FUNC_AND_FIELDS, 64
ER_MULTIPLE_PRI_KEY, 64
ER_NET_ERROR_ON_WRITE, 64
ER_NET_FCNTL_ERROR, 64
ER_NET_PACKETS_OUT_OF_ORDER, 64
ER_NET_PACKET_TOO_LARGE, 64
ER_NET_READ_ERROR, 64
ER_NET_READ_ERROR_FROM_PIPE, 64
ER_NET_READ_INTERRUPTED, 64
ER_NET_UNCOMPRESS_ERROR, 64
ER_NET_WRITE_INTERRUPTED, 64
ER_NEW_ABORTING_CONNECTION, 64
ER_NONEXISTING_GRANT, 64
ER_NONEXISTING_TABLE_GRANT, 64
ER_NONUNIQ_TABLE, 64
ER_NON_UNIQ_ERROR, 64
ER_NOT_ALLOWED_COMMAND, 64
ER_NOT_SUPPORTED_AUTH_MODE, 64
ER_NOT_SUPPORTED_YET, 64
ER_NO_DEFAULT, 64
ER_NO_PERMISSION_TO_CREATE_USER, 64
ER_NO_REFERENCED_ROW, 64
ER_NO_SUCH_INDEX, 64
ER_NO_SUCH_TABLE, 64
ER_NULL_COLUMN_IN_INDEX, 64
ER_OPERAND_COLUMNS, 64
ER_OUTOFMEMORY, 64
ER_OUT_OF_SORTMEMORY, 64
ER_PARSE_ERROR, 64
ER_PASSWORD_ANONYMOUS_USER, 64
ER_PASSWORD_NOT_ALLOWED, 64
ER_PASSWORD_NO_MATCH, 64
ER_PRIMARY_CANT_HAVE_NULL, 64
ER_READ_ONLY_TRANSACTION, 64
ER_REGEXP_ERROR, 64
ER_REQUIRES_PRIMARY_KEY, 64
ER_ROW_IS_REFERENCED, 64
ER_SELECT_REDUCE, 64
ER_SERVER_SHUTDOWN, 64
ER_SPATIAL_CANT_HAVE_NULL, 64
ER_SUBQUERY_NO_1_ROW, 64
ER_SYNTAX_ERROR, 64

ER_TABLEACCESS_DENIED_ERROR, 64
ER_TABLENAME_NOT_ALLOWED_HERE, 64
ER_TABLE_CANT_HANDLE_AUTO_INCREMENT, 64
ER_TABLE_CANT_HANDLE_BLOB, 64
ER_TABLE_EXISTS_ERROR, 64
ER_TABLE_MUST_HAVE_COLUMNS, 64
ER_TOO_BIG_FIELDLENGTH, 64
ER_TOO_BIG_ROWSIZE, 64
ER_TOO_BIG_SELECT, 64
ER_TOO_LONG_IDENT, 64
ER_TOO_LONG_KEY, 64
ER_TOO_LONG_STRING, 64
ER_TOO_MANY_KEYS, 64
ER_TOO_MANY_KEY_PARTS, 64
ER_TOO_MANY_ROWS, 64
ER_TOO_MANY_USER_CONNECTIONS, 64
ER_UNKNOWN_CHARACTER_SET, 64
ER_UNKNOWN_COM_ERROR, 64
ER_UNKNOWN_PROCEDURE, 64
ER_UNKNOWN_STORAGE_ENGINE, 64
ER_UNKNOWN_TABLE, 64
ER_UNSUPPORTED_EXTENSION, 64
ER_USER_LIMIT_REACHED, 64
ER_WARN_DATA_OUT_OF_RANGE, 64
ER_WARN_DATA_TRUNCATED, 64
ER_WARN_NULL_TO_NOTNULL, 64
ER_WARN_TOO_FEW_RECORDS, 64
ER_WARN_TOO_MANY_RECORDS, 64
ER_WRONG_AUTO_KEY, 64
ER_WRONG_COLUMN_NAME, 64
ER_WRONG_DB_NAME, 64
ER_WRONG_FIELD_SPEC, 64
ER_WRONG_FIELD_TERMINATORS, 64
ER_WRONG_FIELD_WITH_GROUP, 64
ER_WRONG_FK_DEF, 64
ER_WRONG_GROUP_FIELD, 64
ER_WRONG_KEY_COLUMN, 64
ER_WRONG_NAME_FOR_CATALOG, 64
ER_WRONG_NAME_FOR_INDEX, 64
ER_WRONG_NUMBER_OF_COLUMNS_IN_SELECT, 64
ER_WRONG_OUTER_JOIN, 64
ER_WRONG_PARAMCOUNT_TO_PROCEDURE, 64
ER_WRONG_SUM_SELECT, 64
ER_WRONG_TABLE_NAME, 64
ER_WRONG_TYPE_FOR_VAR, 64
ER_WRONG_VALUE_COUNT, 64
ER_WRONG_VALUE_COUNT_ON_ROW, 64
ER_WRONG_VALUE_FOR_VAR, 64

F

failover
 Java clients, 87

J

J2EE
 connection pooling, 83
 load balancing, 90
JBoss application server, 103
JDBC
 and MySQL data types, 55
 background information for Connector/J, 73
 character sets, 58
 CLASSPATH, 7
 code examples, 17
 compatibility, 53
 configuration properties, 22
 driver for MySQL, 1
 SQLState codes, 64
 troubleshooting, 111, 119
 versions supported, 3

K

known issues
 Connector/J, 119

L

limitations
 Connector/J, 119
load balancing
 with Connector/J, 90, 93

M

multi-host connections
 with Connector/J, 87

N

named Pipes, 63

P

PAM authentication
 with Connector/J, 64

R

replication
 with Connector/J, 93

S

Spring framework, 105
SQLState error codes, 64
SSL, 59

T

Tomcat application server, 101
troubleshooting

Connector/J, 111
JDBC SQLState codes, 64

U

Unicode
 with Connector/J, 58
Unix domain socket, 63

X

X DevAPI
 client-side failover, 90

