

X DevAPI User Guide

Abstract

User documentation for developers using X DevAPI.

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

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

Document generated on: 2025-08-04 (revision: 83188)

Table of Contents

Preface and Legal Notices	ix
1 Overview	1
2 Connection and Session Concepts	3
2.1 Database Connection Example	3
2.2 Connecting to a Session	6
2.2.1 Connecting to a Single MySQL Server	6
2.2.2 Connecting to One of Multiple Hosts and Connection Failover	9
2.2.3 Connecting to a Single MySQL Server Using Connection Pooling	9
2.2.4 Connections Using DNS SRV Records	12
2.2.5 Connection Option Summary	13
2.2.6 Connection Attributes	13
2.3 Working with a Session Object	14
2.4 Using SQL with Session	17
2.5 Setting the Current Schema	20
2.6 Dynamic SQL	21
3 CRUD Operations	25
3.1 CRUD Operations Overview	25
3.2 Method Chaining	26
3.3 Synchronous versus Asynchronous Execution	28
3.4 Parameter Binding	30
3.5 MySQL Shell Automatic Code Execution	36
4 Working with Collections	39
4.1 Basic CRUD Operations on Collections	39
4.2 Collection Objects	43
4.2.1 Creating a Collection	43
4.2.2 Working with Existing Collections	44
4.3 Collection CRUD Function Overview	45
4.3.1 Collection.add()	45
4.3.2 Collection.find()	47
4.3.3 Collection.modify()	54
4.3.4 Collection.remove()	62
4.4 Indexing Collections	64
4.5 Single Document Operations	67
4.6 JSON Schema Validation	68
5 Working with Documents	73
5.1 Creating Documents	73
5.2 Working with Document IDs	75
5.3 Understanding Document IDs	77
6 Working with Relational Tables	81
6.1 Syntax of the SQL CRUD Functions	84
7 Working with Relational Tables and Documents	87
7.1 Collections as Relational Tables	87
8 Statement Execution	89
8.1 Transaction Handling	89
8.1.1 Processing Warnings	92
8.1.2 Error Handling	97
8.2 Working with Savepoints	101
8.3 Working with Locking	102
8.4 Working with Prepared Statements	103
9 Working with Result Sets	105
9.1 Result Set Classes	105

9.2 Working with <code>AUTO_INCREMENT</code> Values	106
9.3 Working with Data Sets	106
9.4 Fetching All Data Items at Once	110
9.5 Working with SQL Result Sets	112
9.6 Working with Metadata	122
9.7 Support for Language Native Iterators	123
10 Building Expressions	125
10.1 Expression Strings	125
10.1.1 Boolean Expression Strings	125
10.1.2 Value Expression Strings	125
11 CRUD EBNF Definitions	127
11.1 Session Objects and Functions	127
11.2 Schema Objects and Functions	129
11.3 Collection CRUD Functions	132
11.4 Collection Index Management Functions	134
11.5 Table CRUD Functions	134
11.6 Result Functions	136
11.7 Other EBNF Definitions	139
12 Expressions EBNF Definitions	147
13 Implementation Notes	163
13.1 MySQL Connector Notes	163
13.2 MySQL Shell X DevAPI extensions	163
13.3 MySQL Connector/Node.js Notes	163
13.4 MySQL Connector/J Notes	163

List of Figures

2.1 Connection URI	13
3.1 Database Object - Class Diagram	26
6.1 Table.insert() Syntax Diagram	85
6.2 Table.select() Syntax Diagram	86
6.3 Table.update() Syntax Diagram	86
6.4 Table.delete() Syntax Diagram	86
9.1 Result Classes	106
11.1 Session	128
11.2 SqlExecute	128
11.3 SQLPlaceholderValues	129
11.4 SQLPlaceholderName	129
11.5 SQLLiteral	129
11.6 Schema	130
11.7 Collection	131
11.8 Table	132
11.9 CollectionFindFunction	132
11.10 CollectionModifyFunction	133
11.11 CollectionAddFunction	133
11.12 CollectionRemoveFunction	134
11.13 CollectionCreateIndexFunction	134
11.14 Collection DropIndex	134
11.15 TableSelectFunction	135
11.16 TableInsertFunction	135
11.17 TableUpdateFunction	136
11.18 TableDeleteFunction	136
11.19 Result	137
11.20 DocResult	137
11.21 RowResult	137
11.22 Column	138
11.23 SqlResult	139
11.24 SearchConditionStr	139
11.25 SearchExprStrList	139
11.26 ProjectedDocumentExprStr	140
11.27 ProjectedSearchExprStrList	140
11.28 SortExprStrList	140
11.29 ExprOrLiteral	140
11.30 ExprOrLiterals	141
11.31 ExprOrLiteralOrOperand	141
11.32 PlaceholderValues	141
11.33 PlaceholderName	141
11.34 DocPath	142
11.35 Literal	142
11.36 Expression	143
11.37 Document	143
11.38 JSONExpression	144
11.39 JSONDocumentExpression	144
11.40 FunctionName	144
11.41 DocumentOrJSON	144
11.42 TableField	145
11.43 TableFields	145
12.1 ident	147

12.2 schemaQualifiedIdent	147
12.3 columnIdent	147
12.4 documentPathLastItem	148
12.5 documentPathItem	148
12.6 documentPath	148
12.7 documentField	148
12.8 argsList	149
12.9 lengthSpec	149
12.10 castType	150
12.11 functionCall	150
12.12 placeholder	150
12.13 groupedExpr	151
12.14 unaryOp	151
12.15 literal	151
12.16 jsonKeyValue	151
12.17 jsonDoc	152
12.18 jsonArray	152
12.19 atomicExpr	152
12.20 INTERVAL_UNIT	154
12.21 interval	154
12.22 intervalExpr	155
12.23 mulDivExpr	155
12.24 addSubExpr	155
12.25 shiftExpr	155
12.26 bitExpr	156
12.27 compExpr	156
12.28 ilriExpr	157
12.29 andExpr	157
12.30 orExpr	158
12.31 expr	158
12.32 DIGIT	158
12.33 FLOAT	158
12.34 INT	159
12.35 QUOTED_ID	159
12.36 ID	159
12.37 WS	160
12.38 SCHAR	161
12.39 STRING_DQ	162
12.40 STRING_SQ	162

List of Tables

2.1 Options for Configuring a Connection Pool	9
4.1 JSON Patch Merge of Document Fields	56
5.1 Different Types of Document Objects, Their Supported Languages, and Their Advantages	73
9.1 The Result Classes and the Information They Provide	105

Preface and Legal Notices

This is the X DevAPI User Guide.

Legal Notices

Copyright © 2015, 2025, Oracle and/or its affiliates.

License Restrictions

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.

Warranty Disclaimer

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.

Restricted Rights Notice

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), 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 embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

Hazardous Applications Notice

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.

Trademark Notice

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

Third-Party Content, Products, and Services Disclaimer

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.

Use of This Documentation

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

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

Access to Oracle Support for Accessibility

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

This guide explains how to use the X DevAPI and provides examples of its functionality. The X DevAPI is implemented by MySQL Shell and MySQL Connectors that support X Protocol. For more background information and instructions on how to install and get started using X DevAPI, see [Using MySQL as a Document Store](#). For quick-start tutorials introducing you to X DevAPI, see [JavaScript Quick-Start Guide: MySQL Shell for Document Store](#) and [Python Quick-Start Guide: MySQL Shell for Document Store](#). In addition to this documentation, there is developer documentation for all X DevAPI methods in the API references, available from [Connectors and APIs](#).

This section introduces the X DevAPI and provides an overview of the features available when using it to develop applications.



Important

The X DevAPI implementation in MySQL Shell can differ from the implementation in the Connector products. This guide provides an overview of using the concepts in all X DevAPI implementations.

The X DevAPI wraps powerful concepts in a simple API.

- A new high-level session concept enables you to write code that can transparently scale from single MySQL Server to a multiple server environment. See [Chapter 2, Connection and Session Concepts](#).
- Read operations are simple and easy to understand.
- Non-blocking, asynchronous calls follow common host language patterns.

The X DevAPI introduces a new, modern, and easy-to-learn way to work with your data.

- Documents are stored in Collections and have their dedicated CRUD operation set. See [Chapter 4, Working with Collections](#) and [Chapter 5, Working with Documents](#).
- Work with your existing domain objects or generate code based on structure definitions for strictly typed languages. See [Chapter 5, Working with Documents](#).
- Focus is put on working with data via CRUD operations. See [Section 3.1, “CRUD Operations Overview”](#).
- Modern practices and syntax styles are used to get away from traditional SQL-String-Building. See [Chapter 10, Building Expressions](#) for details.

Chapter 2 Connection and Session Concepts

Table of Contents

2.1 Database Connection Example	3
2.2 Connecting to a Session	6
2.2.1 Connecting to a Single MySQL Server	6
2.2.2 Connecting to One of Multiple Hosts and Connection Failover	9
2.2.3 Connecting to a Single MySQL Server Using Connection Pooling	9
2.2.4 Connections Using DNS SRV Records	12
2.2.5 Connection Option Summary	13
2.2.6 Connection Attributes	13
2.3 Working with a Session Object	14
2.4 Using SQL with Session	17
2.5 Setting the Current Schema	20
2.6 Dynamic SQL	21

This section explains the concepts of connections and sessions as used by the X DevAPI. Code examples for connecting to a MySQL Document Store (see [Using MySQL as a Document Store](#)) and using sessions are provided.

An X DevAPI session is a high-level database session concept that is different from working with traditional low-level MySQL connections. Sessions can encapsulate one or more actual MySQL connections when using the X Protocol. Use of this higher abstraction level decouples the physical MySQL setup from the application code. Sessions provide full support of X DevAPI and limited support of SQL.

For MySQL Shell, when a low-level MySQL connection to a single MySQL instance is needed this is still supported by using a `ClassicSession`, which provides full support of SQL.

Before looking at the concepts in more detail, the following examples show how to connect using a session.

2.1 Database Connection Example

The code that is needed to connect to a MySQL document store looks a lot like the traditional MySQL connection code, but now applications can establish logical sessions to MySQL server instances running the X Plugin. Sessions are produced by the `mysqlx` factory, and the returned sessions can encapsulate access to one or more MySQL server instances running X Plugin. Applications that use `Session` objects by default can be deployed on both single server setups and database clusters with no code changes.

Create an X DevAPI session using the `mysqlx.getSession(connection)` method. You pass in the connection parameters to connect to the MySQL server, such as the hostname and user, very much like the code in one of the classic APIs. The connection parameters can be specified as either a URI type string, for example `user:@localhost:33060`, or as a data dictionary, for example `{user: myuser, password: mypassword, host: example.com, port: 33060}`. See [Connecting to the Server Using URI-Like Strings or Key-Value Pairs](#) for more information.

The MySQL user account used for the connection should use either the `mysql_native_password` or `caching_sha2_password` authentication plugin, see [Pluggable Authentication](#). The server you are connecting to should have encrypted connections enabled, the default in MySQL 8.0 and later. This ensures that the client uses the X Protocol `PLAIN` password mechanism which works with user accounts that use either of the authentication plugins. If you try to connect to a server instance which does not have encrypted connections enabled, for user accounts that use the

`mysql_native_password` plugin authentication is attempted using `MYSQL41` first, and for user accounts that use `caching_sha2_password` authentication falls back to `SHA256_MEMORY`.

The following example code shows how to connect to a MySQL server and get a document from the `my_collection` collection that has the field `name` starting with `L`. The example assumes that a schema called `test` exists, and the `my_collection` collection exists. To make the example work, replace `user` with your username, and `password` with your password. If you are connecting to a different host or through a different port, change the host from `localhost` and the port from `33060`.

MySQL Shell JavaScript Code

```
var mysqlx = require('mysqlx');

// Connect to server on localhost
var mySession = mysqlx.getSession( {
    host: 'localhost', port: 33060,
    user: 'user', password: 'password' } );

var myDb = mySession.getSchema('test');

// Use the collection 'my_collection'
var myColl = myDb.getCollection('my_collection');

// Specify which document to find with Collection.find() and
// fetch it from the database with .execute()
var myDocs = myColl.find('name like :param').limit(1).
    bind('param', 'L%').execute();

// Print document
print(myDocs.fetchOne());

mySession.close();
```

MySQL Shell Python Code

```
from mysqlsh import mysqlx

# Connect to server on localhost
mySession = mysqlx.get_session( {
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password' } )

myDb = mySession.get_schema('test')

# Use the collection 'my_collection'
myColl = myDb.get_collection('my_collection')

# Specify which document to find with Collection.find() and
# fetch it from the database with .execute()
myDocs = myColl.find('name like :param').limit(1).bind('param', 'L%').execute()

# Print document
document = myDocs.fetch_one()
print(document)

mySession.close()
```

Node.js JavaScript Code

```
var mysqlx = require('@mysql/xdevapi');

// Connect to server on localhost
mysqlx
    .getSession({
```

```
        user: 'user',
        password: 'password',
        host: 'localhost',
        port: '33060'
    })
    .then(function (session) {
        var db = session.getSchema('test');
        // Use the collection 'my_collection'
        var myColl = db.getCollection('my_collection');
        // Specify which document to find with Collection.find() and
        // fetch it from the database with .execute()
        return myColl
            .find('name like :param')
            .limit(1)
            .bind('param', 'L%')
            .execute(function (doc) {
                console.log(doc);
            });
    })
    .catch(function (err) {
        // Handle error
    });
});
```

C# Code

```
// Connect to server on localhost
var mySession = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password;");

var myDb = mySession.GetSchema("test");

// Use the collection "my_collection"
var myColl = myDb.GetCollection("my_collection");

// Specify which document to find with Collection.Find() and
// fetch it from the database with .Execute()
var myDocs = myColl.Find("name like :param").Limit(1)
    .Bind("param", "L%").Execute();

// Print document
Console.WriteLine(myDocs.FetchOne());

mySession.Close();
```

Python Code

```
import mysqlx

# Connect to server on localhost
my_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

my_schema = my_session.get_schema('test')

# Use the collection 'my_collection'
my_coll = my_schema.get_collection('my_collection')

# Specify which document to find with Collection.find() and
# fetch it from the database with .execute()
docs = my_coll.find('name like :param').limit(1).bind('param', 'L%').execute()

# Print document
doc = docs.fetch_one()
print(doc)

my_session.close()
```

Java Code

```
import com.mysql.cj.xdevapi.*;

// Connect to server on localhost
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

Schema myDb = mySession.getSchema("test");

// Use the collection 'my_collection'
Collection myColl = myDb.getCollection("my_collection");

// Specify which document to find with Collection.find() and
// fetch it from the database with .execute()
DocResult myDocs = myColl.find("name like :param").limit(1).bind("param", "L%").execute();

// Print document
System.out.println(myDocs.fetchOne());

mySession.close();
```

C++ Code

```
#include <mysqlx/xdevapi.h>

// Scope controls life-time of objects such as session or schema
{
    Session sess("localhost", 33060, "user", "password");
    Schema db= sess.getSchema("test");
    // or Schema db(sess, "test");

    Collection myColl = db.getCollection("my_collection");
    // or Collection myColl(db, "my_collection");

    DocResult myDocs = myColl.find("name like :param")
        .limit(1)
        .bind("param", "L%").execute();

    cout << myDocs.fetchOne();
}
```

2.2 Connecting to a Session

There are several ways of using a session to connect to MySQL depending on the specific setup in use. This section explains the different methods available.

2.2.1 Connecting to a Single MySQL Server

In this example a connection to a local MySQL Server instance running X Plugin on the default TCP/IP port 33060 is established using the MySQL user account `user` with its password. As no other parameters are set, default values are used.

MySQL Shell JavaScript Code

```
// Passing the parameters in the { param: value } format
var dictSession = mysqlx.getSession( {
    host: 'localhost', 'port': 33060,
    user: 'user', password: 'password' } )

var dbl = dictSession.getSchema('test')
```

```
// Passing the parameters in the URI format
var uriSession = mysqlx.getSession('user:password@localhost:33060')

var db2 = uriSession.getSchema('test')
```

MySQL Shell Python Code

```
# Passing the parameters in the { param: value } format
dictSession = mysqlx.get_session( {
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password' } )

db1 = dictSession.get_schema('test')

# Passing the parameters in the URI format
uriSession = mysqlx.get_session('user:password@localhost:33060')

db2 = uriSession.get_schema('test')
```

The following example shows how to connect to a single MySQL Server instance by providing a TCP/IP address “localhost” and the same user account as before. You are prompted to enter the user name and password in this case.

MySQL Shell JavaScript Code

```
// Passing the parameters in the { param: value } format
// Query the user for the account information
print("Please enter the database user information.");
var usr = shell.prompt("Username: ", {defaultValue: "user"});
var pwd = shell.prompt("Password: ", {type: "password"});

// Connect to MySQL Server on a network machine
mySession = mysqlx.getSession( {
    host: 'localhost', 'port': 33060,
    user: usr, password: pwd} );

myDb = mySession.getSchema('test');
```

MySQL Shell Python Code

```
# Passing the parameters in the { param: value } format
# Query the user for the account information
print("Please enter the database user information.")
usr = shell.prompt("Username: ", {'defaultValue': "user"})
pwd = shell.prompt("Password: ", {'type': "password"})

# Connect to MySQL Server on a network machine
mySession = mysqlx.get_session( {
    'host': 'localhost', 'port': 33060,
    'user': usr, 'password': pwd} )

myDb = mySession.get_schema('test')
```

Node.js JavaScript Code

```
// Passing the parameters in the { param: value } format
mysqlx.getSession({ host: 'localhost', port: 33060, user: 'user', password: 'password' })
  .then(function (dictSession) {
    var dbl = dictSession.getSchema('test')
  })
// Passing the parameters in the URI format
mysqlx.getSession('user:password@localhost:33060')
  .then(function (uriSession) {
```

```
    var db2 = uriSession.getSchema('test')
})
```

C# Code

```
// Query the user for the user information
Console.WriteLine("Please enter the database user information.");
Console.Write("Username: ");
var usr = Console.ReadLine();
Console.Write("Password: ");
var pwd = Console.ReadLine();

// Connect to server on localhost using a connection URI
var mySession = MySQLX.GetSession(string.Format("mysqlx://localhost:33060/test?user={0}&password={1}", usr, pw

var myDb = mySession.getSchema("test");
```

Python Code

```
# Passing the parameters in the { param: value } format
dict_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

my_schema_1 = dict_session.get_schema('test')

# Passing the parameters in the URI format
uri_session = mysqlx.get_session('user:password@localhost:33060')

my_schema_2 = uri_session.get_schema('test')
```

Java Code

```
import com.mysql.cj.xdevapi.*;

// Connect to server on localhost using a connection URI
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

Schema myDb = mySession.getSchema("test");
```

C++ Code

```
// This code sample assumes that we have function prompt() defined somewhere.

string usr = prompt("Username:");
string pwd = prompt("Password:");

// Connect to MySQL Server on a network machine
Session mySession(SessionOption::HOST, "localhost",
                  SessionOption::PORT, 33060,
                  SessionOption::USER, usr,
                  SessionOption::PWD, pwd);

// An alternative way of defining session settings.

SessionSettings settings(SessionOption::HOST,"localhost",
                        SessionOption::PORT, 33060);

settings.set(SessionOption::USER, usr);
settings.set(SessionOption::PWD, pwd);

Session mySession(settings);

Schema myDb= mySession.getSchema("test");
```

2.2.2 Connecting to One of Multiple Hosts and Connection Failover

You can provide multiple MySQL Router or server endpoints (as host and an optional port, or Unix sockets) when creating a session. You specify a URI-like string which contains multiple hosts, ports, and optional priority. This enables the connector to perform automatic connection failover selection when any of the endpoints are not available. When multiple endpoints are available, the choice of which server is used for the session depends on whether you specify priority. If a priority is specified for each endpoint, the available endpoint with the highest priority is used. If you do not specify priority, a random available endpoint is used. This section describes how to configure multiple endpoints in a URI-like string. For more information on URI-like connection strings, see [Connecting Using URI-Like Connection Strings](#).



Note

The user and password you supply in the URI-like string apply to all of the possible endpoints and therefore the same MySQL account must exist on each of the endpoints.

To specify multiple hosts without priority, the URI-like connection string is formatted as:

```
user:password@[ (address=[host]:[port]), (address=[host]:[port]) ... ]
```

The collection of endpoints is listed as either host, port, or host and port. When setting up multiple endpoints without priority, any of the available endpoints is chosen for the connection.

To specify priorities for the endpoints, the URI-like connection string is formatted as:

```
user:password@[ (address=[host]:[port], priority=value), (address=[host]:[port], priority=value) ... ]
```

This sets up endpoints with explicit priority, for example to failover the connection to a specific endpoint in the event of another endpoint not being available. Specified priorities can range from `0` (lowest priority) to `100` (highest priority). When two endpoints share the same priority, one of them is chosen randomly.

For example, suppose you connect using an URI-like string such as:

```
mysqlx://user:password@[ (address=example1.com:33060,priority=99), (address=example2.com:33060,priority=100) ... ]
```

In this case, there are two possible endpoints, and `example2:33060` has a higher priority. When both endpoints are available, the connector connects to `example2.com:33060`. In the event that `example2.com:33060` is not available, the connector connects to `example1:33060`.

2.2.3 Connecting to a Single MySQL Server Using Connection Pooling

X DevAPI supports connection pooling, which can reduce overhead for applications that open many connections to a MySQL Server. Connections are managed as a pool by a Client object. When opening a new Session with a Client, before a new network connection is opened, an attempt is made to retrieve from the pool an existing and currently unused connection, which is then reset and reused.

A connection pool is configured using a single key-value pair (see [Connecting Using Key-Value Pairs](#)) that contains a single key named `pooling`. The value for the `pooling` key is another set of key-value pairs containing any combination of the keys described in the following table:

Table 2.1 Options for Configuring a Connection Pool

Option	Meaning	Default
<code>enabled</code>	Connection pooling enabled. When the option is set to <code>false</code> , a regular, non-pooled connection	true

Option	Meaning	Default
	is returned, and the other connection pool options listed below are ignored.	
maxSize	The maximum number of connections allowed in the pool	25
maxIdleTime	The maximum number of milliseconds a connection is allowed to idle in the queue before being closed. A zero value means infinite.	0
queueTimeout	The maximum number of milliseconds a request is allowed to wait for a connection to become available. A zero value means infinite	0

Closing the Session marks the underlying connection as unused and returns it to the Client object's connection pool.

Closing the Client object closes all connections it manages, invalidates all Sessions the Client has created, and destroys the managed pool.



Note

Connection pooling is not supported by MySQL Shell.

Node.js JavaScript Code

```
var mysqlx = require('@mysql/xdevapi');
var client = mysqlx.getClient(
  { user: 'user', host: 'localhost', port: 33060 },
  { pooling: { enabled: true, maxIdleTime: 30000, maxSize: 25, queueTimeout: 10000 } }
);
client.getSession()
  .then(session => {
    console.log(session.inspect())
    return session.close() // the connection becomes idle in the client pool
  })
  .then(() => {
    return client.getSession()
  })
  .then(session => {
    console.log(session.inspect())
    return client.close() // closes all connections and destroys the pool
  })
})
```

C# Code

```
using (Client client = MySQLX.GetClient("server=localhost;user=user;port=33060;",
  new { pooling = new { Enabled = true, MaxSize = 100, MaxIdleTime=30000, QueueTimeout = 10000 } }))
{
  using (Session session = client.GetSession())
  {
    foreach (Collection coll in session.Schema.GetCollections())
    {
      Console.WriteLine(coll.Name);
    }
  }
}
```

```
    } // session.Dispose() is called and the session becomes idle in the pool
} // client.Dispose() is called then all sessions are closed and pool is destroyed
```

Python Code

```
connection_string = {
    'host': 'localhost',
    'port': 37210,
    'user': 'user',
    'password': 'password'
}
client_options = {
    'pooling': {
        "max_size": 10,
        "max_idle_time": 30000
    }
}
client = mysqlx.get_client(connection_string, client_options)
session1 = client.get_session()
session2 = client.get_session()

# closing all the sessions
client.close()
```

Java Code

```
//Obtain new ClientFactory
ClientFactory cf = new ClientFactory();

//Obtain Client from ClientFactory
Client cli = cf.getClient(this.baseUrl, "{\"pooling\":{\"enabled\":true, \"maxSize\":8,
    \"maxIdleTime\":30000, \"queueTimeout\":10000} }");
Session sess = cli.getSession();

//Use Session as usual

//Close Client after use
cli.close();
```

C++ Code

```
using namespace mysqlx;

Client cli("user:password@host_name/db_name", ClientOption::POOL_MAX_SIZE, 7);
Session sess = cli.getSession();

// use Session sess as usual

cli.close(); // close all Sessions
```

Connector/C++ Code using X DevAPI for C

```
char error_buf[255];
int error_code;

mysqlx_client_t *cli
= mysqlx_get_client_from_url(
    "user:password@host_name/db_name", "{ \"maxSize\": 7 }", error_buf, &error_code
);
mysqlx_session_t *sess = mysqlx_get_session_from_client(cli);

// use sess as before

mysqlx_close_client(cli); // close session sess
```

2.2.4 Connections Using DNS SRV Records

X DevAPI supports the use of DNS SRV records for connecting to MySQL servers. A client that receives a DNS SRV lookup result attempts to connect to the MySQL server on each of the listed hosts in order of preference, based on the priority and weighting assigned to each host by the DNS administrator. A failure to connect occurs only if the client cannot connect to any of the servers. This section focuses on use of DNS SRV within X DevAPI applications. For general information about DNS SRV support in MySQL, see [Connecting to the Server Using DNS SRV Records](#).

MySQL Connectors that implement X DevAPI can request DNS SRV record lookup by specifying `mysqlx+srv` as the scheme element of the URI-like connection string, along with the DNS SRV name. For example:

```
mysqlx+srv://_mysqlx._tcp.example.com/db?options
```

A DNS SRV name consists of a service, protocol, and domain, with the service and protocol each prefixed by an underscore. In the example, `mysqlx` indicates the X Protocol service and `tcp` indicates the TCP protocol.

The X DevAPI `mysqlx.getSession()` method, and the `mysqlx.getClient()` method for connection pooling, validate connection information with this protocol scheme extension, and handle the resulting DNS SRV record as a list of hosts for the purposes of failover behavior and connection pooling. The priority and weighting specified in the DNS SRV record is respected.

MySQL Connectors also have connector-specific options to request DNS SRV record lookup both for X Protocol connections and for classic MySQL protocol connections. For details, see the documentation for individual MySQL Connectors.



Note

MySQL Shell does not currently support DNS SRV records.

When DNS SRV record lookup is used, clients generally must apply these rules for connection requests (there may be connector-specific exceptions):

- The request must specify the full DNS SRV record name, with the service and protocol names prefixed by underscores. For example, this DNS SRV record relates to an X Protocol service implemented over TCP that can be provided by multiple servers in the installation:

Name	TTL	Class	Priority	Weight	Port	Target
_mysqlx._tcp.example.com.	86400	IN	SRV	0	5	33060 server1.example.com.
_mysqlx._tcp.example.com.	86400	IN	SRV	0	10	33060 server2.example.com.
_mysqlx._tcp.example.com.	86400	IN	SRV	10	5	33060 server3.example.com.
_mysqlx._tcp.example.com.	86400	IN	SRV	20	5	33060 server4.example.com.

A client can specify that DNS SRV record using syntax like this:

```
var client = mysqlx.getClient("mysqlx+srv://_mysqlx._tcp.example.com")
```

- The request must not specify multiple host names.
- The request must not specify a port number.
- Only TCP connections are supported. Unix socket files, Windows named pipes, and shared memory cannot be used.

Java Code

```
Session mySession = new
SessionFactory().getSession("mysqlx+srv://user:password@_mysql._tcp.example.com/db");
```

Node.js JavaScript Code

```
mysqlx.getSession({ host: '_mysqlx._tcp.example.com', resolveSrv: true })
```

C# Code

```
var session = MySQLX.GetSession("mysqlx+srv://user:password@_mysqlx._tcp.example.com.");
```

Connector/C++ Code using X DevAPI for C

```
mysqlx::Session sess(
    SessionOption::HOST, "_mysqlx._tcp.example.com",
    SessionOption::DNS_SRV, true,
    SessionOption::USER, "user",
    SessionOption::PWD, "password");
```

Python Code

```
session = mysqlx.get_session(host="tcp.example.com", dns_srv=True)
```

2.2.5 Connection Option Summary

When using an X DevAPI session the following options are available to configure the connection.

Option	Name	Optional	Default	Notes
TCP/IP Host	host	-		localhost, IPv4 host name, no IP-range
TCP/IP Port	port	Yes	33060	Standard X Plugin port is 33060
MySQL user	user	-		MySQL database user
MySQL password	password	-		The MySQL user's password

Supported authentication methods are:

- PLAIN
- MYSQL 4.1

URI elements and format.

Figure 2.1 Connection URI



ConnectURI1 ::= 'user' ':' 'password' '@' 'host' ':' 'port'

2.2.6 Connection Attributes

[Connection attributes](#) are key-value pairs that application programs can pass to the server during connection time to be stored in the PERFORMANCE_SCHEMA tables [session_account_connect_attrs](#) and [session_connect_attrs](#). There are two different kinds of connection attributes:

- Client-defined attributes are reserved key-value mappings implicitly encoded by a client. The set of client-defined attributes differ from client to client—see the X DevAPI references of your client for details. The client-defined attributes are sent to the server by default.
- User-defined attributes are key-value mappings provided by the user or application.

X DevAPI supports MySQL connection attributes through the connection parameter `connection-attributes` (`xdevapi.connection-attributes` for Connector/J), set through the `getSession()` and `getClient()` methods. Here are the different ways to use the parameter:

- These settings for the parameter are equivalent to the default behavior (that is, sending client-defined attributes) when the parameter is not used:
 - `connection-attributes`
 - `connection-attributes=`
 - `connection-attributes=true`
 - `connection-attributes=[]`
- `connection-attributes=false` prevents any connection attributes to be sent, including the client-defined ones.
- `connection-attributes=[key1=value1,key2=value2, ...]` sends user-defined connection attributes as key-value pairs alongside the client-defined connection attributes. When a value is missing from a key-value pair, a null value is set for the attribute

The following is a generic example of a connection string that configures the connection attributes:

```
mysqlx://user:password@localhost:33060?connection-attributes=[key1=value1,key2=value2,key3=]
```

2.3 Working with a Session Object

All previous examples used the `getSchema()` or `getDefaultSchema()` methods of the Session object, which return a Schema object. You use this Schema object to access Collections and Tables. Most examples make use of the X DevAPI ability to chain all object constructions, enabling you to get to the Schema object in one line. For example:

```
schema = mysqlx.getSession(...).getSchema();
```

This object chain is equivalent to the following, with the difference that the intermediate step is omitted:

```
session = mysqlx.getSession();
schema = session.getSchema();
```

There is no requirement to always chain calls until you get a Schema object, neither is it always what you want. If you want to work with the Session object, for example, to call the Session object method `getSchemas()`, there is no need to navigate down to the Schema. For example:

```
session = mysqlx.getSession(); session.getSchemas();
```

In this example the `mysqlx.getSession()` function is used to open a Session. Then the `Session.getschemas()` function is used to get a list of all available schemas and print them to the console.

MySQL Shell JavaScript Code

```
// Connecting to MySQL and working with a Session
var mysqlx = require('mysqlx');

// Connect to a dedicated MySQL server using a connection URI
var mySession = mysqlx.getSession('user:password@localhost');

// Get a list of all available schemas
var schemaList = mySession.getschemas();

print('Available schemas in this session:\n');

// Loop over all available schemas and print their name
for (index in schemaList) {
    print(schemaList[index].name + '\n');
}

mySession.close();
```

MySQL Shell Python Code

```
# Connecting to MySQL and working with a Session
from mysqlsh import mysqlx

# Connect to a dedicated MySQL server using a connection URI
mySession = mysqlx.get_session('user:password@localhost')

# Get a list of all available schemas
schemaList = mySession.get_schemas()

print('Available schemas in this session:\n')

# Loop over all available schemas and print their name
for schema in schemaList:
    print('%s\n' % schema.name)

mySession.close()
```

Node.js JavaScript Code

```
// Connecting to MySQL and working with a Session
var mysqlx = require('@mysql/xdevapi');

// Connect to a dedicated MySQL server using a connection URI
mysqlx
    .getSession('user:password@localhost')
    .then(function (mySession) {
        // Get a list of all available schemas
        return mySession.getschemas();
    })
    .then(function (schemaList) {
        console.log('Available schemas in this session:\n');

        // Loop over all available schemas and print their name
        schemaList.forEach(function (schema) {
            console.log(schema.getName() + '\n');
        });
    });
});
```

C# Code

```
// Connect to a dedicated MySQL server node using a connection URI
var mySession = MySQLX.GetSession("mysqlx://user:password@localhost:33060");

// Get a list of all available schemas
var schemaList = mySession.Getschemas();
```

```
Console.WriteLine("Available schemas in this session:");

// Loop over all available schemas and print their name
foreach (var schema in schemaList)
{
    Console.WriteLine(schema.Name);
}

mySession.Close();
```

Python Code

```
# Connector/Python
# Connecting to MySQL and working with a Session
from mysqlsh import mysqlx

# Connect to a dedicated MySQL server using a connection URI
mySession = mysqlx.get_session('user:password@localhost')

# Get a list of all available schemas
schemaList = mySession.get_schemas()

print('Available schemas in this session:\n')

# Loop over all available schemas and print their name
for schema in schemaList:
    print('%s\n' % schema.name)

mySession.close()
```

Java Code

```
import java.util.List;
import com.mysql.cj.xdevapi.*;

// Connecting to MySQL and working with a Session
// Connect to a dedicated MySQL server using a connection URI
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

// Get a list of all available schemas
List<Schema> schemaList = mySession.getschemas();

System.out.println("Available schemas in this session:");

// Loop over all available schemas and print their name
for (Schema schema : schemaList) {
    System.out.println(schema.getName());
}

mySession.close();
```

C++ Code

```
#include <mysqlx/xdevapi.h>

// Connecting to MySQL and working with a Session

// Connect to a dedicated MySQL server using a connection URI
string url = "mysqlx://localhost:33060/test?user=user&password=password";
{
    Session mySession(url);

    // Get a list of all available schemas
    std::list<Schema> schemaList = mySession.getschemas();

    cout << "Available schemas in this session:" << endl;
```

```
// Loop over all available schemas and print their name
for (Schema schema : schemaList) {
    cout << schema.getName() << endl;
}
```

2.4 Using SQL with Session

In addition to the simplified X DevAPI syntax of the Session object, the Session object has a `sql()` function that takes any SQL statement as a string.

The following example uses a Session to call an SQL Stored Procedure on the specific node.

MySQL Shell JavaScript Code

```
var mysqlx = require('mysqlx');

// Connect to server using a Session
var mySession = mysqlx.getSession('user:password@localhost');

// Switch to use schema 'test'
mySession.sql("USE test").execute();

// In a Session context the full SQL language can be used
mySession.sql("CREATE PROCEDURE my_add_one_procedure " +
    "(INOUT incr_param INT) " +
    "BEGIN " +
    "    SET incr_param = incr_param + 1;" +
    "END;").execute();
mySession.sql("SET @my_var = ?").bind(10).execute();
mySession.sql("CALL my_add_one_procedure(@my_var);").execute();
mySession.sql("DROP PROCEDURE my_add_one_procedure;").execute();

// Use an SQL query to get the result
var myResult = mySession.sql("SELECT @my_var").execute();

// Gets the row and prints the first column
var row = myResult.fetchOne();
print(row[0]);

mySession.close();
```

MySQL Shell Python Code

```
from mysqlsh import mysqlx

# Connect to server using a Session
mySession = mysqlx.get_session('user:password@localhost')

# Switch to use schema 'test'
mySession.sql("USE test").execute()

# In a Session context the full SQL language can be used
sql = """CREATE PROCEDURE my_add_one_procedure
        (INOUT incr_param INT)
        BEGIN
            SET incr_param = incr_param + 1;
        END
"""

mySession.sql(sql).execute()
mySession.sql("SET @my_var = ?").bind(10).execute()
mySession.sql("CALL my_add_one_procedure(@my_var)").execute()
mySession.sql("DROP PROCEDURE my_add_one_procedure").execute()
```

```
# Use an SQL query to get the result
myResult = mySession.sql("SELECT @my_var").execute()

# Gets the row and prints the first column
row = myResult.fetch_one()
print(row[0])

mySession.close()
```

Node.js JavaScript Code

```
var mysqlx = require('@mysql/xdevapi');
var session;

// Connect to server using a Low-Level Session
mysqlx
    .getSession('user:password@localhost')
    .then(function (s) {
        session = s;

        return session.getSchema('test');
    })
    .then(function () {
        return Promise.all([
            // Switch to use schema 'test'
            session.sql('USE test').execute(),
            // In a Session context the full SQL language can be used
            session.sql('CREATE PROCEDURE my_add_one_procedure' +
                ' (INOUT incr_param INT) ' +
                'BEGIN ' +
                '   SET incr_param = incr_param + 1;' +
                'END;').execute(),
            session.executeSql('SET @my_var = ?;', [10]).execute(),
            session.sql('CALL my_add_one_procedure(@my_var)').execute(),
            session.sql('DROP PROCEDURE my_add_one_procedure;').execute()
        ])
    })
    .then(function() {
        // Use an SQL query to get the result
        return session.sql('SELECT @my_var').execute(function (row) {
            // Print result
            console.log(row);
        });
    });
});
```

C# Code

```
// Connect to server using a Session
var mySession = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password;");

// Switch to use schema "test"
mySession.SQL("USE test").Execute();

// In a Session context the full SQL language can be used
mySession.SQL("CREATE PROCEDURE my_add_one_procedure" +
    " (INOUT incr_param INT) " +
    "BEGIN " +
    "   SET incr_param = incr_param + 1;" +
    "END;").Execute();
mySession.SQL("SET @my_var = 10;").Execute();
mySession.SQL("CALL my_add_one_procedure(@my_var);").Execute();
mySession.SQL("DROP PROCEDURE my_add_one_procedure;").Execute();

// Use an SQL query to get the result
var myResult = mySession.SQL("SELECT @my_var").Execute();
```

```
// Gets the row and prints the first column
var row = myResult.FetchOne();
Console.WriteLine(row[0]);

mySession.Close();
```

Python Code

```
# Connector/Python
from mysqlsh import mysqlx

# Connect to server using a Session
mySession = mysqlx.get_session('user:password@localhost')

# Switch to use schema 'test'
mySession.sql("USE test").execute()

# In a Session context the full SQL language can be used
sql = """CREATE PROCEDURE my_add_one_procedure
          (INOUT incr_param INT)
        BEGIN
            SET incr_param = incr_param + 1;
        END
"""

mySession.sql(sql).execute()
mySession.sql("SET @my_var = ?").bind(10).execute()
mySession.sql("CALL my_add_one_procedure(@my_var)").execute()
mySession.sql("DROP PROCEDURE my_add_one_procedure").execute()

# Use an SQL query to get the result
myResult = mySession.sql("SELECT @my_var").execute()

# Gets the row and prints the first column
row = myResult.fetch_one()
print(row[0])

mySession.close()
```

Java Code

```
import com.mysql.cj.xdevapi.*;

// Connect to server on localhost
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

// Switch to use schema 'test'
mySession.sql("USE test").execute();

// In a Session context the full SQL language can be used
mySession.sql("CREATE PROCEDURE my_add_one_procedure " + " (INOUT incr_param INT) " + " BEGIN " + " SET incr_param = incr_param + 1; ").execute();
mySession.sql("SET @my_var = ?").bind(10).execute();
mySession.sql("CALL my_add_one_procedure(@my_var)").execute();
mySession.sql("DROP PROCEDURE my_add_one_procedure").execute();

// Use an SQL query to get the result
SqlResult myResult = mySession.sql("SELECT @my_var").execute();

// Gets the row and prints the first column
Row row = myResult.fetchOne();
System.out.println(row.getInt(0));

mySession.close();
```

C++ Code

```
#include <mysqlx/xdevapi.h>

// Connect to server on localhost
string url = "mysqlx://localhost:33060/test?user=user&password=password";
Session mySession(url);

// Switch to use schema 'test'
mySession.sql("USE test").execute();

// In a Session context the full SQL language can be used
mySession.sql("CREATE PROCEDURE my_add_one_procedure "
    " (INOUT incr_param INT) "
    "BEGIN "
    "   SET incr_param = incr_param + 1;"
    "END;")
.execute();
mySession.sql("SET @my_var = ?;").bind(10).execute();
mySession.sql("CALL my_add_one_procedure(@my_var);").execute();
mySession.sql("DROP PROCEDURE my_add_one_procedure;").execute();

// Use an SQL query to get the result
auto myResult = mySession.sql("SELECT @my_var").execute();

// Gets the row and prints the first column
Row row = myResult.fetchOne();
cout << row[0] << endl;
```

When using literal/verbatim SQL the common API patterns are mostly the same compared to using DML and CRUD operations on Tables and Collections. Two differences exist: setting the current schema and escaping names.

2.5 Setting the Current Schema

A default schema for a session can be specified using the [schema](#) attribute in the [URI-like connection string or key-value pairs](#) when opening a connection session. The [Session](#) class `getDefaultSchema()` method returns the default schema for the [Session](#).

If no default schema has been selected at connection, the [Session](#) class `setCurrentSchema()` function can be used to set a current schema.

MySQL Shell JavaScript Code

```
var mysqlx = require('mysqlx');

// Direct connect with no client-side default schema specified
var mySession = mysqlx.getSession('user:password@localhost');
mySession.setCurrentSchema("test");
```

MySQL Shell Python Code

```
from mysqlsh import mysqlx

# Direct connect with no client-side default schema specified
mySession = mysqlx.get_session('user:password@localhost')
mySession.set_current_schema("test")
```

Node.js JavaScript Code

```
/*
  Connector/Node.js does not support the setCurrentSchema() method.
  One can specify the default schema in the URI-like connection string.
*/
```

C# Code

```
// Direct connect with no client-side default schema specified
var mySession = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password;");
mySession.setCurrentSchema("test");
```

Python Code

```
# Connector/Python
from mysqlsh import mysqlx

# Direct connect with no client-side default schema specified
mySession = mysqlx.get_session('user:password@localhost')
mySession.set_current_schema("test")
```

Java Code

```
/*
  Connector/J does not support the setCurrentSchema() method.
  One can specify the default schema in the URI-like connection string.
*/
```

C++ Code

```
/*
  Connector/C++ does not support the setCurrentSchema() method.
  One can specify the default schema in the URI-like connection string.
*/
```

Notice that `setCurrentSchema()` does not change the session's default schema, which remains unchanged throughout the session, or remains `null` if not set at connection. The schema set by `setCurrentSchema()` can be returned by the `getCurrentSchema()` method.

An alternative way to set the current schema is to use the `Session` class `sql()` method and the `USE db_name` statement.

2.6 Dynamic SQL

A quoting function exists to escape SQL names and identifiers. `Session.quoteName()` escapes the identifier given in accordance to the settings of the current connection.

**Note**

The quoting function must not be used to escape values. Use the value binding syntax of `Session.sql()` instead; see [Section 2.4, “Using SQL with Session”](#) for some examples.

MySQL Shell JavaScript Code

```
function createTestTable(session, name) {

  // use escape function to quote names/identifier
  quoted_name = session.quoteName(name);

  session.sql("DROP TABLE IF EXISTS " + quoted_name).execute();

  var create = "CREATE TABLE ";
  create += quoted_name;
  create += " (id INT NOT NULL PRIMARY KEY AUTO_INCREMENT)";

  session.sql(create).execute();
```

```
    return session.getCurrentSchema().getTable(name);
}

var mysqlx = require('mysqlx');

var session = mysqlx.getSession('user:password@localhost:33060/test');

var default_schema = session.getDefaultSchema().name;
session.setCurrentSchema(default_schema);

// Creates some tables
var table1 = createTestTable(session, 'test1');
var table2 = createTestTable(session, 'test2');
```

MySQL Shell Python Code

```
def createTestTable(session, name):

    # use escape function to quote names/identifier
    quoted_name = session.quote_name(name)
    session.sql("DROP TABLE IF EXISTS " + quoted_name).execute()
    create = "CREATE TABLE "
    create += quoted_name
    create += " (id INT NOT NULL PRIMARY KEY AUTO_INCREMENT)"
    session.sql(create).execute()
    return session.get_current_schema().get_table(name)

from mysqlsh import mysqlx

session = mysqlx.get_session('user:password@localhost:33060/test')

default_schema = session.get_default_schema().name
session.set_current_schema(default_schema)

# Creates some tables
table1 = createTestTable(session, 'test1')
table2 = createTestTable(session, 'test2')
```

Node.js JavaScript Code

```
var mysqlx = require('mysqlx');

function createTestTable(session, name) {
    var create = 'CREATE TABLE ';
    create += name;
    create += ' (id INT NOT NULL PRIMARY KEY AUTO_INCREMENT)';

    return session
        .sql('DROP TABLE IF EXISTS ' + name)
        .execute()
        .then(function () {
            return session.sql(create).execute();
        });
}

var session;

mysqlx
    .getSession({
        user: 'user',
        password: 'password'
    })
    .then(function (s) {
        session = s;
```

```

        return session
            .sql('use myschema')
            .execute()
    })
    .then(function () {
        // Creates some tables
        return Promise.map([
            createTestTable(session, 'test1'),
            createTestTable(session, 'test2')
        ])
    })
    .then(function () {
        session.close();
    })
});
);

```

C# Code

```

var session = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password;" );

session.SQL("use test;").Execute();
session.GetSchema("test");

// Creates some tables
var table1 = CreateTestTable(session, "test1");
var table2 = CreateTestTable(session, "test2");

private Table CreateTestTable(Session session, string name)
{
    // use escape function to quote names/identifier
    string quoted_name = `"` + name + `"";

    session.SQL("DROP TABLE IF EXISTS " + quoted_name).Execute();

    var create = "CREATE TABLE ";
    create += quoted_name;
    create += " (id INT NOT NULL PRIMARY KEY AUTO_INCREMENT)";

    session.SQL(create).Execute();

    return session.Schema.GetTable(name);
}

```

Python Code

```

# Connector/Python
def createTestTable(session, name):

    # use escape function to quote names/identifier
    quoted_name = session.quote_name(name)
    session.sql("DROP TABLE IF EXISTS " + quoted_name).execute()
    create = "CREATE TABLE "
    create += quoted_name
    create += " (id INT NOT NULL PRIMARY KEY AUTO_INCREMENT)"
    session.sql(create).execute()
    return session.get_current_schema().get_table(name)

from mysqlsh import mysqlx

session = mysqlx.get_session('user:password@localhost:33060/test')

default_schema = session.get_default_schema().name
session.set_current_schema(default_schema)

# Creates some tables
table1 = createTestTable(session, 'test1')

```

```
table2 = createTestTable(session, 'test2')
```

Java Code

```
Java does not currently support the quoteName() method.
```

C++ Code

```
#include <mysqlx/xdevapi.h>

// Note: The following features are not yet implemented by
// Connector/C++:
// - DataSoure configuration files,
// - quoteName() method.

Table createTestTable(Session &session, const string &name)
{
    string quoted_name = string(``)
        + session.getDefaultSchemaName()
        + L"`.`" + name + L"`;
    session.sql(string("DROP TABLE IF EXISTS") + quoted_name).execute();

    string create = "CREATE TABLE ";
    create += quoted_name;
    create += L"(id INT NOT NULL PRIMARY KEY AUTO_INCREMENT)";

    session.sql(create).execute();
    return session.getDefaultSchema().getTable(name);
}

Session session(33060, "user", "password");

Table table1 = createTestTable(session, "test1");
Table table2 = createTestTable(session, "test2");
```

Code that uses X DevAPI does not need to escape identifiers. This is true for working with collections and for working with relational tables.

Chapter 3 CRUD Operations

Table of Contents

3.1 CRUD Operations Overview	25
3.2 Method Chaining	26
3.3 Synchronous versus Asynchronous Execution	28
3.4 Parameter Binding	30
3.5 MySQL Shell Automatic Code Execution	36

This section explains how to use the X DevAPI for Create Read, Update, and Delete (CRUD) operations.

MySQL's core domain has always been working with relational tables. X DevAPI extends this domain by adding support for CRUD operations that can be run against collections of documents. This section explains how to use these.

3.1 CRUD Operations Overview


CRUD operations are available as methods, which operate on Schema objects. The available Schema objects consist of Collection objects containing Documents, or Table objects consisting of rows and Collections containing Documents.

The following table shows the available CRUD operations for both Collection and Table objects.

Operation	Document	Relational
Create	Section 4.3.1, "Collection.add()"	<code>Table.insert()</code>
Read	Section 4.3.2, "Collection.find()"	<code>Table.select()</code>
Update	Section 4.3.3, "Collection.modify()"	<code>Table.update()</code>
Delete	Section 4.3.4, "Collection.remove()"	<code>Table.delete()</code>

Database Object Classes

Figure 3.1 Database Object - Class Diagram



3.2 Method Chaining

X DevAPI supports a number of modern practices to make working with CRUD operations easier and to fit naturally into modern development environments. This section explains how to use method chaining instead of working with SQL strings or JSON structures.

The following example shows how method chaining is used instead of an SQL string when working with Session objects. The example assumes that the `test` schema exists and an `employee` table exists.

MySQL Shell JavaScript Code

```

// New method chaining used for executing an SQL SELECT statement
// Recommended way for executing queries
var employees = db.getTable('employee');

var res = employees.select(['name', 'age'])
    .where('name like :param')
    .orderBy(['name'])
    .bind('param', 'm%').execute();

// Traditional SQL execution by passing an SQL string
// It should only be used when absolutely necessary
var result = session.sql('SELECT name, age ' +
    'FROM employee ' +
    'WHERE name like ? ' +
    'ORDER BY name').bind('m%').execute();
  
```

MySQL Shell Python Code

```

# New method chaining used for executing an SQL SELECT statement
# Recommended way for executing queries
employees = db.get_table('employee')

res = employees.select(['name', 'age']) \
    .where('name like :param') \
  
```

```
.order_by(['name']) \
    .bind('param', 'm%').execute()

# Traditional SQL execution by passing an SQL string
# It should only be used when absolutely necessary
result = session.sql('SELECT name, age ' +
    'FROM employee ' +
    'WHERE name like ? ' +
    'ORDER BY name').bind('m%').execute()
```

Node.js JavaScript Code

```
// New method chaining used for executing an SQL SELECT statement
// Recommended way for executing queries
var employees = db.getTable('employee');
var promise = employees.select('name', 'age')
    .where('name like :name')
    .orderBy('name')
    .bind('m%')
    .execute();

// Traditional SQL execution by passing an SQL string
var sqlString = 'SELECT name, age ' +
    'FROM employee ' +
    'WHERE name like ? ' +
    'ORDER BY name';
var promise = db.executeSql(sqlString, 'm%').execute();
```

C# Code

```
// New method chaining used for executing an SQL SELECT statement
// Recommended way for executing queries
var employees = db.GetTable("employee");

var res = employees.Select("name", "age")
    .Where("name like :param")
    .OrderBy("name")
    .Bind("param", "m%").Execute();

// Traditional SQL execution by passing an SQL string
// It should only be used when absolutely necessary
var result = session.SQL("SELECT name, age " +
    "FROM employee " +
    "WHERE name like ? " +
    "ORDER BY name").Bind("m%").Execute();
```

Python Code

```
# Connector/Python
# New method chaining used for executing an SQL SELECT statement
# Recommended way for executing queries
employees = db.get_table('employee')

res = employees.select(['name', 'age']) \
    .where('name like :param') \
    .order_by(['name']) \
    .bind('param', 'm%').execute()

# Traditional SQL execution by passing an SQL string
# It should only be used when absolutely necessary
result = session.sql('SELECT name, age ' +
    'FROM employee ' +
    'WHERE name like ? ' +
    'ORDER BY name').bind('m%').execute()
```

Java Code

```
// New method chaining used for executing an SQL SELECT statement
// Recommended way for executing queries
Table employees = db.getTable("employee");

RowResult res = employees.select("name, age")
    .where("name like :param")
    .orderBy("name")
    .bind("param", "m%").execute();

// Traditional SQL execution by passing an SQL string
// It should only be used when absolutely necessary
SqlResult result = session.sql("SELECT name, age "
    "FROM employee "
    "WHERE name like ? "
    "ORDER BY name").bind("m%").execute();
```

C++ Code

```
// New method chaining used for executing an SQL SELECT statement
// Recommended way for executing queries
Table employees = db.getTable("employee");

RowResult res = employees.select("name", "age")
    .where("name like :param")
    .orderBy("name")
    .bind("param", "m%").execute();

// Traditional SQL execution by passing an SQL string
// It should only be used when absolutely necessary
RowResult result = session.sql("SELECT name, age "
    "FROM employee "
    "WHERE name like ? "
    "ORDER BY name").bind("m%").execute();
```

3.3 Synchronous versus Asynchronous Execution

Traditionally, many MySQL drivers used a synchronous approach when executing SQL statements. This meant that operations such as opening connections and executing queries were blocked until completion, which could take a long time. To allow for parallel execution, a developer had to write a multithreaded application.

Any MySQL client that supports the X Protocol can provide asynchronous execution, either using callbacks, Promises, or by explicitly waiting on a specific result at the moment in time when it is actually needed.

Asynchronous Operations

Using callbacks is a very common way to implement asynchronous operations. When a callback function is specified, the CRUD operation is non-blocking which means that the next statement is called immediately even though the result from the database has not yet been fetched. Only when the result is available is the callback called.

Node.js JavaScript Code

```
var employees = db.getTable('employee');

employees.select('name', 'age')
    .where('name like :name')
    .orderBy('name')
    .bind('name', 'm%')
    .execute(function (row) {
        // do something with a row
    })
```

```
    })
    .catch(err) {
        // Handle error
});
```

C# Code

```
var employees = db.GetTable("employee");

var select = employees.Select("name", "age")
    .Where("name like :name")
    .OrderBy("name")
    .Bind("name", "m%")
    .ExecuteAsync();

select.ContinueWith(t =>
{
    if (t.Exception != null)
    {
        // Handle error
    }
    // Do something with the resultset
});
```

Java Code

```
Table employees = db.getTable("employee");

// execute the query asynchronously, obtain a future
CompletableFuture<RowResult> rowsFuture = employees.select("name", "age")
    .where("name like :name")
    .orderBy("name")
    .bind("name", "m%").executeAsync();

// dependent functions can be attached to the CompletableFuture
```

MySQL Shell JavaScript Code

```
// Asynchronous execution is not implemented
```

MySQL Shell Python Code

```
// Asynchronous execution is not implemented
```

Python Code

```
// Asynchronous execution is not implemented
```

C++ Code

```
// Asynchronous execution is not implemented
```

Asynchronous Operations using Awaits

Some languages can use an `async/await` pattern.

C# Code

```
Task<RowResult> getEmployeesTask = employees.Select("name", "age")
    .Where("name like :name").OrderBy("name")
    .Bind("name", "m%").ExecuteAsync();

// Do something else while the getEmployeesTask is executing in the background

// at this point we are ready to get our results back. If it is not done,
```

```
// this will block until done
RowResult res = await getEmployeesTask;

foreach (var row in res.FetchAll())
{
    // use row object
}
```

Connector/Node.js uses asynchronous operations using Promises for all network operations. See other examples.

Java Code

```
Table employees = db.getTable("employee");

// execute the query asynchronously, obtain a future
CompletableFuture<RowResult> rowsFuture = employees.select("name", "age")
    .where("name like :name")
    .orderBy("name")
    .bind("name", "%m%").executeAsync();

// wait until it's ready
RowResult rows = rowsFuture.get();
```

Syntax Differences

Depending on which language you are using, the X DevAPI may implement a function such as `executeAsync()` in exchange for `execute([mysqlx.Async])` or in addition to `execute([mysqlx.Async])`.

For example, in a Node.js context all executions are asynchronous. Therefore, Connector/Node.js does not need to distinguish between `execute()` and `executeAsync()`. To denote the asynchronous default execution, Connector/Node.js only implements `execute()` which returns JavaScript Promise objects.

Strongly typed programming languages, such as Java or C#, can take advantage of having two distinctly named API calls for synchronous and asynchronous executions. The two calls can have different return types. For example, Connector/J can use `execute()` to return a `RowResult` or `DocResult` and `executeAsync()` to return a `CompletableFuture<T>` where the type parameter is one of the result types.

3.4 Parameter Binding

Instead of using values directly in an expression string it is good practice to separate values from the expression string. This is done using parameters in the expression string and the `bind()` function to bind values to the parameters.

Parameters can be specified in the following ways: anonymous and named.

Parameter Type	Syntax	Example	Allowed in CRUD operations	Allowed in SQL strings
Anonymous	?	'age > ?'	no	yes
Named	:<name>	'age > :age'	yes	no

The following example shows how to use the `bind()` function before an `execute()` function. For each named parameter, provide an argument to `bind()` that contains the parameter name and its value. The order in which the parameter value pairs are passed to `bind()` is of no importance. The example assumes that the `test` schema has been assigned to the variable `db` and that the collection `my_collection` exists.

MySQL Shell and Node.js JavaScript Code

```
// Collection.find() function with fixed values
var myColl = db.getCollection('my_collection');

var myRes1 = myColl.find('age = 18').execute();

// Using the .bind() function to bind parameters
var myRes2 = myColl.find('name = :param1 AND age = :param2').bind('param1','Rohit').bind('param2', 18).execute()

// Using named parameters
myColl.modify('name = :param').set('age', 55).
    bind('param', 'Nadya').execute();

// Binding works for all CRUD statements except add()
var myRes3 = myColl.find('name like :param').
    bind('param', 'R%').execute();
```

When running this with Connector/Node.js be aware that `execute()` returns a Promise. You might want to check the results to avoid errors being lost.

MySQL Shell Python Code

```
# Collection.find() function with hardcoded values
myColl = db.get_collection('my_collection')

myRes1 = myColl.find('age = 18').execute()

# Using the .bind() function to bind parameters
myRes2 = myColl.find('name = :param1 AND age = :param2').bind('param1','Rohit').bind('param2', 18).execute()

# Using named parameters
myColl.modify('name = :param').set('age', 55).bind('param', 'Nadya').execute()

# Binding works for all CRUD statements except add()
myRes3 = myColl.find('name like :param').bind('param', 'R%').execute()
```

C# Code

```
// Collection.Find() function with fixed values
var myColl = db.GetCollection("my_collection");

var myRes1 = myColl.Find("age = 18").Execute();

// Using the .Bind() function to bind parameters
var myRes2 = myColl.Find("name = :param1 AND age = :param2").Bind("param1", "Rohit").Bind("param2", 18).Execute()

// Using named parameters
myColl.Modify("name = :param").Set("age", 55)
    .Bind("param", "Nadya").Execute();

// Binding works for all CRUD statements except Add()
var myRes3 = myColl.Find("name like :param")
    .Bind("param", "R%").Execute();
```

Python Code

```
# Collection.find() function with hardcoded values
my_coll = my_schema.get_collection('my_collection')

my_res_1 = my_coll.find('age = 18').execute()

# Using the .bind() function to bind parameters
my_res_2 = my_coll.find('name = :param1 AND age = :param2').bind('param1', 'Rohit').bind('param2', 18).execute()

# Using named parameters
```

```
my_coll.modify('name = :param').set('age', 55).bind('param', 'Nadya').execute()  
  
# Binding works for all CRUD statements except add()  
my_res_3 = my_coll.find('name like :param').bind('param', 'R%').execute()
```

Java Code

```
// Collection.find() function with fixed values  
Collection myColl = db.getCollection("my_collection");  
  
DocResult myRes1 = myColl.find("age = 18").execute();  
  
// Using the .bind() function to bind parameters  
DocResult myRes2 = myColl.find("name = :param1 AND age = :param2").bind("param1", "Rohit").bind("param2", 18);  
  
// Using named parameters  
myColl.modify("name = :param").set("age", 55)  
    .bind("param", "Nadya").execute();  
  
// Using named parameters with a Map  
Map<String, Object> params = new HashMap<>();  
params.put("name", "Nadya");  
myColl.modify("name = :name").set("age", 55).bind(params).execute();  
  
// Binding works for all CRUD statements except add()  
DocResult myRes3 = myColl.find("name like :param")  
    .bind("param", "R%").execute(); }
```

C++ Code

```
/// Collection.find() function with fixed values  
Collection myColl = db.getCollection("my_collection");  
  
auto myRes1 = myColl.find("age = 18").execute();  
  
// Using the .bind() function to bind parameters  
auto myRes2 = myColl.find("name = :param1 AND age = :param2")  
    .bind("param1", "Rohit").bind("param2", 18)  
    .execute();  
  
// Using named parameters  
myColl.modify("name = :param").set("age", 55)  
    .bind("param", "Nadya").execute();  
  
// Binding works for all CRUD statements except add()  
auto myRes3 = myColl.find("name like :param")  
    .bind("param", "R%").execute();
```

Anonymous placeholders are not supported in X DevAPI. This restriction improves code clarity in CRUD command chains with multiple methods using placeholders. Regardless of the `bind()` syntax variant used there is always a clear association between parameters and placeholders based on the parameter name.

All methods of a CRUD command chain form one namespace for placeholders. In the following example, `modify()` and `set()` (or `find()` and `Fields()`) are chained. Both methods take an expression with placeholders. The placeholders refer to one combined namespace. Both use one placeholder called `:param`. A single call to `bind()` with one name value parameter for `:param` is used to assign a placeholder value to both occurrences of `:param` in the chained methods.

MySQL Shell JavaScript Code

```
// one bind() per parameter  
var myColl = db.getCollection('relatives');  
var juniors = myColl.find('alias = "jr"').execute().fetchAll();
```

```

for (var index in juniors){
    myColl.modify('name = :param').
        set('parent_name', mySQLx.expr(':param')).
        bind('param', juniors[index].name).execute();
}

```

MySQL Shell Python Code

```

# one bind() per parameter
myColl = db.get_collection('relatives')
juniors = myColl.find('alias = "jr"').execute().fetch_all()

for junior in juniors:
    myColl.modify('name = :param'). \
        set('parent_name', mySQLx.expr(':param')). \
        bind('param', junior.name).execute()

```

Node.js JavaScript Code

```

// one bind() per parameter
db
    .getCollection('relatives');
    .find('alias = "jr"')
    .execute(function (junior) {
        return myColl
            .modify('name = :param')
            .set('parent_name', mySQLx.expr(':param'))
            .bind('param', junior.name)
            .execute();
    });

```

C# Code

```

// one bind() per parameter
myColl.Find("a = :param").Fields(":param as b")
    .Bind(new { param = "c" }).Execute();

```

Python Code

```

# one bind() per parameter
my_coll = my_schema.get_collection('relatives')
juniors = my_coll.find('alias = "jr"').execute().fetch_all()

for junior in juniors:
    my_coll.modify('name = :param') \
        .set('parent_name', mySQLx.expr(':param')) \
        .bind('param', junior.name).execute()

```

Java Code

```

# one bind() per parameter
Collection myColl = db.getCollection("relatives");
DocResult juniors = myColl.find("alias = 'Jr.'").execute();
while (juniors.hasNext()){
    myColl.modify("name = :param")
        .set("parent_name", Expression.expr(":param"))
        .bind("param", juniors.next().get("name"))
        .execute();
}

```

C++ Code

```

// one bind() per parameter
Collection myColl = db.getCollection("relatives");
DocResult juniors = myColl.find("alias = 'jr'").execute();

```

```
DbDoc junior;
while ((junior = juniors.fetchOne()))
{
    myColl.modify("name = :param")
        .set("parent_name", expr(":param"))
        .bind("param", junior["name"]).execute();
}
```

It is not permitted for a named parameter to use a name that starts with a digit. For example, `:1one` and `:1` are not allowed.

Preparing CRUD Statements

Instead of directly binding and executing CRUD operations with `bind()` and `execute()` or `execute()` it is also possible to store the CRUD operation object in a variable for later execution.

The advantage of doing so is to be able to bind several sets of variables to the parameters defined in the expression strings and therefore get better performance when executing a large number of similar operations. The example assumes that the `test` schema has been assigned to the variable `db` and that the collection `my_collection` exists.

MySQL Shell JavaScript Code

```
var myColl = db.getCollection('my_collection');

// Only prepare a Collection.remove() operation, but do not run it yet
var myRemove = myColl.remove('name = :param1 AND age = :param2');

// Binding parameters to the prepared function and .execute()
myRemove.bind('param1', 'Leon').bind('param2', 39).execute();
myRemove.bind('param1', 'Johannes').bind('param2', 28).execute();

// Binding works for all CRUD statements but add()
var myFind = myColl.find('name like :param1 AND age > :param2');

var myDocs = myFind.bind('param1', 'L%').bind('param2', 20).execute();
var MyOtherDocs = myFind.bind('param1', 'J%').bind('param2', 25).execute();
```

MySQL Shell Python Code

```
myColl = db.get_collection('my_collection')

# Only prepare a Collection.remove() operation, but do not run it yet
myRemove = myColl.remove('name = :param1 AND age = :param2')

# Binding parameters to the prepared function and .execute()
myRemove.bind('param1', 'Leon').bind('param2', 39).execute()
myRemove.bind('param1', 'Johannes').bind('param2', 28).execute()

# Binding works for all CRUD statements but add()
myFind = myColl.find('name like :param1 AND age > :param2')

myDocs = myFind.bind('param1', 'L%').bind('param2', 20).execute()
MyOtherDocs = myFind.bind('param1', 'J%').bind('param2', 25).execute()
```

Node.js JavaScript Code

```
var myColl = db.getCollection('my_collection');

// Only prepare a Collection.remove() operation, but do not run it yet
var myRemove = myColl.remove('name = :param1 AND age = :param2');

// Binding parameters to the prepared function and .execute()
```

```

myRemove.bind('param1', 'Leon').bind('param2', 39).execute();
myRemove.bind('param1', 'Johannes').bind('param2', 28).execute();

// Binding works for all CRUD statements but add()
var myFind = myColl.find('name like :param1 AND age > :param2');

var myDocs = myFind.bind('param1', 'L%').bind('param2', 20).execute();
var MyOtherDocs = myFind.bind('param1', 'J%').bind('param2', 25).execute();

```

C# Code

```

var myColl = db.GetCollection("my_collection");

// Only prepare a Collection.Remove() operation, but do not run it yet
var myRemove = myColl.Remove("name = :param1 AND age = :param2");

// Binding parameters to the prepared function and .Execute()
myRemove.Bind("param1", "Leon").Bind("param2", 39).Execute();
myRemove.Bind("param1", "Johannes").Bind("param2", 28).Execute();

// Binding works for all CRUD statements but Add()
var myFind = myColl.Find("name like :param1 AND age > :param2");

var myDocs = myFind.Bind("param1", "L%").Bind("param2", 20).Execute();
var MyOtherDocs = myFind.Bind("param1", "J%").Bind("param2", 25).Execute();

```

Python Code

```

my_coll = my_schema.get_collection('my_collection')

# Only prepare a Collection.remove() operation, but do not run it yet
my_remove = my_coll.remove('name = :param1 AND age = :param2')

# Binding parameters to the prepared function and .execute()
my_remove.bind('param1', 'Leon').bind('param2', 39).execute()
my_remove.bind('param1', 'Johannes').bind('param2', 28).execute()

# Binding works for all CRUD statements but add()
my_find = my_coll.find('name like :param1 AND age > :param2')

my_docs = my_find.bind('param1', 'L%').bind('param2', 20).execute()
my_other_docs = my_find.bind('param1', 'J%').bind('param2', 25).execute()

```

Java Code

```

Collection myColl = db.getCollection("my_collection");

// Create Collection.remove() operation, but do not run it yet
RemoveStatement myRemove = myColl.remove("name = :param1 AND age = :param2");

// Binding parameters to the prepared function and .execute()
myRemove.bind("param1", "Leon").bind("param2", 39).execute();
myRemove.bind("param1", "Johannes").bind("param2", 28).execute();

// Binding works for all CRUD statements but add()
FindStatement myFind = myColl.find("name LIKE :name AND age > :age");

Map<String, Object> params = new HashMap<>();
params.put("name", "L%");
params.put("age", 20);
DocResult myDocs = myFind.bind(params).execute();
params.put("name", "J%");
params.put("age", 25);
DocResult myOtherDocs = myFind.bind(params).execute();

```

C++ Code

```
Collection myColl = db.getCollection("my_collection");

// Create Collection.remove() operation, but do not run it yet
auto myRemove = myColl.remove("name = :param1 AND age = :param2");

// Binding parameters to the prepared function and .execute()
myRemove.bind("param1", "Leon").bind("param2", 39).execute();
myRemove.bind("param1", "Johannes").bind("param2", 28).execute();

// Binding works for all CRUD statements but Add()
auto myFind = myColl.find("name like :param1 AND age > :param2");

auto myDocs = myFind.bind("param1", "L%").bind("param2", 20).execute();
auto MyOtherDocs = myFind.bind("param1", "J%").bind("param2", 25).execute();
```

3.5 MySQL Shell Automatic Code Execution

When you use X DevAPI in a programming language that fully specifies the syntax to be used, for example, when executing SQL statements through an X DevAPI session or working with any of the CRUD operations, the actual operation is performed only when the `execute()` function is called. For example:

```
var result = mySession.sql('show databases').execute()
var result2 = myColl.find().execute()
```

The call of the `execute()` function above causes the operation to be executed and returns a Result object. The returned Result object is then assigned to a variable, and the assignment is the last operation executed, which returns no data. Such operations can also return a Result object, which is used to process the information returned from the operation.

Alternatively, MySQL Shell provides the following usability features that make it easier to work with X DevAPI interactively:

- Automatic execution of CRUD and SQL operations.
- Automatic processing of results.

To achieve this functionality MySQL Shell monitors the result of the last operation executed every time you enter a statement. The combination of these features makes using the MySQL Shell interactive mode ideal for prototyping code, as operations are executed immediately and their results are displayed without requiring any additional coding. For more information see [MySQL Shell 8.4](#).

Automatic Code Execution

If MySQL Shell detects that a CRUD operation ready to execute has been returned, it automatically calls the `execute()` function. Repeating the example above in MySQL Shell and removing the assignment operation shows the operation is automatically executed.

```
mysql> mySession.sql('show databases')
mysql> myColl.find()
```

MySQL Shell executes the SQL operation, and as mentioned above, once this operation is executed a Result object is returned.

Automatic Result Processing

If MySQL Shell detects that a Result object is going to be returned, it automatically processes it, printing the result data in the best format possible. There are different types of Result objects and the format changes across them.

```
mysql -js> db.countryInfo.find().limit(1)
[
  {
    "GNP": 828,
    "IndepYear": null,
    "Name": "Aruba",
    "_id": "ABW",
    "demographics": {
      "LifeExpectancy": 78.4000015258789,
      "Population": 103000
    },
    "geography": {
      "Continent": "North America",
      "Region": "Caribbean",
      "SurfaceArea": 193
    },
    "government": {
      "GovernmentForm": "Nonmetropolitan Territory of The Netherlands",
      "HeadOfState": "Beatrix"
    }
  }
]

1 document in set (0.00 sec)
```

Chapter 4 Working with Collections

Table of Contents

4.1 Basic CRUD Operations on Collections	39
4.2 Collection Objects	43
4.2.1 Creating a Collection	43
4.2.2 Working with Existing Collections	44
4.3 Collection CRUD Function Overview	45
4.3.1 Collection.add()	45
4.3.2 Collection.find()	47
4.3.3 Collection.modify()	54
4.3.4 Collection.remove()	62
4.4 Indexing Collections	64
4.5 Single Document Operations	67
4.6 JSON Schema Validation	68

The following section explains how to work with Collections, how to use CRUD operations on Collections and return Documents.

4.1 Basic CRUD Operations on Collections

Working with collections of documents is straightforward when using X DevAPI. The following example shows the basic usage of CRUD operations (see [Section 4.3, “Collection CRUD Function Overview”](#) for more details) when working with documents: After establishing a connection to a MySQL Server instance, a new collection that can hold JSON documents is created and several documents are inserted. Then, a find operation is executed to search for a specific document from the collection. Finally, the collection is dropped again from the database. The example assumes that the `test` schema exists and that the collection `my_collection` does not exist.

MySQL Shell JavaScript Code

```
// Connecting to MySQL Server and working with a Collection

var mysqlx = require('mysqlx');

// Connect to server
var mySession = mysqlx.getSession( {
host: 'localhost', port: 33060,
user: 'user', password: 'password' } );

var myDb = mySession.getSchema('test');

// Create a new collection 'my_collection'
var myColl = myDb.createCollection('my_collection');

// Insert documents
myColl.add({ name: 'Laurie', age: 19 }).execute();
myColl.add({ name: 'Nadya', age: 54 }).execute();
myColl.add({ name: 'Lukas', age: 32 }).execute();

// Find a document
var docs = myColl.find('name like :param1 AND age < :param2').limit(1).
bind('param1', 'L%').bind('param2', 20).execute();

// Print document
print(docs.fetchOne());
```

```
// Drop the collection  
myDb.dropCollection('my_collection');
```

MySQL Shell Python Code

```
# Connecting to MySQL Server and working with a Collection  
from mysqlsh import mysqlx  
  
# Connect to server  
mySession = mysqlx.get_session( {  
    'host': 'localhost', 'port': 33060,  
    'user': 'user', 'password': 'password' } )  
  
myDb = mySession.get_schema('test')  
  
# Create a new collection 'my_collection'  
myColl = myDb.create_collection('my_collection')  
  
# Insert documents  
myColl.add({ 'name': 'Laurie', 'age': 19 }).execute()  
myColl.add({ 'name': 'Nadya', 'age': 54 }).execute()  
myColl.add({ 'name': 'Lukas', 'age': 32 }).execute()  
  
# Find a document  
docs = myColl.find('name like :param1 AND age < :param2') \  
      .limit(1) \  
      .bind('param1', 'L%') \  
      .bind('param2', 20) \  
      .execute()  
  
# Print document  
doc = docs.fetch_one()  
print(doc)  
  
# Drop the collection  
myDb.drop_collection('my_collection')
```

Node.js JavaScript Code

```
// -----  
// Connecting to MySQL Server and working with a Collection  
var mysqlx = require('@mysql/xdevapi');  
var db;  
  
// Connect to server  
mysqlx  
  .getSession({  
    user: 'user',  
    password: 'password',  
    host: 'localhost',  
    port: '33060',  
  })  
  .then(function (session) {  
    db = session.getSchema('test');  
    // Create a new collection 'my_collection'  
    return db.createCollection('my_collection');  
  })  
  .then(function (myColl) {  
    // Insert documents  
    return Promise  
      .all([  
        myColl.add({ name: 'Laurie', age: 19 }).execute(),  
        myColl.add({ name: 'Nadya', age: 54 }).execute(),  
        myColl.add({ name: 'Lukas', age: 32 }).execute()  
      ])
```

```

        .then(function () {
            // Find a document
            return myColl
                .find('name like :name && age < :age')
                .bind({ name: 'L%', age: 20 })
                .limit(1)
                .execute(function (doc) {
                    // Print document
                    console.log(doc);
                });
        })
    .then(function(docs) {
        // Drop the collection
        return db.dropCollection('my_collection');
    })
    .catch(function(err) {
        // Handle error
    });
}

```

C# Code

```

// Connecting to MySQL Server and working with a Collection

// Connect to server
var mySession = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password;");

var myDb = mySession.GetSchema("test");

// Create a new collection "my_collection"
var myColl = myDb.CreateCollection("my_collection");

// Insert documents
myColl.Add(new { name = "Laurie", age = 19 }).Execute();
myColl.Add(new { name = "Nadya", age = 54 }).Execute();
myColl.Add(new { name = "Lukas", age = 32 }).Execute();

// Find a document
var docs = myColl.Find("name like :param1 AND age < :param2").Limit(1)
    .Bind("param1", "L%").Bind("param2", 20).Execute();

// Print document
Console.WriteLine(docs.FetchOne());

// Drop the collection
myDb.DropCollection("my_collection");

```

Python Code

```

# Connecting to MySQL Server and working with a Collection

import mysqlx

# Connect to server
my_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

my_schema = my_session.get_schema('test')

# Create a new collection 'my_collection'
my_coll = my_schema.create_collection('my_collection')

# Insert documents
my_coll.add({ 'name': 'Laurie', 'age': 19 }).execute()
my_coll.add({ 'name': 'Nadya', 'age': 54 }).execute()

```

```

my_coll.add({ 'name': 'Lukas', 'age': 32 }).execute()

# Find a document
docs = my_coll.find('name like :param1 AND age < :param2') \
    .limit(1) \
    .bind('param1', 'L%') \
    .bind('param2', 20) \
    .execute()

# Print document
doc = docs.fetch_one()
print("Name: {}".format(doc['name']))
print("Age: {}".format(doc['age']))

# Drop the collection
my_session.drop_collection('test', 'my_collection')

```

Java Code

```

// Connect to server
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

Schema myDb = mySession.getSchema("test");

// Create a new collection 'my_collection'
Collection myColl = myDb.createCollection("my_collection");

// Insert documents
myColl.add("{} \\"name\":\"Laurie\", \"age\":19 }").execute();
myColl.add("{} \\"name\":\"Nadya\", \"age\":54 }").execute();
myColl.add("{} \\"name\":\"Lukas\", \"age\":32 }").execute();

// Find a document
DocResult docs = myColl.find("name like :name AND age < :age")
    .bind("name", "L%").bind("age", 20).execute();

// Print document
DbDoc doc = docs.fetchOne();
System.out.println(doc);

// Drop the collection
myDb.dropCollection("my_collection");

```

C++ Code

```

// Connecting to MySQL Server and working with a Collection

#include <mysqlx/xdevapi.h>

// Connect to server
Session session(33060, "user", "password");
Schema db = session.getSchema("test");

// Create a new collection 'my_collection'
Collection myColl = db.createCollection("my_collection");

// Insert documents
myColl.add(R"({ "name": "Laurie", "age": 19 })").execute();
myColl.add(R"({ "name": "Nadya", "age": 54 })").execute();
myColl.add(R"({ "name": "Lukas", "age": 32 })").execute();

// Find a document
DocResult docs = myColl.find("name like :param1 AND age < :param2").limit(1)
    .bind("param1", "L%").bind("param2", 20).execute();

// Print document

```

```
cout << docs.fetchOne();

// Drop the collection
db.dropCollection("my_collection");
```

4.2 Collection Objects

Documents of the same type (for example users, products) are grouped together and stored in the database as collections. X DevAPI uses Collection objects to store and retrieve documents.

4.2.1 Creating a Collection

In order to create a new collection call the `createCollection()` function from a Schema object. It returns a Collection object that can be used right away to, for example, insert documents into the database.

Optionally, for Connectors, the field `reuseExistingObject` can be set to true and passed as second parameter to prevent an error being generated if a collection with the same name already exists.

MySQL Shell JavaScript Code

```
// Create a new collection called 'my_collection'
var myColl = db.createCollection('my_collection');
```

MySQL Shell Python Code

```
# Create a new collection called 'my_collection'
myColl = db.create_collection('my_collection')
```

Node.js JavaScript Code

```
// Create a new collection called 'my_collection'
var promise = db.createCollection('my_collection');

// Create a new collection or reuse existing one
var promise = db.createCollection('my_collection', { ReuseExistingObject: true } );
```

C# Code

```
// Create a new collection called "my_collection"
var myColl = db.CreateCollection("my_collection");

// Create a new collection or reuse existing one
var myExistingColl = db.CreateCollection("my_collection", ReuseExistingObject: true);
```

Python Code

```
# Create a new collection called 'my_collection'
my_coll = my_schema.create_collection('my_collection')
# Create a new collection or reuse existing one
my_coll = my_schema.create_collection('my_collection', True)
```

Java Code

```
// Create a new collection called 'my_collection'
Collection myColl = db.createCollection("my_collection");

// Create a new collection or reuse existing one
// Second parameter is: boolean reuseExistingObject
Collection myExistingColl = db.createCollection("my_collection", true);
```

C++ Code

```
// Create a new collection called 'my_collection'
Collection myColl = db.createCollection("my_collection");
```

```
// Create a new collection or reuse existing one
Collection myExistingColl = db.createCollection("my_collection", true);
```

4.2.2 Working with Existing Collections

In order to retrieve a Collection object for an existing collection stored in the database call the `getCollection()` function from a Schema object.

If the collection does not yet exist in the database, any subsequent call of a Collection object function throws an error; for some of the Connectors, you can prevent that and catch the error right at the time of Collection retrieval by setting the `validateExistence` field to true and pass it as second parameter to `db.getCollection()`.

MySQL Shell JavaScript Code

```
// Get a collection object for 'my_collection'
var myColl = db.getCollection('my_collection');
```

MySQL Shell Python Code

```
# Get a collection object for 'my_collection'
myColl = db.get_collection('my_collection')
```

Node.js JavaScript Code

```
// Get a collection object for 'my_collection'
var collection = db.getCollection('my_collection');
```

C# Code

```
// Get a collection object for "my_collection"
var myColl = db.GetCollection("my_collection");

// Get a collection object but also ensure it exists in the database
var myColl2 = db.GetCollection("my_collection", ValidateExistence: true);
```

Python Code

```
# Get a collection object for 'my_collection'
my_coll = my_schema.get_collection('my_collection')
# Get a collection object but also ensure it exists in the database
my_coll = my_schema.get_collection('my_collection', True)
```

Java Code

```
// Get a collection object for 'my_collection'
Collection myColl = db.getCollection("my_collection");

// Get a collection object but also ensure it exists in the database
// Second parameter is: boolean requireExists
Collection myColl = db.getCollection("my_collection", true);
```

C++ Code

```
// Get a collection object for 'my_collection'
Collection myColl = db.getCollection("my_collection");

// Get a collection object but also ensure it exists in the database
Collection myColl = db.getCollection("my_collection", true);
```

The `createCollection()`, together with the `ReuseExistingObject` field set to true, can be used to create a new collection or reuse an existing collection with the given name. See [Section 4.2.1, “Creating a Collection”](#) for details.

**Note**

In most cases it is good practice to create database objects during development time and refrain from creating them on the fly during the production phase of a database project. Therefore it is best to separate the code that creates the collections in the database from the actual user application code.

4.3 Collection CRUD Function Overview

The following section explains the individual functions of the Collection object.

The most common operations to be performed on a Collection are the Create Read Update Delete (CRUD) operations. In order to speed up find operations it is recommended to make proper use of [indexes](#).

**Note**

Beyond the CRUD functions described in this section, X DevAPI also provides additional functions that operate on single documents identified by their document IDs; see [Section 4.5, “Single Document Operations”](#) for details.

4.3.1 Collection.add()

The `Collection.add()` function is for storing documents in a collection, similar to the `INSERT` statement for an SQL database. It takes a single document or a list of documents as its argument, and is executed by the `execute()` function.

The collection needs to be created with the `Schema.createCollection()` function before documents can be inserted. To insert documents into an existing collection use the `Schema.getCollection()` function to retrieve the Collection object.

The following example shows how to use the `Collection.add()` function. The example assumes that the test schema exists and that the collection `my_collection` does not exist.

MySQL Shell JavaScript Code

```
// Create a new collection
var myColl = db.createCollection('my_collection');

// Insert a document
myColl.add({ name: 'Laurie', age: 19 }).execute();

// Insert several documents at once
myColl.add([
  { name: 'Nadya', age: 54 },
  { name: 'Lukas', age: 32 } ]).execute();
```

MySQL Shell Python Code

```
# Create a new collection
myColl = db.create_collection('my_collection')

# Insert a document
myColl.add({ 'name': 'Laurie', 'age': 19 }).execute()

# Insert several documents at once
myColl.add([
  { 'name': 'Nadya', 'age': 54 },
  { 'name': 'Lukas', 'age': 32 } ]).execute()
```

Node.js JavaScript Code

Collection.add()

```
// Create a new collection
db.createCollection('myCollection').then(function (myColl) {
  return Promise.all([
    // Insert a document
    myColl
      .add({ name: 'Laurie', age: 19 })
      .execute(),
    // Insert several documents at once
    myColl
      .add([
        { name: 'Nadya', age: 54 },
        { name: 'Lukas', age: 32 }
      ])
      .execute()
  ])
});
```

C# Code

```
// Assumptions: test schema assigned to db, my_collection collection not exists

// Create a new collection
var myColl = db.CreateCollection("my_collection");

// Insert a document
myColl.Add(new { name = "Laurie", age = 19 }).Execute();

// Insert several documents at once
myColl.Add(new[] {
  new { name = "Nadya", age = 54 },
  new { name = "Lukas", age = 32 } }).Execute();
```

Python Code

```
# Create a new collection
my_coll = my_schema.create_collection('my_collection')

# Insert a document
my_coll.add({ 'name': 'Laurie', 'age': 19 }).execute()

# Insert several documents at once
my_coll.add([
  { 'name': 'Nadya', 'age': 54 },
  { 'name': 'Lukas', 'age': 32 }
]).execute()
```

Java Code

```
// Create a new collection
Collection coll = db.createCollection("payments");

// Insert a document
coll.add("{ \"name\":\"Laurie\", \"age\":19 }).execute();

// Insert several documents at once
coll.add("{ \"name\":\"Nadya\", \"age\":54 }",
         "{ \"name\":\"Lukas\", \"age\":32 }").execute();
```

C++ Code

```
// Create a new collection
Collection coll = db.createCollection("payments");

// Insert a document
coll.add(R"({ "name":"Laurie", "age":19 })").execute();
```

```
// Insert several documents at once
std::list<DbDoc> docs = {
    DbDoc(R"({ "name": "Nadya", "age": 54 })"),
    DbDoc(R"({ "name": "Lukas", "age": 32 })")
};
coll.add(docs).execute();
```

See also [CollectionAddFunction](#) for the syntax of `add()` in EBNF.

4.3.2 Collection.find()

The `find(SearchConditionStr)` function is for searching documents in a collection, similar to the `SELECT` statement for an SQL database. It takes a search condition string (`SearchConditionStr`) as a parameter to specify the documents that should be returned from the database. The `execute()` function triggers the actual execution of the `find()` operation.

The `SearchConditionStr` can be in one of these forms:

- If no `SearchConditionStr` is specified, the `find()` operation returns all the documents in the collection.

```
// Get a collection
var myColl = session.getSchema("world_x").getCollection("countryinfo");
// To return all documents in world_x:
myColl.find().execute();
```

- The most common form for a `SearchConditionStr` is:

```
JSON-path [ operator { value | JSON-path} ]
```

Here are some explanations for the different parts of a `SearchConditionStr`:

- `JSON-path`: A JSON path identifies an element in a JSON document; see [JSON Path Syntax](#) for details. Here is a short summary of the JSON path syntax:

- A JSON path starts with a scope: in MySQL's JSON document implementation, the scope of the path is always the document being operated on, represented as `$`, which is always implicitly assumed, so it can be skipped in most cases; for example, the path `$.geography.Region` is equivalent to `geography.Region`.



Note

In some cases, `$` cannot be omitted; for example:

- When the `**` wildcard is used (for example, `find("$.**.b1")`; see the discussion on wildcards below),
- When the JSON path only contains a literal string if `$` is omitted (for example, `find("$.country_name")` for finding all documents that have a `country_name` field).

- After the scope, a path consists of one or more path legs. A path leg leads from one level of the JSON tree down to the next, and consecutive paths are separated by a period `(.)`. For example: `myColl.find("geography.Continent = 'Africa'")` finds all documents that have the value `Africa` for the field `Continent` under the field `geography`.
- Elements in arrays are represented by `[N]`, where `N` is an array index, which has to be a non-negative integer.

```
myColl.add({ name: 'John', favorNums: [1, 3, 5, 7, 9] }).execute();
myColl.find("favorNums[0] = 1").execute(); //Returns the document just added
```

```
}
```

- The wildcard tokens `*` and `**` can be used in JSON paths as follows:

- `object.*` represents the values of all members under the member `object`. For example, in the `countryinfo` collection in the sample `world_x` schema, `geography.*` represents all members under the object `geography`, and `myColl.find("'Africa' in geography.*")` returns all documents that have the value `Africa` in any of the members under `geography`.
- `array[*]` represents the values of all elements in an array. For example:

```
myColl.add({ name:'John', favorNums: [1, 3, 5, 7, 9] }).execute();
myColl.add({ name:'Jane', favorNums: [2, 4, 6, 8, 10] }).execute();
myColl.find("1 in favorNums[*]").execute(); //Returns the first document added above
myColl.find("2 in favorNums[*]").execute(); //Returns the second document added above
}
```

- `[prefix]**suffix` represents all paths under the document `prefix` that end with `suffix`, regardless of the depth of the path. The following examples illustrate how `**` can be used to return different results:

```
mysql> myColl.find().execute();
{
  "a": "bar",
  "b": {
    "b1": 6,
    "b2": 7,
    "b3": {
      "b1": 99,
      "b2": 98,
      "b3": {
        "b1": 999,
        "b2": 998
      }
    }
  },
  "_id": "000061313aa100000000000000001"
}
{
  "a": "baz",
  "b": {
    "b1": 1,
    "b2": 7
  },
  "_id": "000061313aa100000000000000002"
}
{
  "a": "bbr",
  "c": 37,
  "_id": "0000613247ed000000000000000001"
}
3 documents in set (0.0007 sec)
mysql> myColl.find("$**.b2").execute();
{
  "a": "bar",
  "b": {
    "b1": 6,
    "b2": 7,
    "b3": {
      "b1": 99,
      "b2": 98,
      "b3": {
        "b1": 999,
        "b2": 998
      }
    }
  }
}
```

```

        }
    },
    "_id": "000061313aa10000000000000000001"
}
{
    "a": "baz",
    "b": {
        "b1": 1,
        "b2": 7
    },
    "_id": "000061313aa10000000000000000002"
}
2 documents in set, 1 warning (0.0008 sec)
...
mysql-javascript> myColl.find("$**.b3**.b2").execute();
{
    "a": "bar",
    "b": {
        "b1": 6,
        "b2": 7,
        "b3": {
            "b1": 99,
            "b2": 98,
            "b3": {
                "b1": 999,
                "b2": 998
            }
        }
    },
    "_id": "000061313aa10000000000000000001"
}
1 document in set, 1 warning (0.0011 sec)
...

```

The following requirements apply when using the `**` wildcard:

- `prefix` should be `$` or an element that is a document itself.
- `suffix` should be a path leg and is always required (that is, a path expression may not end in `**`).
- A path expression may not contain the sequence `***`.
- `value` is a value to be compared to an element on the [JSON-path](#). The `%` and `_` wildcard characters can be used in `value` with the `LIKE` operator, just like in a MySQL `WHERE` clause. For example:

```

myColl.find("Name LIKE 'Austra%'")
myColl.find("geography.Continent LIKE 'Asi_'"')

```

- `operator`: The following operators can be used in a `SearchConditionStr`: `OR (||)`, `AND (&&)`, `XOR`, `IS`, `NOT`, `BETWEEN`, `IN`, `LIKE`, `OVERLAPS`, `!=`, `<>`, `>`, `>=`, `<`, `<=`, `&`, `|`, `<<, >>`, `+`, `-`, `*`, `/`, `~`, and `%`. Here are some examples for using the operators:

```

myColl.find("Name = 'Australia'")
myColl.find("demographics.Population >= 1000000" )
myColl.find("demographics.LifeExpectancy BETWEEN 50 AND 60" )

```

```
myColl.find("government.HeadOfState = 'Elizabeth II' AND geography.Region = 'Caribbean'")
```

If no operator and subsequent JSON path is supplied, `find()` returns all documents for which the JSON path supplied points to some non-null elements. For example:

```
myColl.find("demographics.Population").execute();
```

Returns all documents that have a `demographics.Population` element:

```
{
  "GNP": 828,
  "_id": "00005de917d8000000000000000000",
  "Code": "ABW",
  "Name": "Aruba",
  "IndepYear": null,
  "geography": {
    "Region": "Caribbean",
    "Continent": "North America",
    "SurfaceArea": 193
  },
  "government": {
    "HeadOfState": "Beatrix",
    "GovernmentForm": "Nonmetropolitan Territory of The Netherlands"
  },
  "demographics": {
    "Population": 103000,
    "LifeExpectancy": 78.4000015258789
  }
}
{
  "GNP": 5976,
  "_id": "00005de917d80000000000000000001",
...
232 documents in set, 1 warning (0.0013 sec)
Warning (code 3986): Evaluating a JSON value in SQL boolean context does an implicit comparison against JS
if this is not what you want, consider converting JSON to an SQL numeric type with JSON_VALUE RETURNING
```

Use the `IN` operator in the `SearchConditionStr` to check for a value within all the members covered by a wildcard:

```
mysql-js> myColl.find("$**.b1").execute();
{
  "a": "bar",
  "b": {
    "b1": 6,
    "b2": 7,
    "b3": {
      "b1": 99,
      "b2": 98,
      "b3": {
        "b1": 999,
        "b2": 998
      }
    }
  },
  "_id": "000061313aa1000000000000000000001"
}
{
  "a": "baz",
  "b": {
    "b1": 1,
    "b2": 7
  },
  "_id": "000061313aa10000000000000000000002"
}
2 documents in set, 1 warning (0.0012 sec)
```

```

...
mysql-js> myColl.find("99 IN $**.b1").execute();
{
  "a": "bar",
  "b": {
    "b1": 6,
    "b2": 7,
    "b3": {
      "b1": 99,
      "b2": 98,
      "b3": {
        "b1": 999,
        "b2": 998
      }
    }
  },
  "_id": "000061313aa1000000000000000000001"
}
1 document in set (0.0016 sec)
...

```

The `OVERLAPS` operator compares two JSON fragments and returns true (1) if the two fragments have any values in any key-value pair or array element in common. For example:

```

mysql-js> myColl.find("list").execute();
{
  "_id": "1",
  "list": [
    1,
    4
  ]
}
{
  "_id": "2",
  "list": [
    4,
    7
  ]
}
2 documents in set, 1 warning (0.0010 sec)
mysql-js> myColl.find("[1,2,3] OVERLAPS $.list")
{
  "_id": "1",
  "list": [
    1,
    4
  ]
}
1 document in set (0.0006 sec)

```

Several methods such as `fields()`, `sort()`, and `limit()` can be chained to the `find()` function to further refine the result. For example:

```

myColl.find("Name LIKE 'Austra%'").fields("Code")
myColl.find("geography.Continent LIKE 'A%'").limit(10)

```

Parameter binding using `bind()` is also supported. The following example illustrates the use of `bind()` with `find()`:

MySQL Shell JavaScript Code

```

// Use the collection 'my_collection'
var myColl = db.getCollection('my_collection');

// Find a single document that has a field 'name' that starts with 'L'
var docs = myColl.find('name like :param').

```

```
    limit(1).bind('param', 'L%').execute();

print(docs.fetchOne());

// Get all documents with a field 'name' that starts with 'L'
docs = myColl.find('name like :param').
    bind('param','L%').execute();

var myDoc;
while (myDoc = docs.fetchOne()) {
    print(myDoc);
}
```

MySQL Shell Python Code

```
# Use the collection 'my_collection'
myColl = db.get_collection('my_collection')

# Find a single document that has a field 'name' that starts with 'L'
docs = myColl.find('name like :param').limit(1).bind('param', 'L%').execute()

print(docs.fetch_one())

# Get all documents with a field 'name' that starts with 'L'
docs = myColl.find('name like :param').bind('param','L%').execute()

myDoc = docs.fetch_one()
while myDoc:
    print(myDoc)
    myDoc = docs.fetch_one()
```

Node.js JavaScript Code

```
// Use the collection 'my_collection'
var myColl = db.getCollection('my_collection');

// Find a single document that has a field 'name' that starts with 'L'
myColl
    .find('name like :name')
    .bind('name', 'L%')
    .limit(1)
    .execute(function (doc) {
        console.log(doc);
    })
    .then(function () {
        // handle details
    });

// Get all documents with a field 'name' that starts with 'L'
myColl
    .find('name like :name')
    .bind('name', 'L%')
    .execute(function (doc) {
        console.log(doc);
    })
    .then(function () {
        // handle details
    });
```

C# Code

```
// Use the collection "my_collection"
var myColl = db.GetCollection("my_collection");

// Find a single document that has a field "name" that starts with "L"
var docs = myColl.Find("name like :param")
    .Limit(1).Bind("param", "L%").Execute();
```

```

Console.WriteLine(docs.FetchOne());

// Get all documents with a field "name" that starts with "L"
docs = myColl.Find("name like :param")
    .Bind("param", "L%").Execute();

while (docs.Next())
{
    Console.WriteLine(docs.Current);
}

```

Python Code

```

# Use the collection 'my_collection'
my_coll = my_schema.get_collection('my_collection')

# Find a single document that has a field 'name' that starts with 'L'
docs = my_coll.find('name like :param').limit(1).bind('param', 'L%').execute()

print(docs.fetch_one())

# Get all documents with a field 'name' that starts with 'L'
docs = my_coll.find('name like :param').bind('param', 'L%').execute()

doc = docs.fetch_one()
print(doc)

```

Java Code

```

// Use the collection 'my_collection'
Collection myColl = db.getCollection("my_collection");

// Find a single document that has a field 'name' that starts with 'L'
DocResult docs = myColl.find("name like :name").bind("name", "L%").execute();

System.out.println(docs.fetchOne());

// Get all documents with a field 'name' that starts with 'L'
docs = myColl.find("name like :name").bind("name", "L%").execute();

while (docs.hasNext()) {
    DbDoc myDoc = docs.next();
    System.out.println(myDoc);
}

```

C++ Code

```

// Use the collection 'my_collection'
Collection myColl = db.getCollection("my_collection");

// Find a single document that has a field 'name' that starts with 'L'
DocResult docs = myColl.find("name like :param")
    .limit(1).bind("param", "L%").execute();

cout << docs.fetchOne() << endl;

// Get all documents with a field 'name' that starts with 'L'
docs = myColl.find("name like :param")
    .bind("param", "L%").execute();

DbDoc myDoc;
while ((myDoc = docs.fetchOne()))
{
    cout << myDoc << endl;
}

```

See also [CollectionFindFunction](#) for the syntax of `find()` in EBNF.

4.3.3 Collection.modify()

The `modify(SearchConditionStr)` function is for modifying documents in a collection, similar to an [UPDATE](#) statement for an SQL database. It takes a search condition string (`SearchConditionStr`) as a parameter to specify the documents that are to be modified—a detailed discussion on the `SearchConditionStr` can be found in [Section 4.3.2, “Collection.find\(\)”](#).

If one or more documents are matched by the search condition string, they are modified by any of these methods that are chained after the `modify()` method. They can be chained one after another and for multiple times:



Notes

- The `_id` of a document cannot be modified or removed by the methods below.
- For any methods below that take a `DocPath` expression as one of its arguments, the following rules apply:
 - Within the `DocPath` expression, any field names containing a space or a special character must be quoted; for example, `set("name.'last name'", "Smith")`, `unset("name.'last%name'")`
 - The `DocPath` expression *cannot* contain a wildcard token (either `*` or `**`).
 - The `DocPath` expression *cannot* be null or empty.

- `set("DocPath", ExprOrLiteral)`: Set the elements matched by the Document Path (`DocPath`) expression with the value represented by the Expression or Literal (`ExprOrLiteral`) expression.

The `DocPath` expression is a [JSON path](#) expression identifying one or more JSON elements in the documents found by the `modify()` function. See discussions on the [JSON path](#) in [Section 4.3.2, “Collection.find\(\)”](#). If the element specified by `DocPath` does not exist, it is added to the document as a new element.

`ExprOrLiteral` specifies the value to be set for the element represented by `DocPath`. It can be any of the following:

- A literal value. For example, `10` or `"John"`.
- Any X DevAPI [Expression](#), wrapped in the `expr()` function (or `mysqlx.expr()` for MySQL Shell and some Connectors), so that it is not taken as a literal value. Here are some examples, which do not exhaust the possibilities of using `mysqlx.expr(Expression)` for `ExprOrLiteral`:
 - Another `DocPath` selecting a value from the document that is being modified (for example, `set("favorNums[0]", mysqlx.expr("favorNums[1]"))`, or `set("name", mysqlx.expr("$.last name"))`).
 - A functional expression that involves one or more [Expressions](#) (for example, `set("favorNums[0]", mysqlx.expr("abs(favorNums[1]))")`).
 - One or more [Expressions](#) connected by operators (for example, `set("favorNums[0]", mysqlx.expr("favorNums[1]+favorNums[2]+favorNums[3]+3"))`, or `set("SameValueOrNot", mysqlx.expr("favorNums[1] = favorNums[2]"))`).

- A JSON document (for example, `set("Greeting", mysqlx.expr("{ 'season': 'winter', 'phrase': 'Happy Holiday' }"))`)



Note

`set("$", mysqlx.expr("json_document"))` replaces all documents matched by `modify()` with the supplied `json_document`, except for the original `_id` field, which is inalterable once set at document creation.

- `unset("DocPath[, DocPath] ...")`: Delete one or more fields or array elements represented by a list of one or more `DocPath` (for example, `unset("name")`, `unset("name.'last name'", "name.'first name'")`, or `unset("favorNums[0]")`).

An error is returned if no `DocPath` is supplied, or if `DocPath` is \$ (use `remove()` instead if you want to delete a whole document).



Warning

Notice that when multiple array elements are unset or deleted, they are being removed one after another, and the same array index in a statement might, therefore, refer to different elements for each unset action. Take that into consideration when removing array elements. For example, for the document:

```
mysqljs> myColl.find("name = 'Ann'");
{
  "_id": "00006239f74a000000000000000004",
  "name": "Ann",
  "favorNums": [
    1,
    2,
    3,
    4,
    5
  ]
}
```

The following statement does not remove the first and second elements of the array as one might expect:

```
mysqljs> myColl.modify("name = 'Ann'").unset("favorNums[0]", "favorNums[1]");
Query OK, 1 item affected (0.0038 sec)

mysqljs> myColl.find("name = 'Ann'").fields("favorNums");
{
  "favorNums": [
    2,
    4,
    5
  ]
}
1 document in set (0.0007 sec)
```

Instead, it removed the first and third elements of the array. To delete the first two elements, you can do the following:

```
mysqljs> myColl.modify("name = 'Ann'").unset("favorNums[0]", "favorNums[0]");
Query OK, 1 item affected (0.0108 sec)

Rows matched: 1  Changed: 1  Warnings: 0
mysqljs > myColl.find("name = 'Ann'").fields("favorNums");
{
```

```

        "favorNums": [
            3,
            4,
            5
        ]
    }
1 document in set (0.0005 sec)

```

- `patch(Document)`: Performs a merge patch for any documents matched by `modify()` and the JSON `Document` supplied as its parameter. The operation follows the RFC 7396 specification for JSON merge patch created by the Internet Engineering Task Force (IETF). The following table explains the action on a field, which depends on the field statuses in the two documents (notice that this is a recursive operation):

Table 4.1 JSON Patch Merge of Document Fields

Field Status in the Original Document	Field Status in Patch Document	Action to be Taken on the Field in the Original Document
Any value	Value is Null	Remove field
Value B	Value A (not Null)	If either Value A or B is a scalar, replace Value B with Value A If both Value A and B are JSON objects, they are merged using the same rules described in this table (that is, the merge is applied recursively for JSON documents).
Field does not exist	Value A (not Null)	Add field with Value A
Value C	Field does not exist	No change to field

Here is a simple example of a merge using `patch()`:

```

mysql-js> myColl.find("name = 'John Doe'");
{
    "DOB": "1970-01-01",
    "_id": "0000626028c30000000000000000002",
    "name": "John Doe",
    "Phone": 1234567,
    "Standing": "Good",
    "favorNums": {
        "a": 1,
        "b": 2
    }
}
1 document in set (0.0009 sec)

mysql-js> myColl.modify("name = 'John Doe'")
    .patch({ name: "Jane Doe", DOB: null, Phone: 9876543, favorNums: { a: 3, b: 4 } });
Query OK, 1 item affected (0.0413 sec)

Rows matched: 1  Changed: 1  Warnings: 0

mysql-js> myColl.find("name = 'Jane Doe'");
{
    "_id": "0000626028c30000000000000000002",
    "name": "Jane Doe",
    "Phone": 9876543,
    "Standing": "Good",

```

```

    "favorNums": {
        "a": 3,
        "b": 4
    }
}
1 document in set (0.0008 sec)

```

- **arrayInsert(*DocPath*, *ExprOrLiteral*)**: Insert an *ExprOrLiteral* (see explanations above) into an array at the location identified by *DocPath*, shifting any following values in the array to the right. For example: `arrayInsert("favorNums[1]", 7), arrayInsert("favorNums[1]", {even: 2, odd: 3, irrational: 'pi'})`. The following rules apply:
 - If *DocPath* does not identify an array element, an error is returned.
 - If *DocPath* identifies an array position past the end of an array, the value is inserted at the end of the array.
- **arrayAppend(*DocPath*, *ExprOrLiteral*)**: Append a value represented by *ExprOrLiteral* to the end of an array identified by *DocPath*. For example, `arrayAppend("favorNums", 555)`.

Notice that if *DocPath* points to a scalar or a document value, that value is autowrapped within an array and the value represented by *ExprOrLiteral* is added to that array. For example:

```

mysql-js> myColl.find("name='Jane Doe'");
{
    "_id": "000062b0faf9000000000000000001",
    "name": "Jane Doe",
    "favorNum": 2
}
1 document in set (0.0011 sec)

mysql-js> myColl.modify("name='Jane Doe'").arrayAppend("favorNum",3);
Query OK, 1 item affected (0.0094 sec)

Rows matched: 1  Changed: 1  Warnings: 0
mysql-js> myColl.find("name='Jane Doe'");
{
    "_id": "000062b0faf9000000000000000001",
    "name": "Jane Doe",
    "favorNum": [
        2,
        3
    ]
}
1 document in set (0.0006 sec)

```

The following methods can be chained to the modification methods described above to configure the modification:

- **sort(*sortCriteriaList*)**: Sort the order in which documents are to be modified according to *sortCriteriaList*, which is either a comma-separated list or an array of *sortCriteria*. Each *sortCriteria* consists of a component name and a search order (*asc* for ascending, or *desc* for descending). For example:
 - `sort('name asc', 'age desc')`
 - `sort(['name asc', 'age desc'])`

The method is used in combination with the `limit()` method to determine which of the documents matched by `modify(SearchConditionStr)` are to be modified.

- `limit(int)`: Limits the number of documents to be modified to `int`. When chained after `sort()`, only the first `int` of documents in the sorted list are modified.

This is an example of using `sort().limit()` to limit modifications to the documents:

```
mysql> myColl.find("name like '%Doe%'");
{
  "_id": "000062b0faf900000000000000000001",
  "name": "Jane Doe",
  "favorNum": [
    2,
    3
  ]
}
{
  "_id": "000062b372f8000000000000000000001",
  "name": "Bob Doe",
  "favorNum": [
    1,
    2
  ]
}
{
  "_id": "000062b372f8000000000000000000002",
  "name": "Mark Doe",
  "favorNum": [
    7,
    8
  ]
}
{
  "_id": "000062b372f8000000000000000000003",
  "name": "John Doe",
  "favorNum": [
    0,
    4
  ]
}
mysql> myColl.modify("name like '%Doe%'").unset("favorNum").sort("name asc").limit(2);
Query OK, 2 items affected (0.0082 sec)

Rows matched: 2  Changed: 2  Warnings: 0
mysql> myColl.find("name like '%Doe%'").sort('name asc');
{
  "_id": "000062b372f8000000000000000000001",
  "name": "Bob Doe"
}
{
  "_id": "000062b0faf90000000000000000000001",
  "name": "Jane Doe"
}
{
  "_id": "000062b372f80000000000000000000003",
  "name": "John Doe",
  "favorNum": [
    0,
    4
  ]
}
{
  "_id": "000062b372f8000000000000000000002",
  "name": "Mark Doe",
  "favorNum": [
    7,
    8
  ]
}
```

```
4 documents in set (0.0068 sec)
```

Parameter binding using `bind()` is also supported. The `execute()` function triggers the actual execution of the `modify()` operation. The following example illustrates the use of `modify()`:

MySQL Shell JavaScript Code

```
// Use the collection 'my_collection'
var myColl = db.getCollection('my_collection');

// Add a new document to the collection
myColl.add({ name:"John Doe", DOB:"1970-01-01", Phone:1234567, Standing: "Good" }).execute();

// Patch the added document, adding, removing, and changing some fields
myColl.modify("name = 'John Doe'").patch({ name: "Jane Doe", DOB: null, Phone: 9876543, favorNums: [1,2,3,4] });

//Modify fields with different methods
myColl.modify("name like :param").set("Standing", "Bad").bind("param", "J%Doe").execute();
myColl.modify("name like :param").unset("Phone").bind("param", "J%Doe").execute();
myColl.modify("name like :param").arrayInsert("favorNums[1]", 7).bind("param", "J%Doe").execute();
myColl.modify("name like :param").arrayAppend("favorNums", 99).bind("param", "J%Doe").execute();
myColl.modify("name like :param").unset("favorNums[2]").bind("param", "J%Doe").execute();

var doc = myColl.find('name like :param').limit(1).bind('param', 'J%Doe').execute();

print(doc.fetchOne());
/* The output looks like:
{
  "Standing": "Bad",
  "_id": "0000626718c10000000000000000002",
  "favorNums": [
    1,
    7,
    3,
    4,
    5,
    99
  ],
  "name": "Jane Doe"
} */
```

MySQL Shell Python Code

```
# Use the collection 'my_collection'
myColl = db.get_collection('my_collection')
# Add a new document to the collection
myColl.add({ "name": "John Doe", "DOB": "1970-01-01", "Phone": 1234567, "Standing": "Good" }).execute()

# Patch the added document, adding, removing, and changing some fields
myColl.modify("name = 'John Doe'").patch({ "name": "Jane Doe", "DOB": None, "Phone": 9876543, "favorNums": [1,2,3,4] });

# Modify fields with different methods
myColl.modify("name like :param").set("Standing", "Bad").bind("param", "J%Doe").execute()
myColl.modify("name like :param").unset("Phone").bind("param", "J%Doe").execute()
myColl.modify("name like :param").array_insert("favorNums[1]", 7).bind("param", "J%Doe").execute()
myColl.modify("name like :param").array_append("favorNums", 99).bind("param", "J%Doe").execute()
myColl.modify("name like :param").unset("favorNums[2]").bind("param", "J%Doe").execute()

doc = myColl.find('name like :param').limit(1).bind('param', 'J%Doe').execute()

print(doc.fetch_one())

# The output looks like:
# {"Standing": "Bad", "_id": "0000626718c10000000000000000005", "favorNums": [1, 7, 3, 4, 5, 99], "name": "Jane Doe"}
```

Node.js JavaScript Code

```

const mysqlx = require('@mysql/xdevapi');

mysqlx.getSession('mysqlx://localhost:33060')
    .then(session=> {
        const collection = session.getSchema('mySchema').getCollection('myCollection');
        // Add a new document to the collection
        return collection.add({ name:"John Doe", DOB:"1970-01-01", Phone:1234567, Standing: "Good" })
            .execute()
        // Patch the added document, adding, removing, and changing some fields
        .then(()=> {
            return collection.modify("name = 'John Doe'")
                .patch({ name: "Jane Doe", DOB: null, Phone: 9876543, favorNums: [1,2,3,4,5] })
                .execute();
        })
        // Modify fields with different methods
        .then(()=> {
            return collection.modify("name like :param")
                .set("Standing", "Bad")
                .bind("param", "J%Doe")
                .execute();
        })
        .then(()=> {
            return collection.modify("name like :param")
                .unset("Phone").bind("param", "J%Doe")
                .bind("param", "J%Doe")
                .execute();
        })
        .then(()=> {
            return collection.modify("name like :param")
                .arrayInsert("favorNums[1]", 7)
                .bind("param", "J%Doe")
                .execute();
        })
        .then(()=> {
            return collection.modify("name like :param")
                .arrayAppend("favorNums", 99)
                .bind("param", "J%Doe")
                .execute();
        })
        .then(()=> {
            return collection.modify("name like :param")
                .unset("favorNums[2]")
                .bind("param", "J%Doe")
                .execute();
        })
        .then(()=> {
            return collection.find('name like :param')
                .limit(1).bind('param', 'J%Doe')
                .execute();
        })
        .then(res => {
            console.log(res.fetchOne());
            /* The output looks like:
            {
                _id: '00006272fb89000000000000000004',
                name: 'Jane Doe',
                Standing: 'Bad',
                favorNums: [ 1, 7, 3, 4, 5, 99 ]
            } */
        });
    });
}

```

C# Code

Collection.modify()

```
// Use the collection "my_collection"
var myColl = myDb.GetCollection("my_collection");

// Insert a document
myColl.Add(new { name = "John Doe", DOB = "1970-01-01", Phone = 1234567, Standing = "Good" }).Execute();

//Patch the added document, adding, removing, and changing some fields
myColl.Modify("name = 'John Doe'")
    .Patch(new { name = "Jane Doe", DOB = (string)null, Phone = 9876543, favorNums = new[] { 1, 2, 3, 4, 5 } })
    .Execute();

//Modify fields with different methods
myColl.Modify("name like :param").Set("Standing", "Bad").Bind("param", "J%Doe").Execute();
myColl.Modify("name like :param").Unset("Phone").Bind("param", "J%Doe").Execute();
myColl.Modify("name like :param").ArrayInsert("favorNums[1]", 7).Bind("param", "J%Doe").Execute();
myColl.Modify("name like :param").ArrayAppend("favorNums", 99).Bind("param", "J%Doe").Execute();
myColl.Modify("name like :param").Unset("favorNums[2]").Bind("param", "J%Doe").Execute();

var docs = myColl.Find("name like :param").Limit(1).Bind("param", "J%Doe").Execute();
Console.WriteLine(docs.FetchOne());

/* The output looks like:
{
    "_id": "00006274151a00000000000000000004",
    "name": "Jane Doe",
    "Standing": "Bad",
    "favorNums": [
        1,
        7,
        3,
        4,
        5,
        99
    ]
}
*/

```

Python Code

```
# Use the collection 'my_collection'
my_coll = my_schema.get_collection('my_collection')

# Add a new document to the collection
my_coll.add({ "name": "John Doe", "DOB": "1970-01-01", "Phone": 1234567, "Standing": "Good" }).execute()

# Patch the added document, adding, removing, and changing some fields
my_coll.modify("name = 'John Doe'").patch({ "name": "Jane Doe", "DOB": None, "Phone": 9876543, "favorNums": [] })

# Modify fields with different methods
my_coll.modify("name like :param").set("Standing", "Bad").bind("param", "J%Doe").execute()
my_coll.modify("name like :param").unset("Phone").bind("param", "J%Doe").execute()
my_coll.modify("name like :param").array_insert("favorNums[1]", 7).bind("param", "J%Doe").execute()
my_coll.modify("name like :param").array_append("favorNums", 99).bind("param", "J%Doe").execute()
my_coll.modify("name like :param").unset("favorNums[2]").bind("param", "J%Doe").execute()

docs = my_coll.find('name like :param').limit(1).bind('param', 'J%Doe').execute()

print(docs.fetch_one())

# The output looks like:
#[{"_id": "0000627ac9ac000000000000000007", "name": "Jane Doe", "Standing": "Bad", "favorNums": [1, 7, 3, 4, 5]}]
```

Java Code

```
// Use the collection 'my_collection'
```

Collection.remove()

```
Collection myColl = db.getCollection("my_collection");

// Add a new document to the collection
myColl.add("{ \"name\":\"John Doe\", \"DOB\" : \"1970-01-01\", \"Phone\" : 1234567, \"Standing\" : \"Good\" }");

// Patch the added document, adding, removing, and changing some fields
myColl.modify("name = 'John Doe'").patch("{ \"name\": \"Jane Doe\", \"DOB\" : null, \"Phone\" : 9876543, \"fav");

//Modify fields with different methods
myColl.modify("name like :param").set("Standing", "Bad").bind("param", "J%Doe").execute();
myColl.modify("name like :param").unset("Phone").bind("param", "J%Doe").execute();
myColl.modify("name like :param").arrayInsert("favorNums[1]", 7).bind("param", "J%Doe").execute();
myColl.modify("name like :param").arrayAppend("favorNums", 99).bind("param", "J%Doe").execute();
myColl.modify("name like :param").unset("favorNums[2]").bind("param", "J%Doe").execute();

DocResult docs = myColl.find("name like : param").limit(1).bind("param", "J%Doe").execute();

while (docs.hasNext()) {
    DbDoc myDoc = docs.next();
    System.out.println(myDoc);
}

//The output looks like:
//{"Standing":"Bad","_id":"0000627eadcb0000000000000001","favorNums":[1,7,3,4,5,99],"name":"Jane Doe"}
```

C++ Code

```
Collection myColl = db.getCollection("my_collection");

// Add a new document to the collection
myColl.add(R"({ "name": "John Doe", "DOB" : "1970-01-01", "Phone" : 1234567, "Standing" : "Good" })").execute()

// Patch the added document, adding, removing, and changing some fields
myColl
    .modify("name = 'John Doe'")
    .patch(R"({
        "name": "Jane Doe", "DOB" : null, "Phone" : 9876543, "favorNums" : [1,2,3,4,5]
    })")
    .execute();

//Modify fields with different methods
myColl.modify("name like :param").set("Standing", "Bad").bind("param", "J%Doe").execute();
myColl.modify("name like :param").unset("Phone").bind("param", "J%Doe").execute();
myColl.modify("name like :param").arrayInsert("favorNums[1]", 7).bind("param", "J%Doe").execute();
myColl.modify("name like :param").arrayAppend("favorNums", 99).bind("param", "J%Doe").execute();
myColl.modify("name like :param").unset("favorNums[2]").bind("param", "J%Doe").execute();

DocResult docs = myColl.find("name like :param").limit(1).bind("param", "J%Doe").execute();

DbDoc myDoc;
while ((myDoc = docs.fetchOne()))
{
    cout << myDoc << endl;
}

//The output looks like:
// {"_id": "0000627ea0a80000000000000000006", "name": "Jane Doe", "Standing": "Bad", "favorNums": [1, 7, 3, 4, 5,
```

See also [CollectionModifyFunction](#) for the syntax of `add()` in EBNF.

4.3.4 Collection.remove()

The `Collection.remove()` function is for removing documents in a collection, similar to the `DELETE` statement for an SQL database. It takes a search condition string (`SearchConditionStr`) as a parameter

to specify the documents that should be removed from the collection (a detailed explanation of the `SearchConditionStr` can be found in [Section 4.3.2, “Collection.find\(\)”](#)). `remove()` returns an error if no search condition string is provided, or if an empty string is provided. All documents in the collection are removed if any expression that evaluates to true without matching any document (for example, “`true`” or “`_id IS NOT NULL`”) is passed as the search condition string.

The following methods can be chained to the `remove()` method to configure the deletion:

- `limit(int)`: Limits the number of documents to be deleted to `int`.
- `sort(sortCriteriaList)`: Sort the order in which documents are to be deleted according to `sortCriteriaList`, which is either a comma-separated list or an array of `sortCriteria`. Each `sortCriteria` consists of a component name and a search order (`asc` for ascending, or `desc` for descending). For example:
 - `sort('name asc', 'age desc')`
 - `sort(['name asc', 'age desc'])`

The method is usually used in combination with the `limit()` method to determine which of the documents matched by the search condition string are to be deleted.

Parameter binding using `bind()` is also supported, and the `execute()` function triggers the actual execution of the remove operation. The following example shows how to use the `Collection.remove()` function. It assumes some documents have been added to the collection as illustrated by the code example in [Section 4.3.1, “Collection.add\(\)”](#):

MySQL Shell JavaScript Code

```
// Use the collection 'my_collection'
var myColl = db.getCollection('my_collection');

// Remove documents by criteria
myColl.remove('name like :name AND age < :age') \
  .limit(1).bind('name', 'N%').bind('age', 60).execute();
```

MySQL Shell Python Code

```
# Use the collection 'my_collection'
myColl = db.get_collection('my_collection')

# Remove documents by criteria
myColl.remove('name like :name AND age < :age') \
  .limit(1).bind('name', 'N%').bind('age', 60).execute()
```

Node.js JavaScript Code

```
// Use the collection 'my_collection'
var myColl = db.getCollection('my_collection');

// Remove documents by criteria
myColl
  .remove('name like :name && age < :age')
  .limit(1)
  .bind({ name: 'N%', age: 60 })
  .execute();
```

C# Code

```
// Use the collection "my_collection"
var myColl = db.GetCollection("my_collection");
```

```
// Remove documents by criteria
myColl.Remove("name like :name AND age < :age").Limit(1).Bind("name", "N%").Bind("age", 60).Execute();
```

Python Code

```
# Use the collection "my_collection"
my_coll = my_schema.get_collection('my_collection')

# Remove documents by criteria
my_coll.remove("name like :name AND age < :age").limit(1).bind("name", "N%").bind("age", 60).execute();
```

Java Code

```
// Use the collection 'my_collection'
Collection myColl = db.getCollection("my_collection");

// Remove documents by criteria
myColl.remove("name like :name AND age < :age").limit(1)
    .bind("name", "N%").bind("age", 60).execute();
```

C++ Code

```
// Use the collection 'my_collection'
Collection myColl = db.getCollection("my_collection");

// Remove documents by criteria
myColl.remove("name like :name AND age < :age").limit(1)
    .bind("name", "N%").bind("age", 60).execute();
```

See also [CollectionRemoveFunction](#) for the syntax of `add()` in EBNF.

4.4 Indexing Collections

To make large collections of documents more efficient to navigate you can create an index based on one or more fields found in the documents in the collection. This section describes how to index a collection.

Creating an Index

Collection indexes are ordinary MySQL indexes on virtual columns that extract data from the documents in the collection. Because MySQL cannot index JSON values directly, to enable indexing of a collection, you provide a JSON document that specifies the document's fields to be used by the index. You pass the JSON document defining the index as the `IndexDefinition` parameter to the `Collection.createIndex(name, IndexDefinition)` method. This generic example (actual syntax might vary for different programming languages) shows how to create a mandatory integer type index based on the field `count`:

```
myCollection.createIndex("count", {fields:[{"field": "$.count", "type":"INT", required:true}]});
```

This example shows how to create an index based on a text field: a zip code in this case. For a text field, you must specify a prefix length for the index, as required by MySQL Server:

```
myCollection.createIndex("zip", {fields: [{field: "$.zip", type: "TEXT(10)"}]})
```

See [Defining an Index](#) for information on the format of `IndexDefinition` and on the supported field types for indexing.

The `Collection.createIndex()` method fails with an error if an index with the same name already exists or if the index definition is not correctly formed. The name parameter is required and must be a valid index name as accepted by the SQL statement `CREATE INDEX`.

To remove an existing index use the `collection.dropIndex(string name)` method. This would delete the index with the passed name, and the operation silently succeeds if the named index does not exist.

The indexes of a collection are stored as virtual columns. To verify a created index use the `SHOW INDEX` statement. For example to use this SQL from MySQL Shell:

```
session.runSql('SHOW INDEX FROM mySchema.myCollection');
```

Defining an Index

To create an index based on the documents in a collection you need to create an `IndexDefinition` JSON document. This section explains the valid fields you can use in such a JSON document to define an index.

To define a document field to index a collection on, the type of that field must be uniform across the whole collection. In other words, the type must be consistent. The JSON document used for defining an index, such as `{fields: [{field: '$.username', type: 'TEXT'}]}`, can contain the following:

- `fields`: an array of at least one `IndexField` object, each of which describes a JSON document field to be included in the index.

A single `IndexField` description consists of the following fields:

- `field`: a string with the full document path to the document member or field to be indexed
- `type`: a string for one of the supported column types to map the field to (see [Field Data Types](#)). For numeric types, the optional `UNSIGNED` keyword can follow. For the `TEXT` type you must define the length to consider for indexing (the prefix length).
- `required`: an optional boolean that should be set to `true` if the field is required to exist in the document. Defaults to `false` for all types except `GEOJSON`, which defaults to `true`.
- `options`: an optional integer that is used as a special option flag when decoding `GEOJSON` data (see the description for `ST_GeomFromGeoJSON()` for details).
- `srid`: an optional integer to be used as the srid value when decoding `GEOJSON` data (see the description for `ST_GeomFromGeoJSON()` for details).
- `array`: An optional boolean that is set to `true` if the field contains arrays. The default value is `false`. See [Indexing Array Fields](#) for details.



Important

For MySQL 8.0.16 and earlier, fields that are JSON arrays are not supported in the index; specifying a field that contains array data does not generate an error from the server, but the index does not function correctly.

- `type`: an optional string that defines the type of index. Value is one of `INDEX` or `SPATIAL`. The default is `INDEX` and can be omitted.

Including any other fields in an `IndexDefinition` or `IndexField` JSON document which is not described above causes `collection.createIndex()` to fail with an error.

If index type is not specified or is set to `INDEX` then the resulting index is created in the same way as it would be created by issuing `CREATE INDEX`. If index type is set to `SPATIAL` then the created index is the

same as it would be created by issuing `CREATE INDEX` with the `SPATIAL` keyword, see [SPATIAL Index Optimization](#) and [Creating Spatial Indexes](#). For example:

```
myCollection.createIndex('myIndex', //  
{fields: [{field: '$.myGeoJsonField', type: 'GEOJSON', required: true}], type:'SPATIAL'})
```



Important

When using the `SPATIAL` type of index the `required` field cannot be set to `false` in `IndexField` entries.

This is an example to create an index based on multiple fields:

```
myCollection.createIndex('myIndex', {fields: [{field: '$.myField', type: 'TEXT'}, //  
{field: '$.myField2', type: 'TEXT(10)'}, {field: '$.myField3', type: 'INT'}]})
```

The values of indexed fields are converted from JSON to the type specified in the `IndexField` description using standard MySQL type conversions (see [Type Conversion in Expression Evaluation](#)), except for the `GEOJSON` type, which uses the `ST_GeomFromGeoJSON()` function for conversion. That means when using a numeric type in an `IndexField` description, an actual field value that is non-numeric is converted to 0.

The `options` and `srid` fields in `IndexField` can only be present if `type` is set to `GEOJSON`. If present, they are used as parameters for `ST_GeomFromGeoJSON()` when converting `GEOJSON` data into MySQL native `GEOMETRY` values.

Field Data Types

The following data types are supported for document fields. Type names are case-insensitive when used in the `type` field.

- `INT` [UNSIGNED]
- `TINYINT` [UNSIGNED]
- `SMALLINT` [UNSIGNED]
- `MEDIUMINT` [UNSIGNED]
- `INTEGER` [UNSIGNED]
- `BIGINT` [UNSIGNED]
- `REAL` [UNSIGNED]
- `FLOAT` [UNSIGNED]
- `DOUBLE` [UNSIGNED]
- `DECIMAL` [UNSIGNED]
- `NUMERIC` [UNSIGNED]
- `DATE`
- `TIME`
- `TIMESTAMP`
- `DATETIME`
- `TEXT(length)`

- [GEOJSON](#) (extra options: options, srid)

Indexing Array Fields

X DevAPI supports creating indexes based on array fields by setting the boolean `array` field in the `IndexField` description to `true`. For example, to create an index on the `emails` array field:

```
collection.createIndex("emails_idx", //  
    {fields: [{"field": "$.emails", "type": "CHAR(128)", "array": true}]})
```

The following restrictions apply to creating indexes based on arrays:

- For each index, only one indexed field can be an `array`
- Data types for which index on arrays can be created:
 - Numeric types: `INTEGER [UNSIGNED]` (`INT` is NOT supported)
 - Fixed-point types: `DECIMAL(m, n)` (the precision and scale values are mandatory)
 - Date and time types: `DATE`, `TIME`, and `DATETIME`
 - String types: `CHAR(n)` and `BINARY(n)`; the character or byte length `n` is mandatory (`TEXT` is NOT supported)

4.5 Single Document Operations

The CRUD commands described at [Section 4.3, “Collection CRUD Function Overview”](#) all act on a group of documents in a collection that match a filter. X DevAPI also provides the following operations, which work on single documents that are identified by their document IDs:

- `Collection.getOne(string id)` returns the document with the given `id`. This is a shortcut for `Collection.find("_id = :id").bind("id", id).execute().fetchOne()`.
- `Collection.replaceOne(string id, Document doc)` updates or replaces the document identified by `id`, if it exists, with the provided document.
- `Collection.addOrReplaceOne(string id, Document doc)` adds the given document; however, if the `id` or any other field that has a unique index on it already exists in the collection, the operation updates the matching document instead.
- `Collection.removeOne(string id)` removes the document with the given `id`. This is a shortcut for `Collection.remove("_id = :id").bind("id", id).execute()`.

Using these operations you can reference a document by its ID (see [Section 5.2, “Working with Document IDs”](#)), making operations on single documents simpler by following a “load, modify, and save” pattern such as the following:

```
doc = collection.getOne(id); // Load document of the specified id into a temporary document called doc  
doc["address"] = "123 Long Street"; //Modify the "address" field of doc  
collection.replaceOne(id, doc); // Save doc into the document with the specified id
```

Syntax of Single Document Operations

The syntax of the single document operations is as follows:

- `Document getOne(string id)`, where `id` is the document ID of the document to be retrieved. This operation returns the document, or `NULL` if no match is found. Searches for the document that has the given `id` and returns it.

- `Result replaceOne(string id, Document doc)`, where `id` is the document ID of the document to be replaced, and `doc`, which can contain expressions, is the new document for replacing the document identified by `id`. If `doc` itself contains an `_id` value and it is different from `id`, the operation fails. The operation also fails if the new document contains a unique key value that conflicts with any other document in the collection. The operation returns a `Result` object, which indicates the number of affected documents (1 or 0). If no matches are found for `id`, the function returns normally with no changes being made.
- `Result addOrReplaceOne(string id, Document doc)`, where `id` is the document ID of the document to be replaced or added (if no match can be found for the `id`), and `doc`, which can contain expressions, is the new document used for replacement or addition. If `doc` itself contains an `_id` value and it is different from `id`, the operation fails. The operation also fails if the new document contains a unique key value that conflicts with any other document in the collection. This operation returns a `Result` object, which indicates the number of affected documents (1 or 0).
- `Result removeOne(string id)`, where `id` is the document ID of the document to be removed. This operation returns a `Result` object, which indicates the number of removed documents (1 or 0, if none).

4.6 JSON Schema Validation

Collections can be configured to verify documents against a JSON schema. This enables you to require that documents have a certain structure before they can be inserted or updated in a collection. You specify a JSON schema as described at <http://json-schema.org>. Schema validation is performed by the server, which returns an error message if a document in a collection does not validate against the assigned JSON schema. For more information on JSON schema validation in MySQL, see [JSON Schema Validation Functions](#). This section describes how to configure a collection to validate documents against a JSON schema.



Note

Connector/J implements JSON schema validation very differently from the model described in this section. See [Schema Validation](#) in the [MySQL Connector/J Developer Guide](#) for details.

To enable or modify JSON schema validation, you supply to a collection a `validation` JSON object like the following:

```
{
  validation: {
    level: "off|strict",
    schema: "json-schema"
  }
}
```

Here, `validation` is a JSON object that contains the keys you can use to configure JSON schema validation. The first key is `level`, which can take the value `strict` or `off`. The second key, `schema`, is a JSON schema, as defined at <http://json-schema.org>. If the `level` key is set to `strict`, documents are validated against the `json-schema` when they are added to the collection or, if they are already in the collection, when they are updated by some operations. If a document does not validate, the server generates an error and the operation fails. If the `level` key is set to `off`, documents are not validated against the `json-schema`.

Creating a Validated Collection

To enable JSON schema validation when you create a new collection, supply a `validation` JSON object as described above. For example, to create a collection that holds longitude and latitude values and require validating those values as numbers:

MySQL Shell JavaScript Code

```
var coll = schema.createCollection("longlang", {
  validation: {
    level: "strict",
    schema: {
      "id": "http://json-schema.org/geo",
      "$schema": "http://json-schema.org/draft-06/schema#",
      "description": "A geographical coordinate",
      "type": "object",
      "properties": {
        "latitude": {
          "type": "number"
        },
        "longitude": {
          "type": "number"
        }
      },
      "required": ["latitude", "longitude"]
    }
  }
})
```

MySQL Shell Python Code

```
coll = schema.create_collection("longlang", validation={
  "level": "strict",
  "schema": {
    "id": "http://json-schema.org/geo",
    "$schema": "http://json-schema.org/draft-06/schema#",
    "description": "A geographical coordinate",
    "type": "object",
    "properties": {
      "latitude": {
        "type": "number"
      },
      "longitude": {
        "type": "number"
      }
    },
    "required": ["latitude", "longitude"]
  }
})
```

Node.js JavaScript Code

```
var coll = schema.createCollection("longlang", {
  validation: {
    level: "strict",
    schema: {
      "id": "http://json-schema.org/geo",
      "$schema": "http://json-schema.org/draft-06/schema#",
      "description": "A geographical coordinate",
      "type": "object",
      "properties": {
        "latitude": {
          "type": "number"
        },
        "longitude": {
          "type": "number"
        }
      },
      "required": ["latitude", "longitude"]
    }
  }
})
```

C# Code

```
var collOptions = CreateCollectionOptions() {
    reuseExistingObject = false,
    validation = Validation() {
        level = ValidationLevel.Strict,
        schema = "{\"id\": \"http://json-schema.org/geo\", "
            + "\">$schema\": \"http://json-schema.org/draft-06/schema#\", "
            + "description\": \"A geographical coordinate\", "
            + "type\": \"object\", "
            + "properties\": { "
            + "    \"latitude\": { "
            + "        type\": \"number\" "
            + "    }, "
            + "    \"longitude\": { "
            + "        type\": \"number\" "
            + "    } "
            + "}, "
            + "required\": [\"latitude\", \"longitude\"] "
        }
    }
};

var coll = schema.CreateCollection("longlang", collOptions);
```

Python Code

```
coll = schema.create_collection("longlang", validation={
    "level": "strict",
    "schema": {
        "id": "http://json-schema.org/geo",
        "$schema": "http://json-schema.org/draft-06/schema#",
        "description": "A geographical coordinate",
        "type": "object",
        "properties": {
            "latitude": {
                "type": "number"
            },
            "longitude": {
                "type": "number"
            }
        },
        "required": ["latitude", "longitude"]
    }
})
```

Modifying Collection Validation

You can modify a collection to control the JSON schema validation of documents. For example you can enable or disable validation, or change the JSON schema that documents are validated against.

In order to modify the JSON schema validation of a collection, supply a `validation` JSON object when calling the `Collection.modify()` method. For example, to modify a collection to disable JSON schema validation, the `validation` object would be:

```
{
    validation: {
        "level": "off"
    }
}
```

When modifying the JSON schema validation, you can supply the `level` option alone to change just the level of schema validation. For example, pass the JSON object shown above to disable JSON schema validation. This makes no change to the JSON schema previously specified and does not remove the

JSON schema from the collection. Alternatively, you can modify the schema only by passing just a new JSON schema object.

Chapter 5 Working with Documents

Table of Contents

5.1 Creating Documents	73
5.2 Working with Document IDs	75
5.3 Understanding Document IDs	77

5.1 Creating Documents

Once a collection has been created, it can store JSON documents. You store documents by passing a JSON data structure to the `Collection.add()` function. Some languages have direct support for JSON data, others have an equivalent syntax to represent that data. MySQL Connectors that implement X DevAPI aim to implement support for all JSON methods that are native to the Connectors' specific languages.

In addition, in some MySQL Connectors the generic `DbDoc` objects can be used. The most convenient way to create them is by calling the `Collection.newDoc()`. `DbDoc` is a data type to represent JSON documents and how it is implemented is not defined by X DevAPI. Languages implementing X DevAPI are free to follow an object-oriented approach with getter and setter methods, or use a C struct style with public members.

For strictly-typed languages it is possible to create class files based on the document structure definition of collections. MySQL Shell can be used to create those files.

Table 5.1 Different Types of Document Objects, Their Supported Languages, and Their Advantages

Document Objects	Supported languages	Advantages
Native JSON	Scripting languages (JavaScript, Python)	Easy to use
JSON equivalent syntax	C# (Anonymous Types, ExpandoObject)	Easy to use
DbDoc	All languages	Unified across languages
Generated Doc Classes	Strictly typed languages (C#)	Natural to use

The following example shows the different methods of inserting documents into a collection.

MySQL Shell JavaScript Code

```
// Create a new collection 'my_collection'
var myColl = db.createCollection('my_collection');

// Insert JSON data directly
myColl.add({_id: '8901', name: 'Mats', age: 21}).execute();

// Inserting several docs at once
myColl.add([ {_id: '8902', name: 'Lotte', age: 24},
  {_id: '8903', name: 'Vera', age: 39} ]).execute();
```

MySQL Shell Python Code

```
// Create a new collection 'my_collection'
var myColl = db.createCollection('my_collection');
```

```
// Insert JSON data directly
myColl.add({_id: '8901', name: 'Mats', age: 21}).execute();

// Inserting several docs at once
myColl.add([ {_id: '8902', name: 'Lotte', age: 24},
    {_id: '8903', name: 'Vera', age: 39} ]).execute();
```

Node.js JavaScript Code

```
// Create a new collection 'my_collection'
db.createCollection('my_collection').then(function (myColl) {

    // Add a document to insert
    var insert = myColl.add({ name: 'Mats', age: 21 });

    // Add multiple documents to insert
    insert.add([
        { name: 'Lotte', age: 24 },
        { name: 'Vera', age: 39 }
    ]);

    // Add one more document to insert
    var myDoc = {};
    myDoc.name = 'Jamie';
    myDoc.age = 47;
    insert.add(myDoc);

    // run the operation
    return insert.execute();
});
```

C# Code

```
// Create a new collection "my_collection"
var myColl = db.CreateCollection("my_collection");

// Insert JSON data directly
myColl.Add(new { name = "Mats", age = 21 }).Execute();

// Inserting several docs at once
myColl.Add(new[] {new { name = "Lotte", age = 24},
    new { name = "Vera", age = 39} }).Execute();

// Insert Documents using DbDoc
var myDoc = new DbDoc();
myDoc.SetValue("name", "Jamie");
myDoc.SetValue("age", 47);
myColl.Add(myDoc).Execute();

//Fetch all docs
var docResult = myColl.Find().Execute();
var docs = docResult.FetchAll();
```

Python Code

```
# Create a new collection 'my_collection'
my_coll = my_schema.create_collection('my_collection')

# Insert JSON data directly
my_coll.add({'name': 'Mats', 'age': 21}).execute()

# Inserting several docs at once
my_coll.add([
    {'name': 'Lotte', 'age': 24},
    {'name': 'Vera', 'age': 39}
]).execute()
```

Java Code

```
// Create a new collection 'my_collection'
Collection coll = db.createCollection("my_collection");

// Insert JSON data directly
coll.add("{\"name\":\"Mats\", \"age\":21}).execute();

// Insert several documents at once
coll.add("{\"name\":\"Lotte\", \"age\":24}",
    "{\"name\":\"Vera\", \"age\":39}).execute();

// Insert Documents using DbDoc
DbDoc myDoc = new coll.newDoc();

myDoc.add("name", new JsonString().setValue("Jamie"));
myDoc.add("age", new JsonNumber().setValue("47"));
coll.add(myDoc).execute();
```

C++ Code

```
// Create a new collection 'my_collection'
Collection myColl = db.createCollection("my_collection");

// Insert JSON data directly
myColl.add(R"({ "name": "Mats", "age": 21})").execute();

// Inserting several docs at once
std::list<DbDoc> docs = {
    DbDoc(R"({ "name": "Lotte", "age": 24})"),
    DbDoc(R"({ "name": "Vera", "age": 39})")
};
myColl.add(docs).execute();
```

5.2 Working with Document IDs

This section describes what a document ID is and how to work with it.

Every document has a unique identifier called the document ID, which can be thought of as the equivalent of a table's primary key. The document ID value is usually automatically generated by the server when the document is added, but can also be manually assigned. The assigned document ID is returned in the `generatedIds` property of the `Result` (`AddResult` for Connector/J) object for the `collection.add()` operation and can be accessed using the `getGeneratedIds()` method. See [Section 5.3, “Understanding Document IDs”](#) for more background information on document IDs.

The following example in JavaScript code shows adding a document to a collection, retrieving the added document's IDs and testing that duplicate IDs cannot be added.

```
mysql-js > var result = mycollection.add({test:'demo01'}).execute()
mysql-js > print(result.generatedIds)
[
    "00006075f68100000000000000000006"
]
mysql-js > var result = mycollection.add({test:'demo02'}).add({test:'demo03'}).execute()
mysql-js > print(result.generatedIds)
[
    "00006075f68100000000000000000007",
    "00006075f68100000000000000000008"
]

mysql-js > mycollection.find()
{
    "_id": "00006075f68100000000000000000006",
```

```
"test": "demo01"
}
{
  "_id": "00006075f68100000000000000000007",
  "test": "demo02"
}
{
  "_id": "00006075f68100000000000000000008",
  "test": "demo03"
}
3 documents in set (0.0102 sec)
mysql-js > var result = mycollection.add({_id:'00006075f68100000000000000000008', test:'demo04'}).execute()
Document contains a field value that is not unique but required to be (MySQL Error 5116)
```

As shown in the example above, the document ID is stored in the `_id` field of a document. The document ID is a `VARBINARY()` with a maximum length of 32 characters. If an `_id` is provided when a document is created, it is honored; if no `_id` is provided, one is automatically assigned to the document.

The following example illustrates how the `_id` value can either be provided or autogenerated. It is assumed that the `test` schema exists and is assigned to the variable `db`, that the collection `my_collection` exists and that `custom_id` is unique.

MySQL Shell JavaScript Code

```
// If the _id is provided, it will be honored
var result = myColl.add( { _id: 'custom_id', a : 1 } ).execute();
var document = myColl.find("a = 1").execute().fetchOne();
print("User Provided Id:", document._id);

// If the _id is not provided, one will be automatically assigned
result = myColl.add( { b: 2 } ).execute();
print("Autogenerated Id:", result.getGeneratedIds()[0]);
```

MySQL Shell Python Code

```
# If the _id is provided, it will be honored
result = myColl.add( { '_id': 'custom_id', 'a' : 1 } ).execute()
document = myColl.find('a = 1').execute().fetch_one()
print("User Provided Id: %s" % document._id)

# If the _id is not provided, one will be automatically assigned
result = myColl.add( { 'b': 2 } ).execute()
print("Autogenerated Id: %s" % result.get_generated_ids()[0])
```

Node.js JavaScript Code

```
// If the _id is provided, it will be honored
myColl.add({ _id: 'custom_id', a : 1 }).execute().then(function () {
  myColl.getOne('custom_id').then(function (doc) {
    console.log('User Provided Id:', doc._id);
  });
});

// If the _id is not provided, one will be automatically assigned
myColl.add({ b: 2 }).execute().then(function (result) {
  console.log('Autogenerated Id:', result.getGeneratedIds()[0]);
});
```

C# Code

```
// If the _id is provided, it will be honored
var result = myColl.Add(new { _id = "custom_id", a = 1 }).Execute();
Console.WriteLine("User Provided Id:", result.AutoIncrementValue);
```

```
// If the _id is not provided, one will be automatically assigned
result = myColl.Add(new { b = 2 }).Execute();
Console.WriteLine("Autogenerated Id:", result.AutoIncrementValue);
```

Python Code

```
# If the _id is provided, it will be honored
result = my_coll.add({'_id': 'custom_id', 'a': 1}).execute()
print("User Provided Id: {}".format(result.get_last_document_id()))

# If the _id is not provided, one will be automatically assigned
result = my_coll.add({'b': 2}).execute()
print("Autogenerated Id: {}".format(result.get_last_document_id()))
```

Java Code

```
// If the _id is provided, it will be honored
AddResult result = coll.add("{\"_id\":\"custom_id\", \"a\":1}").execute();
System.out.println("User Provided Id:" + ((JsonString) coll.getOne("custom_id").get("_id")).getString());

// If the _id is not provided, one will be automatically assigned
result = coll.add("{\"b\":2}").execute();
System.out.println("Autogenerated Id:" + result.getGeneratedIds().get(0));
```

C++ Code

```
// If the _id is provided, it will be honored
Result result = myColl.add(R"({ "_id": "custom_id", "a" : 1 })").execute();
std::vector<string> ids = result.getGeneratedIds();
if (ids.empty())
    cout << "No Autogenerated Ids" << endl;

// If the _id is not provided, one will be automatically assigned
result = myColl.add(R"({ "b": 2 })").execute();
ids = result.getGeneratedIds();
cout << "Autogenerated Id:" << ids[0] << endl;
```

Some documents have a natural unique key. For example, a collection that holds a list of books is likely to include the International Standard Book Number (ISBN) for each document that represents a book. The ISBN is a string with a length of 13 characters, which is well within the length limit of 32 characters for the `_id` field.

```
// using a book's unique ISBN as the object ID
myColl.add(
    _id: "978-1449374020",
    title: "MySQL Cookbook: Solutions for Database Developers and Administrators"
).execute();
```

Use `find()` to fetch the newly inserted book from the collection by its document ID.

```
var book = myColl.find('_id = "978-1449374020"').execute();
```

Currently, X DevAPI does not support using any document field other than the implicit `_id` as the document ID—there is no way to define another key to perform the same function.

5.3 Understanding Document IDs

This sections describes in detail how document IDs are generated and how to interpret them. X DevAPI relies on server-based document ID generation, which results in sequentially increasing document IDs across all clients. `InnoDB` uses the document ID as a primary key, resulting in efficient page splits and tree reorganizations.

This section describes the properties and format of the automatically generated document IDs.

Document ID Properties

The `_id` field of a document behaves in the same way as any other fields of the document during queries, except that its value cannot be changed once it has been inserted to the collection. The `_id` field is used as the primary key of the collection . It is possible to override the automatic generation of document IDs by manually including an ID in an inserted document.



Important

X Plugin is not aware of the data inserted into the collection, including any manual document IDs you use. When using manual document IDs, you must ensure that they do not clash with any IDs that might ever be generated automatically by the server (see [Document ID Generation](#) for details), in order to avoid any errors due to primary key duplication.

Whenever an `_id` field value is not present in an inserted document, the server generates an `_id` value. The generated `_id` value used for a document is returned to the client as part of the [Result \(Result](#) for Connector/J) object of the `add()` operation. If you are using X DevAPI on an InnoDB Cluster, the automatically generated `_id` must be unique within the whole cluster. By setting the `mysqlx_document_id_unique_prefix` to a unique value per cluster instance, you can ensure document IDs are unique across all the instances.

The `_id` field must be sequential (always incrementing) for optimal InnoDB insertion performance (at least within a single server). The sequential nature of `_id` values is maintained across server restarts.

In a multi-primary Group Replication or InnoDB Cluster environment, the generated `_id` values of a table are unique across instances to avoid primary key conflicts and minimize transaction certification.

Document ID Generation

This section describes how document IDs are formatted.

The format of automatically generated document ID is:

unique_prefix	start_timestamp	serial
4 bytes	8 bytes	16 bytes

Where:

- `unique_prefix` is a value assigned by InnoDB Cluster to the instance, which is used to make the document ID unique across all instances from the same cluster. The range of `unique_prefix` is from 0 to $2^{16}-1$, which is hex encoded. Default value is 0, if it is neither set by InnoDB Cluster nor by the `mysqlx_document_id_unique_prefix` system variable.
- `start_timestamp` is the time stamp of the startup time of the server instance, which is hex encoded. In the unlikely event that the value of `serial` overflows, the `start_timestamp` is incremented by 1 and the `serial` value then restarts at 0.
- `serial` is a per-instance automatically incremented integer serial number value, which is hex encoded and has a range of 0 to $2^{64}-1$. The initial value of `serial` is set to the `auto_increment_offset` system variable, and the increment of the value is set by the `auto_increment_increment` system variable.

This document ID format ensures that:

- The primary key value monotonically increments for inserts originating from a single server instance, although the interval between values is not uniform within a table.
- When using multi-primary Group Replication or InnoDB Cluster, inserts to the same table from different instances do not have conflicting primary key values, as long as the instances have the `auto_increment_offset` and the `auto_increment_increment` system variables configured properly (see descriptions of the variables for details).

Chapter 6 Working with Relational Tables

Table of Contents

6.1 Syntax of the SQL CRUD Functions	84
--------------------------------------------	----

The X DevAPI SQL CRUD functions allow you to work with relational tables in manners similar to using traditional SQL statements. The following code sample shows how to use the `add()` and `select()` methods of the X DevAPI SQL CRUD functions, which are similar to running `INSERT` and `SELECT` statements on a table with an SQL client. Compare this with the examples found in [Section 4.3, “Collection CRUD Function Overview”](#) to see the differences and similarities between the CRUD functions for tables and collections in the X DevAPI.

MySQL Shell JavaScript Code

```
// Working with Relational Tables
var mysqlx = require('mysqlx');

// Connect to server using a connection URL
var mySession = mysqlx.getSession( {
  host: 'localhost', port: 33060,
  user: 'user', password: 'password' } )

var myDb = mySession.getSchema('test');

// Accessing an existing table
var myTable = myDb.getTable('my_table');

// Insert SQL Table data
myTable.insert(['name', 'birthday', 'age']).
  values('Laurie', mysqlx.dateValue(2000, 5, 27), 19).execute();

// Find a row in the SQL Table
var myResult = myTable.select(['_id', 'name', 'birthday']).
  where('name like :name AND age < :age').
  bind('name', 'L%').bind('age', 30).execute();

// Print result
print(myResult.fetchOne());
```

MySQL Shell Python Code

```
# Working with Relational Tables
from mysqlsh import mysqlx

# Connect to server using a connection URL
mySession = mysqlx.get_session( {
  'host': 'localhost', 'port': 33060,
  'user': 'user', 'password': 'password' } )

myDb = mySession.get_schema('test')

# Accessing an existing table
myTable = myDb.get_table('my_table')

# Insert SQL Table data
myTable.insert(['name','birthday','age']) \
  .values('Laurie', mysqlx.date_value(2000, 5, 27), 19).execute()
```

```

# Find a row in the SQL Table
myResult = myTable.select(['_id', 'name', 'birthday']) \
    .where('name like :name AND age < :age') \
    .bind('name', 'L%') \
    .bind('age', 30).execute()

# Print result
print(myResult.fetch_all())

```

Node.js JavaScript Code

```

// Working with Relational Tables
var mysqlx = require('@mysql/xdevapi');
var myTable;

// Connect to server using a connection URL
mysqlx
    .getSession({
        user: 'user',
        password: 'password',
        host: 'localhost',
        port: 33060
    })
    .then(function (session) {
        // Accessing an existing table
        myTable = session.getSchema('test').getTable('my_table');

        // Insert SQL Table data
        return myTable
            .insert(['name', 'birthday', 'age'])
            .values(['Laurie', '2000-5-27', 19])
            .execute()
    })
    .then(function () {
        // Find a row in the SQL Table
        return myTable
            .select(['_id', 'name', 'birthday'])
            .where('name like :name && age < :age')
            .bind('name', 'L%')
            .bind('age', 30)
            .execute();
    })
    .then(function (myResult) {
        console.log(myResult.fetchAll());
    });

```

C# Code

```

// Working with Relational Tables

// Connect to server using a connection
var db = MySQLX.GetSession(
    "server=localhost;port=33060;user=user;password=password")
    .GetSchema("test");

// Accessing an existing table
var myTable = db.GetTable("my_table");

// Insert SQL Table data
myTable.Insert("name", "age")
    .Values("Laurie", "19").Execute();

// Find a row in the SQL Table
var myResult = myTable.Select("_id, name, age")
    .Where("name like :name AND age < :age")

```

```
.Bind(new { name = "L%", age = 30 }).Execute();

// Print result
PrintResult(myResult.FetchAll());
```

Python Code

```
# Working with Relational Tables
import mysqlx

# Connect to server using a connection URL
my_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

my_schema = my_session.get_schema('test')

# Accessing an existing table
my_table = my_schema.get_table('my_table')

# Insert SQL Table data
my_table.insert(['name', 'birthday', 'age']) \
    .values('Laurie', mysqlx.date_value(2000, 5, 27), 19).execute()

# Find a row in the SQL Table
result = my_table.select(['_id', 'name', 'birthday']) \
    .where('name like :name AND age < :age') \
    .bind('name', 'L%') \
    .bind('age', 30).execute()

# Print result
print(result.fetch_all())
```

Java Code

```
// Working with Relational Tables
import com.mysql.cj.xdevapi.*;

// Connect to server using a connection URL
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");
Schema db = mySession.getSchema("test");

// Accessing an existing table
Table myTable = db.getTable("my_table");

// Insert SQL Table data
myTable.insert("name", "birthday", "age").values("Laurie", "2000-05-27", 19).execute();

// Find a row in the SQL Table
RowResult myResult = myTable.select("_id, name, birthday")
    .where("name like :name AND age < :age")
    .bind("name", "L%").bind("age", 30).execute();

// Print result
myResult.forEach(r ->
    System.out.println(r.getString(1) + ":" + r.getDate(2)));
```

C++ Code

```
// Working with Relational Tables
#include <mysqlx/xdevapi.h>

// Connect to server using a connection URL
Session mySession(33060, "user", "password");
```

```

Schema myDb = mySession.getSchema("test");

// Accessing an existing table
Table myTable = myDb.getTable("my_table");

// Insert SQL Table data
myTable.insert("name", "birthday", "age")
    .values("Laurie", "2000-5-27", 19).execute();

// Find a row in the SQL Table
RowResult myResult = myTable.select("_id", "name", "birthday")
    .where("name like :name AND age < :age")
    .bind("name", "L%").bind("age", 30).execute();

// Print result
Row row = myResult.fetchOne();
cout << "    _id: " << row[0] << endl;
cout << "    name: " << row[1] << endl;
cout << "birthday: " << row[2] << endl;

```

6.1 Syntax of the SQL CRUD Functions

The following SQL CRUD functions are available in X DevAPI.

Table.insert()

The `Table.insert()` method works like an `INSERT` statement in SQL. It is used to store data in a relational table in the database. It is executed by the `execute()` function.

The following example shows how to use the `Table.insert() function`. The example assumes that the `test` schema exists and is assigned to the variable `db`, and that an empty table called `my_table` exists.

MySQL Shell JavaScript Code

```

// Accessing an existing table
var myTable = db.getTable('my_table');

// Insert a row of data.
myTable.insert(['id', 'name'])
    .values(1, 'Imani')
    .values(2, 'Adam')
    .execute();

```

MySQL Shell Python Code

```

# Accessing an existing table
myTable = db.get_table('my_table')

# Insert a row of data.
myTable.insert(['id', 'name']).values(1, 'Imani').values(2, 'Adam').execute()

```

Node.js JavaScript Code

```

// Accessing an existing table
var myTable = db.getTable('my_table');

// Insert a row of data.
myTable.insert(['id', 'name'])
    .values(1, 'Imani')
    .values(2, 'Adam')
    .execute();

```

Table.select()

```
execute();
```

C# Code

```
// Assumptions: test schema assigned to db, empty my_table table exists

// Accessing an existing table
var myTable = db.GetTable("my_table");

// Insert a row of data.
myTable.Insert("id", "name")
    .Values(1, "Imani")
    .Values(2, "Adam")
    .Execute();
```

Python Code

```
# Accessing an existing table
my_table = db.get_table('my_table')

# Insert a row of data.
my_table.insert(['id', 'name']).values(1, 'Imani').values(2, 'Adam').execute()
```

Java Code

```
// Accessing an existing table
Table myTable = db.getTable("my_table");

// Insert a row of data.
myTable.insert("id", "name")
    .values(1, "Imani")
    .values(2, "Adam")
    .execute();
```

C++ Code

```
// Accessing an existing table
var myTable = db.getTable("my_table");

// Insert a row of data.
myTable.insert("id", "name")
    .values(1, "Imani")
    .values(2, "Adam")
    .execute();
```

Figure 6.1 Table.insert() Syntax Diagram




Table.select()

The `Table.select()` method works like a `SELECT` statement in SQL. Notice that `Table.select()` and `collection.find()` use different methods for sorting results: `Table.select()` uses the method `orderBy()`, reminiscent of the `ORDER BY` keyword in SQL, while the `sort()` method is used to sort the results returned by `Collection.find()`.

Table.update()

Figure 6.2 Table.select() Syntax Diagram




Table.update()

The `Table.update()` method works like an `UPDATE` statement in SQL.

Figure 6.3 Table.update() Syntax Diagram





Table.delete()

The `Table.delete()` method works like a `DELETE` statement in SQL.

Figure 6.4 Table.delete() Syntax Diagram



Chapter 7 Working with Relational Tables and Documents

Table of Contents

7.1 Collections as Relational Tables	87
--------------------------------------------	----

After seeing how to work with documents and how to work with relational tables, this section explains how to combine the two and work with both at the same time.

It can be beneficial to use documents for very specific tasks inside an application and rely on relational tables for other tasks. Or a very simple document only application can outgrow the document model and incrementally integrate or move to a more powerful relational database. This way the advantages of both documents and relational tables can be combined. SQL tables contribute strictly typed value semantics, predictable and optimized storage. Documents contribute type flexibility, schema flexibility and non-scalar types.

7.1 Collections as Relational Tables

Applications that seek to store standard SQL columns with Documents can cast a collection to a table. In this case a collection can be fetched as a Table object with the `Schema.getCollectionAsTable()` function. From that moment on it is treated as a regular table. Document values can be accessed in SQL CRUD operations using the following syntax:

```
doc->$.field'
```

`doc->$.field'` is used to access the document top level fields. More complex paths can be specified as well.

```
doc->$.some.field.like[3].this'
```

Once a collection has been fetched as a table with the `Schema.getCollectionAsTable()` function, all SQL CRUD operations can be used. Using the syntax for document access, you can select data from the Documents of the Collection and the extra SQL columns.

The following example shows how to insert a JSON document string into the `doc` field.

MySQL Shell JavaScript Code

```
// Get the customers collection as a table
var customers = db.getCollectionAsTable('customers');
customers.insert('doc').values('{"_id":"001", "name": "Ana", "last_name": "Silva"}').execute();

// Now do a find operation to retrieve the inserted document
var result = customers.select(["doc->$.name'", "doc->$.last_name'"]).where("doc->$.id' = '001'").execute();

var record = result.fetchOne();

print ("Name : " + record[0]);
print ("Last Name : " + record[1]);
```

MySQL Shell Python Code

```
# Get the customers collection as a table
customers = db.get_collection_as_table('customers')
customers.insert('doc').values('{"_id":"001", "name": "Ana", "last_name": "Silva"}').execute()

# Now do a find operation to retrieve the inserted document
```

```
result = customers.select(["doc->$.name'", "doc->$.last_name'"]).where("doc->$.id' = '001'").execute()

record = result.fetch_one()

print("Name : %s\n" % record[0])
print("Last Name : %s\n" % record[1])
```

Node.js JavaScript Code

```
// Get the customers collection as a table
var customers = db.getCollectionAsTable('customers');
customers.insert('doc').values({'_id":"001", "name": "Ana"}).execute();
```

C# Code

```
// Get the customers collection as a table
var customers = db.GetCollectionAsTable("customers");
customers.Insert("doc").Values("{ \"_id\": 1, \"name\": \"Ana\" }").Execute();
```

Python Code

```
# Get the customers collection as a table
customers = db.get_collection_as_table("customers")
customers.insert('doc').values({'_id':'001', 'name': 'Ana', 'last_name': 'Silva'}).execute()

# Now do a find operation to retrieve the inserted document
result = customers.select(["doc->$.name'", "doc->$.last_name'"]).where("doc->$.id' = '001'").execute()

record = result.fetch_one()

print('Name : {}'.format(record[0]))
print('Last Name : {}'.format(record[1]))
```

Java Code

```
// Get the customers collection as a table
Table customers = db.getCollectionAsTable("customers");
customers.insert("doc").values("{ \"name\": \"Ana\" }").execute();
```

C++ Code

```
// Get the customers collection as a table
Table customers = db.getCollectionAsTable("customers");
customers.insert("doc")
    .values(R"({"_id":"001", "name": "Ana", "last_name": "Silva"})").execute();

// Now do a find operation to retrieve the inserted document
RowResult result = customers.select("doc->$.name'", "doc->$.last_name' ")
    .where("doc->$.id' = '001'").execute();

Row record = result.fetchOne();
cout << "Name : " << record[0] << endl;
cout << "Last Name : " << record[1] << endl;
```

Chapter 8 Statement Execution

Table of Contents

8.1 Transaction Handling	89
8.1.1 Processing Warnings	92
8.1.2 Error Handling	97
8.2 Working with Savepoints	101
8.3 Working with Locking	102
8.4 Working with Prepared Statements	103

This section explains statement execution, with information on how to handle transactions and errors.

8.1 Transaction Handling

Transactions can be used to group operations into an atomic unit. Either all operations of a transaction succeed when they are committed, or none. It is possible to roll back a transaction as long as it has not been committed.

Transactions can be started in a session using the `startTransaction()` method, committed with `commitTransaction()` and cancelled or rolled back with `rollbackTransaction()`. This is illustrated in the following example. The example assumes that the `test` schema exists and that the collection `my_collection` does not exist.

MySQL Shell JavaScript Code

```
var mysqlx = require('mysqlx');

// Connect to server
var session = mysqlx.getSession( {
  host: 'localhost', port: 33060,
  user: 'user', password: 'password' } );

// Get the Schema test
var db = session.getSchema('test');

// Create a new collection
var myColl = db.createCollection('my_collection');

// Start a transaction
session.startTransaction();
try {
  myColl.add({name: 'Rohit', age: 18, height: 1.76}).execute();
  myColl.add({name: 'Misaki', age: 24, height: 1.65}).execute();
  myColl.add({name: 'Leon', age: 39, height: 1.9}).execute();

  // Commit the transaction if everything went well
  session.commit();

  print('Data inserted successfully.');
}

catch (err) {
  // Rollback the transaction in case of an error
  session.rollback();

  // Printing the error message
  print('Data could not be inserted: ' + err.message);
}
```

MySQL Shell Python Code

```
from mysqlsh import mysqlx

# Connect to server
mySession = mysqlx.get_session( {
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password' } )

# Get the Schema test
myDb = mySession.get_schema('test')

# Create a new collection
myColl = myDb.create_collection('my_collection')

# Start a transaction
mySession.start_transaction()
try:
    myColl.add({'name': 'Rohit', 'age': 18, 'height': 1.76}).execute()
    myColl.add({'name': 'Misaki', 'age': 24, 'height': 1.65}).execute()
    myColl.add({'name': 'Leon', 'age': 39, 'height': 1.9}).execute()
    # Commit the transaction if everything went well
    mySession.commit()
    print('Data inserted successfully.')
except Exception as err:
    # Rollback the transaction in case of an error
    mySession.rollback()

    # Printing the error message
    print('Data could not be inserted: %s' % str(err))
```

C# Code

```
// Connect to server
var session = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password");

// Get the Schema test
var db = session.GetSchema("test");

// Create a new collection
var myColl = db.CreateCollection("my_collection");

// Start a transaction
session.StartTransaction();
try
{
    myColl.Add(new { name = "Rohit", age = 18, height = 1.76}).Execute();
    myColl.Add(new { name = "Misaki", age = 24, height = 1.65}).Execute();
    myColl.Add(new { name = "Leon", age = 39, height = 1.9}).Execute();

    // Commit the transaction if everything went well
    session.Commit();

    Console.WriteLine("Data inserted successfully.");
}
catch(Exception err)
{
    // Rollback the transaction in case of an error
    session.Rollback();

    // Printing the error message
    Console.WriteLine("Data could not be inserted: " + err.Message);
}
```

Python Code

```

import mysqlx

# Connect to server
my_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

# Get the Schema test
my_schema = my_session.get_schema('test')

# Create a new collection
my_coll = my_schema.create_collection('my_collection')

# Start a transaction
session.start_transaction()
try:
    my_coll.add({'name': 'Rohit', 'age': 18, 'height': 1.76}).execute()
    my_coll.add({'name': 'Misaki', 'age': 24, 'height': 1.65}).execute()
    my_coll.add({'name': 'Leon', 'age': 39, 'height': 1.9}).execute()

    # Commit the transaction if everything went well
    my_session.commit()

    print('Data inserted successfully.')
except Exception as err:
    # Rollback the transaction in case of an error
    my_session.rollback()

    # Printing the error message
    print('Data could not be inserted: {}'.format(str(err)))

```

Java Code

```

import com.mysql.cj.xdevapi.*;

// Connect to server
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

Schema db = mySession.getSchema("test");

// Create a new collection
Collection myColl = db.createCollection("my_collection");

// Start a transaction
mySession.startTransaction();
try {
    myColl.add("{\"name\":\"Rohit\", \"age\":18}", "{\"name\":\"Misaki\", \"age\":24}", "{\"name\":\"Leon\"}");

    mySession.commit();
    System.out.println("Data inserted successfully.");
} catch (Exception err) {
    // Rollback the transaction in case of an error
    mySession.rollback();

    // Printing the error message
    System.out.println("Data could not be inserted: " + err.getMessage());
}

```

C++ Code

```

// Connect to server
Session session(SessionOption::HOST, "localhost",
               SessionOption::PORT, 33060,
               SessionOption::USER, "user",
               SessionOption::PWD, "password");

```

```
// Get the Schema test
Schema db = session.getSchema("test");

// Create a new collection
Collection myColl = db.createCollection("my_collection");

// Start a transaction
session.startTransaction();
try {
    myColl.add(R"({"name": "Rohit", "age": 18, "height": 1.76})").execute();
    myColl.add(R"({"name": "Misaki", "age": 24, "height": 1.65})").execute();
    myColl.add(R"({"name": "Leon", "age": 39, "height": 1.9})").execute();

    // Commit the transaction if everything went well
    session.commit();

    cout << "Data inserted successfully." << endl;
}
catch (const Error &err) {
    // Rollback the transaction in case of an error
    session.rollback();

    // Printing the error message
    cout << "Data could not be inserted: " << err << endl;
}
```

8.1.1 Processing Warnings

Similar to the execution of single statements committing or rolling back a transaction can also trigger warnings. To be able to process these warnings the replied result object of `session.commit()` or `Session.rollback()` needs to be checked.

This is shown in the following example. The example assumes that the test schema exists and that the collection `my_collection` does not exist.

MySQL Shell JavaScript Code

```
var mysqlx = require('mysqlx');

// Connect to server
var mySession = mysqlx.getSession( {
    host: 'localhost', port: 33060,
    user: 'user', password: 'password' } );

// Get the Schema test
var myDb = mySession.getSchema('test');

// Create a new collection
var myColl = myDb.createCollection('my_collection');

// Start a transaction
mySession.startTransaction();
try {
    myColl.add({'name': 'Rohit', 'age': 18, 'height': 1.76}).execute();
    myColl.add({'name': 'Misaki', 'age': 24, 'height': 1.65}).execute();
    myColl.add({'name': 'Leon', 'age': 39, 'height': 1.9}).execute();

    // Commit the transaction if everything went well
    var reply = mySession.commit();

    // handle warnings
    if (reply.warningCount){
        var warnings = reply.getWarnings();
    }
}
```

```

        for (index in warnings){
            var warning = warnings[index];
            print ('Type ['+ warning.level + '] (Code ' + warning.code + '): ' + warning.message + '\n');
        }

        print ('Data inserted successfully.');
    }
catch(err)
{
    // Rollback the transaction in case of an error
    reply = mySession.rollback();

    // handle warnings
    if (reply.warningCount){
        var warnings = reply.getWarnings();
        for (index in warnings){
            var warning = warnings[index];
            print ('Type ['+ warning.level + '] (Code ' + warning.code + '): ' + warning.message + '\n');
        }
    }

    // Printing the error message
    print ('Data could not be inserted: ' + err.message);
}

```

MySQL Shell Python Code

```

from mysqlsh import mysqlx

# Connect to server
mySession = mysqlx.get_session( {
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password' } )

# Get the Schema test
myDb = mySession.get_schema('test')

# Create a new collection
myColl = myDb.create_collection('my_collection')

# Start a transaction
mySession.start_transaction()
try:
    myColl.add({'name': 'Rohit', 'age': 18, 'height': 1.76}).execute()
    myColl.add({'name': 'Misaki', 'age': 24, 'height': 1.65}).execute()
    myColl.add({'name': 'Leon', 'age': 39, 'height': 1.9}).execute()

    # Commit the transaction if everything went well
    reply = mySession.commit()

    # handle warnings
    if reply.warning_count:
        for warning in result.get_warnings():
            print('Type [%s] (Code %s): %s\n' % (warning.level, warning.code, warning.message))

    print('Data inserted successfully.')
except Exception as err:
    # Rollback the transaction in case of an error
    reply = mySession.rollback()

    # handle warnings
    if reply.warning_count:
        for warning in result.get_warnings():
            print('Type [%s] (Code %s): %s\n' % (warning.level, warning.code, warning.message))

```

```
# Printing the error message
print('Data could not be inserted: %s' % str(err))
```

C# Code

```
// Connect to server
var session = MySQLX.GetSession("server=localhost;port=33060;user=user;password=password;" );

// Get the Schema test
var db = session.GetSchema("test");

// Create a new collection
var myColl = db.CreateCollection("my_collection");

// Start a transaction
session.StartTransaction();
int warningCount = 0;
try
{
    var result = myColl.Add(new { name = "Rohit", age = 18, height = 1.76 }).Execute();
    warningCount += result.Warnings.Count;
    result = myColl.Add(new { name = "Misaki", age = 24, height = 1.65 }).Execute();
    warningCount += result.Warnings.Count;
    result = myColl.Add(new { name = "Leon", age = 39, height = 1.9 }).Execute();
    warningCount += result.Warnings.Count;

    // Commit the transaction if everything went well
    session.Commit();
    if(warningCount > 0)
    {
        // handle warnings
    }

    Console.WriteLine("Data inserted successfully.");
}
catch (Exception err)
{
    // Rollback the transaction in case of an error
    session.Rollback();
    if(warningCount > 0)
    {
        // handle warnings
    }

    // Printing the error message
    Console.WriteLine("Data could not be inserted: " + err.Message);
}
```

Python Code

```
import mysqlx

# Connect to server
my_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

# Get the Schema test
my_schema = my_session.get_schema('test')

# Create a new collection
my_coll = my_schema.create_collection('my_collection')

# Start a transaction
my_session.start_transaction()
```

```

try:
    my_coll.add({'name': 'Rohit', 'age': 18, 'height': 1.76}).execute()
    my_coll.add({'name': 'Misaki', 'age': 24, 'height': 1.65}).execute()
    my_coll.add({'name': 'Leon', 'age': 39, 'height': 1.9}).execute()

    # Commit the transaction if everything went well
    result = my_session.commit()

    # handle warnings
    if result.get_warnings_count() > 0:
        for warning in result.get_warnings():
            print('Type [{0}] (Code {1}): {2}'.format(warning['level'], warning['code'], warning['msg']))

    print('Data inserted successfully.')
except Exception as err:
    # Rollback the transaction in case of an error
    result = my_session.rollback()

    # handle warnings
    if result.get_warnings_count() > 0:
        for warning in result.get_warnings():
            print('Type [{0}] (Code {1}): {2}'.format(warning['level'], warning['code'], warning['msg']))

    # Printing the error message
    print('Data could not be inserted: {}'.format(err))

```

Java Code

```
// c.f. "standard transaction handling"
```

C++ Code

```
/*
 * Connector/C++ does not yet provide access to transaction warnings
 * -- Session methods commit() and rollback() do not return a result object.
 */
```

By default all warnings are sent from the server to the client. If an operation is known to generate many warnings and the warnings are of no value to the application then sending the warnings can be suppressed. This helps to save bandwidth. `session.setFetchWarnings()` controls whether warnings are discarded at the server or are sent to the client. `session.getFetchWarnings()` is used to learn the currently active setting.

MySQL Shell JavaScript Code

```

var mysqlx = require('mysqlx');

function process_warnings(result){
    if (result.getWarningCount()){
        var warnings = result.getWarnings();
        for (index in warnings){
            var warning = warnings[index];
            print ('Type [' + warning.level + '] (Code ' + warning.code + '): ' + warning.message + '\n');
        }
    } else{
        print ("No warnings were returned.\n");
    }
}

// Connect to server
var mySession = mysqlx.getSession( {
    host: 'localhost', port: 33060,
    user: 'user', password: 'password' } );

```

```
// Disables warning generation
mySession.setFetchWarnings(false);
var result = mySession.sql('drop schema if exists unexisting').execute();
process_warnings(result);

// Enables warning generation
mySession.setFetchWarnings(true);
var result = mySession.sql('drop schema if exists unexisting').execute();
process_warnings(result);
```

MySQL Shell Python Code

```
from mysqlsh import mysqlx

def process_warnings(result):
    if result.get_warnings_count():
        for warning in result.get_warnings():
            print('Type [%s] (Code %s): %s\n' % (warning.level, warning.code, warning.message))
    else:
        print("No warnings were returned.\n")

    # Connect to server
    mySession = mysqlx.get_session( {
        'host': 'localhost', 'port': 33060,
        'user': 'user', 'password': 'password' } );

    # Disables warning generation
    mySession.set_fetch_warnings(False)
    result = mySession.sql('drop schema if exists unexisting').execute()
    process_warnings(result)

    # Enables warning generation
    mySession.set_fetch_warnings(True)
    result = mySession.sql('drop schema if exists unexisting').execute()
    process_warnings(result)
```

Java Code

```
// Connect to server
Session mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

Schema db = mySession.getSchema("test");

// Create a new collection
Collection myColl = db.createCollection("my_collection");

// Start a transaction
mySession.startTransaction();
try {
    Result res = myColl.add("{"name": "Rohit", "age": 18}, {"name": "Misaki", "age": 24}, {"name": "John", "age": 30});
    System.out.println(res.getWarningsCount());

    Iterator<Warning> warnings = res.getWarnings();
    while (warnings.hasNext()) {
        Warning warn = warnings.next();
        System.out.println(warn.getCode() + ", " + warn.getLevel() + ", " + warn.getMessage());
    }
}

mySession.commit();
System.out.println("Data inserted successfully.");
} catch (Exception err) {
    // Rollback the transaction in case of an error
    mySession.rollback();
```

```
// Printing the error message
System.out.println("Data could not be inserted: " + err.getMessage());
}
```

8.1.2 Error Handling

When writing scripts for MySQL Shell you can often simply rely on the exception handling done by MySQL Shell. For all other languages either proper exception handling is required to catch errors or the traditional error handling pattern needs to be used if the language does not support exceptions.

The default error handling can be changed by creating a custom [SessionContext](#) and passing it to the `mysqlx.getSession()` function. This enables switching from exceptions to result based error checking.

The following example shows how to perform proper error handling. The example assumes that the test schema exists and that the collection `my_collection` exists.

MySQL Shell JavaScript Code

```
var mysqlx = require('mysqlx');

var mySession;

try {
    // Connect to server on localhost
    mySession = mysqlx.getSession( {
        host: 'localhost', port: 33060,
        user: 'user', password: 'password' } );
}
catch (err) {
    print('The database session could not be opened: ' + err.message);
}

try {
    var myDb = mySession.getSchema('test');

    // Use the collection 'my_collection'
    var myColl = myDb.getCollection('my_collection');

    // Find a document
    var myDoc = myColl.find('name like :param').limit(1)
        .bind('param','L%').execute();

    // Print document
    print(myDoc.first());
}
catch (err) {
    print('The following error occurred: ' + err.message);
}
finally {
    // Close the session in any case
    mySession.close();
}
```

MySQL Shell Python Code

```
from mysqlsh import mysqlx

mySession

try:
    # Connect to server on localhost
    mySession = mysqlx.get_session( {
```

```
'host': 'localhost', 'port': 33060,
'username': 'user', 'password': 'password' } )}

except Exception as err:
    print('The database session could not be opened: %s' % str(err))

try:
    myDb = mySession.get_schema('test')

    # Use the collection 'my_collection'
    myColl = myDb.get_collection('my_collection')

    # Find a document
    myDoc = myColl.find('name like :param').limit(1).bind('param','L%').execute()

    # Print document
    print(myDoc.first())
except Exception as err:
    print('The following error occurred: %s' % str(err))
finally:
    # Close the session in any case
    mySession.close()
```

Node.js JavaScript Code

```
var mysqlx = require('@mysql/xdevapi');

// Connect to server on localhost
mysqlx
    .getSession({
        host: 'localhost',
        port: 33060,
        user: 'user',
        password: 'password'
    })
    .then(function (mySession) {
        // This can't throw an error as we check existence at a later operation only
        var myDb = mySession.getSchema('test');

        // Use the collection 'my_collection'
        // This can't throw an error as we check existence at a later operation only
        var myColl = myDb.getCollection('my_collection');

        // Find a document
        return myColl
            .find('name like :param')
            .limit(1)
            .bind('param','L%')
            .execute(function (row) {
                console.log(row);
            })
            .then(function () {
                return session.close();
            })
            .catch(function (err) {
                console.log('The following error occurred: ' + err.message);
            });
    })
    .catch (err) {
        console.log('The database session could not be opened: ' + err.message);
    });
}
```

C# Code

```
Session mySession = null;
try
```

```
{
    // Connect to server on localhost
    mySession = MySQLX.GetSession("mysqlx://user:password@localhost:33060");

    try
    {
        Schema myDb = mySession.GetSchema("test");

        // Use the collection 'my_collection'
        Collection myColl = myDb.GetCollection("my_collection");

        // Find a document
        DocResult myDoc = myColl.Find("name like :param").Limit(1).Bind("param", "L%").Execute();

        // Print document
        Console.WriteLine(myDoc.FetchOne());
    }
    catch (Exception err)
    {
        Console.WriteLine("The following error occurred: " + err.Message);
    }
    finally
    {
        // Close the session in any case
        mySession.Close();
    }
}
catch (Exception err)
{
    Console.WriteLine("The database session could not be opened: " + err.Message);
}
```

Python Code

```
import mysqlx

# Connect to server
my_session = mysqlx.get_session({
    'host': 'localhost', 'port': 33060,
    'user': 'user', 'password': 'password'
})

# Get the Schema test
my_schema = my_session.get_schema('test')

# Create a new collection
my_coll = my_schema.create_collection('my_collection')

# Start a transaction
my_session.start_transaction()
try:
    my_coll.add({'name': 'Rohit', 'age': 18, 'height': 1.76}).execute()
    my_coll.add({'name': 'Misaki', 'age': 24, 'height': 1.65}).execute()
    my_coll.add({'name': 'Leon', 'age': 39, 'height': 1.9}).execute()

    # Commit the transaction if everything went well
    result = my_session.commit()

    # handle warnings
    if result.get_warnings_count() > 0:
        for warning in result.get_warnings():
            print('Type [{0}] (Code {1}): {2}'.format(warning['level'], warning['code'], warning['msg']))

    print('Data inserted successfully.')
except Exception as err:
    # Rollback the transaction in case of an error
    my_session.rollback()
```

```
# handle warnings
if reply.get_warnings_count() > 0:
    for warning in result.get_warnings():
        print('Type [{0}] (Code {1}): {2}'.format(warning['level'], warning['code'], warning['msg']))

# Printing the error message
print('Data could not be inserted: {0}'.format(err))
```

Java Code

```
import com.mysql.cj.xdevapi.*;

Session mySession;

try {
    // Connect to server on localhost
    mySession = new SessionFactory().getSession("mysqlx://localhost:33060/test?user=user&password=password");

    try {
        Schema myDb = mySession.getSchema("test");

        // Use the collection 'my_collection'
        Collection myColl = myDb.getCollection("my_collection");

        // Find a document
        DocResult myDoc = myColl.find("name like :param").limit(1).bind("param", "L%").execute();

        // Print document
        System.out.println(myDoc.fetchOne());
    } catch (XDevAPIError err) { // special exception class for server errors
        System.err.println("The following error occurred: " + err.getMessage());
    } finally {
        // Close the session in any case
        mySession.close();
    }
} catch (Exception err) {
    System.err.println("The database session could not be opened: " + err.getMessage());
}
```

C++ Code

```
#include <mysqlx/xdevapi.h>

try
{
    // Connect to server on localhost
    Session session(33060, "user", "password");

    try
    {
        Schema db = session.getSchema("test");

        // Use the collection 'my_collection'
        Collection myColl = db.getCollection("my_collection");

        // Find a document
        auto myDoc = myColl.find("name like :param").limit(1)
                    .bind("param", "L%").execute();

        // Print document
        cout << myDoc.fetchOne() << endl;
    }

    // Exit with success code
    exit(0);
}
```

```
catch (const Error &err)
{
    cout << "The following error occurred: " << err << endl;
    exit(1);
}

// Note: session is closed automatically when session object
// is destructed.
}
catch (const Error &err)
{
    cout << "The database session could not be opened: " << err << endl;

    // Exit with error code
    exit(1);
}
```

8.2 Working with Savepoints

X DevAPI supports savepoints, which enable you to set a named point within a transaction that you can revert to. By setting savepoints within a transaction, you can later use the rollback functionality to undo any statements issued after setting the savepoint. Savepoints can be released if you no longer require them. This section documents how to work with savepoints in X DevAPI. See [SAVEPOINT](#) for background information.

Setting a Savepoint

Savepoints are identified by a string name. The string can contain any character allowed for an identifier. To create a savepoint, use the `session.setSavepoint()` operation, which maps to the SQL statement `SAVEPOINT name`. If you do not specify a `name`, one is automatically generated. For example by issuing:

```
session.setSavepoint()
```

a transaction savepoint is created with an automatically generated name and a string is returned with the name of the savepoint. This name can be used with the `session.rollbackTo()` or `session.releaseSavepoint()` operations. The `session.setSavepoint()` operation can be called multiple times within a session and each time a unique savepoint name is generated.

It is also possible to manually define the name of the savepoint by passing in a string `name`. For example issuing:

```
session.setSavepoint('name')
```

results in a transaction savepoint with the specified `name`, which is returned by the operation as a string. The `session.setSavepoint('name')` operation can be called multiple times in this way, and if the `name` has already been used for a savepoint then the previous savepoint is deleted and a new one is set.

Rolling Back to a Savepoint

When a session has transaction savepoints, you can undo any subsequent transactions using the `session.rollbackTo()` operation, which maps to the `ROLLBACK TO name` statement. For example, issuing:

```
session.rollbackTo('name')
```

rolls back to the transaction savepoint `name`. This operation succeeds as long as the given savepoint has not been released. Rolling back to a savepoint which was created prior to other savepoints results in the subsequent savepoints being either released or rolled back. For example:

```

session.beginTransaction()
(some data modifications occur...)

session.setSavepoint('point1')      <---- succeeds
(some data modifications occur...)

session.setSavepoint('point2')      <---- succeeds
(some data modifications occur...)

session.rollbackTo('point1')        <---- succeeds
session.rollbackTo('point1')        <---- still succeeds, but position stays the same
session.rollbackTo('point2')        <---- generates an error because lines above already cleared point2
session.rollbackTo('point1')        <---- still succeeds

```

Releasing a Savepoint

To cancel a savepoint, for example when it is no longer needed, use `releaseSavepoint()` and pass in the name of the savepoint you want to release. For example, issuing:

```
session.releaseSavepoint('name')
```

releases the savepoint `name`.

Savepoints and Implicit Transaction Behavior

The exact behavior of savepoints is defined by the server, and specifically how autocommit is configured. See [autocommit, Commit, and Rollback](#).

For example, consider the following statements with no explicit `BEGIN`, `session.beginTransaction()` or similar call:

```

session.setSavepoint('testsavepoint');
session.releaseSavepoint('testsavepoint');

```

If autocommit mode is enabled on the server, these statements result in an error because the savepoint named `testsavepoint` does not exist. This is because the call to `session.setSavepoint()` creates a transaction, then the savepoint and directly commits it. The result is that savepoint does not exist by the time the call to `releaseSavepoint()` is issued, which is instead in its own transaction. In this case, for the savepoint to survive you need to start an explicit transaction block first.

8.3 Working with Locking

X DevAPI supports MySQL locking through the `lockShared()` and `lockExclusive()` methods for the `Collection.find()` and `Table.select()` methods. This enables you to control row locking to ensure safe, transactional document updates on collections and to avoid concurrency problems, for example when using the `modify()` method. This section describes how to use the `lockShared()` and `lockExclusive()` methods for both the `Collection.find()` and `Table.select()` methods. For more background information on locking, see [Locking Reads](#).

The `lockShared()` and `lockExclusive()` methods have the following properties, whether they are used with a `Collection` or a `Table`.

- Multiple calls to the lock methods are permitted. If a locking statement executes while a different transaction holds the same lock, it blocks until the other transaction releases it. If multiple calls to the lock methods are made, the last called lock method takes precedence. In other words `find().lockShared().lockExclusive()` is equivalent to `find().lockExclusive()`.
- `lockShared()` has the same semantics as `SELECT ... LOCK IN SHARE MODE`. Sets a shared mode lock on any rows that are read. Other sessions can read the rows, but cannot modify them until

your transaction commits. If any of these rows were changed by another transaction that has not yet committed, your query waits until that transaction ends and then uses the latest values.

- `lockExclusive()` has the same semantics as `SELECT ... FOR UPDATE`. For any index records the search encounters, it locks the rows and any associated index entries, in the same way as if you issued an `UPDATE` statement for those rows. Other transactions are blocked from updating those rows, from doing `SELECT ... LOCK IN SHARE MODE`, or from reading the data in certain transaction isolation levels. Consistent reads ignore any locks set on the records that exist in the read view. Old versions of a record cannot be locked; they are reconstructed by applying undo logs on an in-memory copy of the record.
- Locks are held for as long as the transaction which they were acquired in exists. They are immediately released after the statement finishes unless a transaction is open or autocommit mode is turned off.

Both locking methods support the `NOWAIT` and `SKIP LOCKED` InnoDB locking modes. For more information see [Locking Read Concurrency with NOWAIT and SKIP LOCKED](#). To use these locking modes with the locking methods, pass in one of the following:

- `NOWAIT` - if the function encounters a row lock it aborts and generates an `ER_LOCK_NOWAIT` error
- `SKIP_LOCKED` - if the function encounters a row lock it skips the row and continues
- `DEFAULT` - if the function encounters a row lock it waits until there is no lock. The equivalent of calling the lock method without a mode.

Locking considerations

When working with locking modes note the following:

- `autocommit` mode means that there is always a transaction open, which is committed automatically when an SQL statement executes.
- By default sessions are in autocommit mode.
- You disable autocommit mode implicitly when you call `startTransaction()`.
- When in autocommit mode, if a lock is acquired, it is released after the statement finishes. This could lead you to conclude that the locks were not acquired, but that is not the case.
- Similarly, if you try to acquire a lock that is already owned by someone else, the statement blocks until the other lock is released.

8.4 Working with Prepared Statements

X DevAPI improves performance for each CRUD statement that is executed repeatedly by using a server-side prepared statement for its second and subsequent executions. This happens internally—applications do not need to do anything extra to utilize the feature, as long as the same operation object is reused.

When a statement is executed for a second time with changes only in data values or in values that refine the execution results (for example, different `offset()` or `limit()` values), the server prepares the statement for subsequent executions, so that there is no need to reparse the statement when it is being run again. New values for re-executions of the prepared statement are provided with parameter binding. When the statement is modified by chaining to it a method that refines the result (for example, `sort()`, `limit()`, or `offset()`), the statement is reprepared. The following pseudocode and the comments on them demonstrate the feature:

```
var f = coll.find("field = :field");
```

```
f.bind("field", 1).execute(); // Normal execution
f.bind("field", 2).execute(); // Same statement executed with a different parameter value triggers statement preparation
f.bind("field", 3).execute(); // Prepared statement executed with a new value
f.bind("field", 3).limit(10).execute(); // Statement reprepared as it is modified with limit()
f.bind("field", 4).limit(20).execute(); // Reprepared statement executed with new parameters
```

Notice that to take advantage of the feature, the same operation object must be reused in the repetitions of the statement. Look at this example:

```
for (i=0; i<100; ++i) {
    coll.find("field = :field").bind("field", i).execute();
}
```

This loop cannot take advantage of the prepared statement feature, because the operation object of `coll.find()` is recreated at each iteration of the `for` loop. Now, look at this example:

```
var f = coll.find("field = :field");

for (i=0; i<100; ++i) {
    f.bind("field", i).execute();
}
```

The repeated statement is prepared once and then reused, as the same operation of `coll.find()` is re-executed for each iteration of the `for` loop.

Prepared statements are part of a `Session`. When a `Client` resets the `Session` (by using, for example, `Mysqlx.Session.Reset`), the prepared statements are dropped.

Chapter 9 Working with Result Sets

Table of Contents

9.1 Result Set Classes	105
9.2 Working with <code>AUTO-INCREMENT</code> Values	106
9.3 Working with Data Sets	106
9.4 Fetching All Data Items at Once	110
9.5 Working with SQL Result Sets	112
9.6 Working with Metadata	122
9.7 Support for Language Native Iterators	123

This section explains how to work with result sets returned by database operations.


9.1 Result Set Classes

All database operations return a result. The type of result returned depends on the operation that was executed. The different types of results are outlined in the following table.

Table 9.1 The Result Classes and the Information They Provide

Result Class	Returned By	Provides
<code>Result</code>	<code>add().execute()</code> , <code>insert().execute()</code> , <code>modify().execute()</code> , <code>update().execute()</code> , <code>remove().execute()</code> , <code>delete().execute()</code>	Number of rows affected by the operation, auto generated document IDs, last auto-generated <code>AUTO INCREMENT</code> column values, or warnings, depending on the operation for which <code>Result</code> is returned.
<code>SqlResult</code>	<code>session.sql().execute()</code>	Number of rows affected by the operation, auto generated document IDs, last auto-generated <code>AUTO INCREMENT</code> column values, warnings, or fetched data set, depending on the operation for which <code>SqlResult</code> is returned.
<code>DocResult</code>	<code>find().execute()</code>	The fetched data set
<code>RowResult</code>	<code>select.execute()</code>	The fetched data set

The following class diagram gives a basic overview of the result classes and their functions.

Figure 9.1 Result Classes

9.2 Working with AUTO-INCREMENT Values

`AUTO_INCREMENT` columns can be used in MySQL for generating primary key or `id` values, but are not limited to these uses. This section explains how to retrieve `AUTO_INCREMENT` values when adding rows using X DevAPI. For more background information, see [Using AUTO_INCREMENT](#).

X DevAPI provides the `getAutoIncrementValue()` method to return the first `AUTO_INCREMENT` column value that was successfully inserted by the operation, taken from the return value of `table.insert()`. In the following example it is assumed that the table contains a `PRIMARY KEY` column for which the `AUTO_INCREMENT` attribute is set:

```
res = myTable.insert(['name']).values('Mats').values('Otto').execute();
print(res.getAutoIncrementValue());
```

This `table.insert()` operation inserted multiple rows. `getAutoIncrementValue()` returns the `AUTO_INCREMENT` column value generated for the first inserted row only, so in this example, for the row containing “Mats”. The reason for this is to make it possible to reproduce easily the same operation against some other server.

9.3 Working with Data Sets

Operations that fetch data items return a cursor that can be used to consume those data items from the result set. Data items can be read from the database using `Collection.find()`, `Table.select()` and `Session.sql().Collection.find()` returns a data set with documents and `Table.select()` respectively `Session.sql()` return a data set with rows.

All result sets implement a unified way of iterating their data items. The unified syntax supports fetching items one by one using `fetchOne()` or retrieving a list of all items using `fetchAll()`. `fetchOne()` and

`fetchAll()` follow forward-only iteration semantics. Connectors implementing the X DevAPI can offer more advanced iteration patterns on top to match common native language patterns.

The following example shows how to access the documents returned by a `Collection.find()` operation by using `fetchOne()` to loop over all documents.

The first call to `fetchOne()` returns the first document found. All subsequent calls increment the internal data item iterator cursor by one position and return the item found making the second call to `fetchOne()` return the second document found, if any. When the last data item has been read and `fetchOne()` is called again, a `NULL` value is returned. This ensures that the basic while loop shown works with all languages that support such an implementation.

When using `fetchOne()`, it is not possible to reset the internal data item cursor to the first data item to start reading the data items again. A data item (here a Document) that has been fetched once using `fetchOne()` can be discarded by the Connector. The data item's life time is decoupled from the data set. From a Connector perspective items are consumed by the caller as they are fetched. This example assumes that the test schema exists.

MySQL Shell JavaScript Code

```
var myColl = db.getCollection('my_collection');

var res = myColl.find('name like :name').bind('name', 'L%').
    execute();

var doc;
while (doc = res.fetchOne()) {
    print(doc);
}
```

MySQL Shell Python Code

```
myColl = db.get_collection('my_collection')

res = myColl.find('name like :name').bind('name', 'L%').execute()

doc = res.fetch_one()
while doc:
    print(doc)
    doc = res.fetch_one()
```

C# Code

```
var myColl = db.GetCollection("my_collection");

var res = myColl.Find("name like :name").Bind("name", "L%")
    .Execute();

DbDoc doc;
while ((doc = res.FetchOne()) != null)
{
    Console.WriteLine(doc);
}
```

Python Code

```
my_coll = db.get_collection('my_collection')

res = my_coll.find('name like :name').bind('name', 'L%').execute()

doc = res.fetch_one()
```

```
while doc:
    print(doc)
    doc = res.fetch_one()
```

Java Code

```
Collection myColl = db.getCollection("my_collection");

DocResult res = myColl.find("name like :name").bind("name", "L%")
    .execute();

DbDoc doc;
while ((doc = res.fetchOne()) != null) {
    System.out.println(doc);
}
```

C++ Code

```
Collection myColl = db.getCollection("my_collection");

DocResult res = myColl.find("name like :name").bind("name", "L%").execute();
DbDoc doc;
while ((doc = res.fetchOne()))
{
    cout << doc << endl;
}
```

Node.js JavaScript Code

```
var myColl = db.getCollection('my_collection');

myColl.find('name like :name')
    .bind('name', 'L%')
    .execute()
    .then(res => {
        while (doc = res.fetchOne()) {
            console.log(doc);
        }
    });

myColl.find('name like :name')
    .bind('name', 'L%')
    .execute(function (doc) {
        console.log(doc);
    });
});
```

When using Node.js, results can also be returned to a callback function, which is passed to `execute()` in an asynchronous manner whenever results from the server arrive.

```
var myColl = db.getCollection('my_collection');

myColl.find('name like :name')
    .bind('name', 'L%')
    .execute(function (doc) {
        console.log(doc);
    });
});
```

The following example shows how to directly access the rows returned by a `Table.select()` operation. The basic code pattern for result iteration is the same. The difference between the following and the previous example is in the data item handling. Here, `fetchOne()` returns Rows. The exact syntax to access the column values of a Row is language dependent. Implementations seek to provide a language native access pattern. The example assumes that the `test` schema exists and that the employee table exists in `myTable`.

MySQL Shell JavaScript Code

```
var myRows = myTable.select(['name', 'age']).  
    where('name like :name').bind('name', 'L%').  
    execute();  
  
var row;  
while (row = myRows.fetchOne()) {  
    // Accessing the fields by array  
    print('Name: ' + row['name'] + '\n');  
  
    // Accessing the fields by dynamic attribute  
    print(' Age: ' + row.age + '\n');  
}
```

MySQL Shell Python Code

```
myRows = myTable.select(['name', 'age']).where('name like :name').bind('name', 'L%').execute()  
  
row = myRows.fetch_one()  
while row:  
    # Accessing the fields by array  
    print('Name: %s\n' % row[0])  
    # Accessing the fields by dynamic attribute  
    print('Age: %s\n' % row.age)  
    row = myRows.fetch_one()
```

Node.js JavaScript Code

```
var myRows = myTable  
.select(['name', 'age'])  
.where('name like :name')  
.bind('name', 'L%')  
.execute(function (row) {  
    // Connector/Node.js does not support referring to row columns by their name yet.  
    // One needs to access fields by their array index.  
    console.log('Name: ' + row[0]);  
    console.log(' Age: ' + row[1]);  
});
```

Alternatively, you can use callbacks:

```
myTable.select(['name', 'age'])  
.where('name like :name')  
.bind('name', 'L%')  
.execute()  
.then(myRows => {  
    while (var row = myRows.fetchOne()) {  
        // Accessing the fields by array  
        console.log('Name: ' + row[0] + '\n');  
        console.log('Age: ' + row[1] + '\n');  
    }  
});
```

C# Code

```
var myRows = myTable.Select("name", "age")  
.Where("name like :name").Bind("name", "L%")  
.Execute();  
  
Row row;  
while ((row = myRows.FetchOne()) != null)  
{  
    // Accessing the fields by array  
    Console.WriteLine("Name: " + row[0]);
```

```
// Accessing the fields by name
Console.WriteLine("Age: " + row["age"]);
}
```

Python Code

```
rows = my_table.select(['name', 'age']).where('name like :name').bind('name', 'L%').execute()

row = rows.fetch_one()
while row:
    # Accessing the fields by array
    print('Name: {}'.format(row[0]))

    # Accessing the fields by dynamic attribute
    print('Age: {}'.format(row['age']))

    row = rows.fetch_one()
```

Java Code

```
RowResult myRows = myTable.select("name, age")
    .where("name like :name").bind("name", "L%")
    .execute();

Row row;
while ((row = myRows.fetchOne()) != null) {
    // Accessing the fields
    System.out.println(" Age: " + row.getInt("age") + "\n");
}
```

C++ Code

```
RowResult myRows = myTable.select("name", "age")
    .where("name like :name")
    .bind("name", "L%")
    .execute();

Row row;
while ((row = myRows.fetchOne()))
{
    // Connector/C++ does not support referring to row columns by their name yet.
    cout << "Name: " << row[0] << endl;
    cout << " Age: " << row[1] << endl;
    int age = row[1];
    // One needs explicit .get<int>() as otherwise operator<>() is ambiguous
    bool uforty = row[age].get<int>() < 40;
    // Alternative formulation
    bool uforty = (int)row[age] < 40;
}
```

9.4 Fetching All Data Items at Once

In addition to the pattern of using `fetchOne()` explained at [Section 9.3, “Working with Data Sets”](#), which enables applications to consume data items one by one, X DevAPI also provides a pattern using `fetchAll()`, which passes all data items of a data set as a list to the application. The different X DevAPI implementations use appropriate data types for their programming language for the list. Because different data types are used, the language's native constructs are supported to access the list elements. The following example assumes that the `test` schema exists and that the `employee` table exists in `myTable`.

MySQL Shell JavaScript Code

```
var myResult = myTable.select(['name', 'age'])
    .where('name like :name').bind('name', 'L%').
```

```

execute();

var myRows = myResult.fetchAll();

for (index in myRows){
    print (myRows[index].name + " is " + myRows[index].age + " years old.");
}

```

MySQL Shell Python Code

```

myResult = myTable.select(['name', 'age']) \
    .where('name like :name').bind('name', 'L%') \
    .execute()

myRows = myResult.fetch_all()

for row in myRows:
    print("%s is %s years old." % (row.name, row.age))

```

Node.js JavaScript Code

```

myTable.select(['name', 'age'])
    .where('name like :name')
    .bind('name', 'L%')
    .execute()
    .then(myResult => {
        var myRows = myResult.fetchAll();

        myRows.forEach(row => {
            console.log(` ${row[0]} is ${row[1]} years old.`);
        });
    });

```

C# Code

```

var myRows = myTable.Select("name", "age")
    .Where("name like :name").Bind("name", "L%")
    .Execute();
var rows = myRows.FetchAll();

```

Python Code

```

result = myTable.select(['name', 'age']) \
    .where('name like :name').bind('name', 'L%') \
    .execute()

rows = result.fetch_all()

for row in rows:
    print("{0} is {1} years old.".format(row["name"], row["age"]))

```

Java Code

```

RowResult myRows = myTable.select("name, age")
    .where("name like :name").bind("name", "L%")
    .execute();

List<Row> rows = myRows.fetchAll();
for (Row row : rows) {
    // Accessing the fields
    System.out.println(" Age: " + row.getInt("age") + "\n");
}

```

C++ Code

```
RowResult myRows = myTable.select("name, age")
    .where("name like :name")
    .bind("name", "L%")
    .execute();

std::list<Row> rows = myRows.fetchAll();
for (Row row : rows)
{
    cout << row[1] << endl;
}

// Directly iterate over rows, without storing them in a container

for (Row row : myRows.fetchAll())
{
    cout << row[1] << endl;
}
```

When mixing `fetchOne()` and `fetchAll()` to read from one data set keep in mind that every call to `fetchOne()` or `fetchAll()` consumes the data items returned. Items consumed cannot be requested again. If, for example, an application calls `fetchOne()` to fetch the first data item of a data set, then a subsequent call to `fetchAll()` returns the second to last data item. The first item is not part of the list of data items returned by `fetchAll()`. Similarly, when calling `fetchAll()` again for a data set after calling it previously, the second call returns an empty collection.

The use of `fetchAll()` forces a Connector to build a list of all items in memory before the list as a whole can be passed to the application. The life time of the list is independent from the life of the data set that has produced it.

9.5 Working with SQL Result Sets

When you execute an SQL operation on a Session using the `sql()` method, an `SqlResult` is returned. Iterating over an `SqlResult` is identical to working with results from CRUD operations. The following example assumes that the `users` table exists.

MySQL Shell JavaScript Code

```
var res = mySession.sql('SELECT name, age FROM users').execute();

var row;
while (row = res.fetchOne()) {
    print('Name: ' + row['name'] + '\n');
    print('Age: ' + row.age + '\n');
}
```

MySQL Shell Python Code

```
res = mySession.sql('SELECT name, age FROM users').execute()

row = res.fetch_one()

while row:
    print('Name: %s\n' % row[0])
    print('Age: %s\n' % row.age)
    row = res.fetch_one()
```

Node.js JavaScript Code

```
mySession.sql('SELECT name, age FROM users')
    .execute()
    .then(res => {
```

```
        while (row = res.fetchOne()) {
            console.log('Name: ' + row[0] + '\n');
            console.log(' Age: ' + row[1] + '\n');
        }
    });
}
```

Alternatively, you can use callbacks:

```
mySession.sql('SELECT name, age FROM users')
    .execute(function (row) {
        console.log('Name: ' + row[0] + '\n');
        console.log(' Age: ' + row[1] + '\n');
    });
}
```

C# Code

```
var res = Session.SQL("SELECT name, age FROM users").Execute();

while (res.Next())
{
    Console.WriteLine("Name: " + res.Current["name"]);
    Console.WriteLine("Age: " + res.Current["age"]);
}
```

Python Code

```
# Connector/Python
res = mySession.sql('SELECT name, age FROM users').execute()

row = res.fetch_one()

while row:
    print('Name: %s\n' % row[0])
    print(' Age: %s\n' % row.age)
    row = res.fetch_one()
```

Java Code

```
SqlResult res = mySession.sql("SELECT name, age FROM users").execute();

Row row;
while ((row = res.fetchOne()) != null) {
    System.out.println(" Name: " + row.getString("name") + "\n");
    System.out.println(" Age: " + row.getInt("age") + "\n");
}
```

C++ Code

```
SqlResult res = mysession.sql("SELECT name, age FROM users").execute();

Row row;
while ((row = res.fetchOne())) {
    cout << "Name: " << row[0] << endl;
    cout << " Age: " << row[1] << endl;
}
```

`SqlResult` differs from results returned by CRUD operations in the way how result sets and data sets are represented. An `SqlResult` combines a result set produced by, for example, `INSERT`, and a data set, produced by, for example, `SELECT` in one. Unlike with CRUD operations, there is no distinction between the two types for `SqlResult`. An `SqlResult` instance exports methods for accessing data and to retrieving the last inserted ID or number of affected rows.

Use the `hasData()` method to learn whether an `SqlResult` is a data set or a result. The method is useful when code is to be written that has no knowledge about the origin of an `SqlResult`. This can

be the case when writing a generic application function to print query results or when processing stored procedure results. If `hasData()` returns `true`, then the `SqlResult` originates from a `SELECT` or similar command that can return rows.

A return value of `true` does not indicate whether the data set contains any rows. The data set can be empty if, for example, `fetchOne()` returns `NUL` or `fetchAll()` returns an empty list. And if multiple result sets are returned, any of the result sets may be empty too. The following example assumes that the procedure `my_proc` exists.

MySQL Shell JavaScript Code

```
var res = mySession.sql('CALL my_proc()').execute();

if (res.hasData()){

    var row = res.fetchOne();
    if (row){
        print('List of rows available for fetching.');
        do {
            print(row);
        } while (row = res.fetchOne());
    }
    else{
        print('Empty list of rows.');
    }
}
else {
    print('No row result.');
}
```

MySQL Shell Python Code

```
res = mySession.sql('CALL my_proc()').execute()

if res.has_data():

    row = res.fetch_one()
    if row:
        print('List of rows available for fetching.')
        while row:
            print(row)
            row = res.fetch_one()
    else:
        print('Empty list of rows.')
else:
    print('No row result.)
```

Node.js JavaScript Code

```
mySession.sql('CALL my_proc()')
.execute()
.then(function (res) {
    if (!res.hasData()) {
        return console.log('No row result.');
    }

    var row = res.fetchOne();

    if (!row) {
        return console.log('Empty list of rows.');
    }

    console.log('List of rows available for fetching.');
```

```
        do {
            console.log(row);
        } while (row = res.fetchOne());
    })
}
```

C# Code

```
var res = Session.SQL("CALL my_proc()").Execute();

if (res.HasData)
{
    var row = res.FetchOne();
    if (row != null)
    {
        Console.WriteLine("List of rows available for fetching.");
        do
        {
            PrintResult(row);
        } while ((row = res.FetchOne()) != null);
    }
    else
    {
        Console.WriteLine("Empty list of rows.");
    }
}
else
{
    Console.WriteLine("No row result.");
}
```

Python Code

```
# Connector/Python
res = mySession.sql('CALL my_proc()').execute()

if res.has_data():

    row = res.fetch_one()
    if row:
        print('List of rows available for fetching.')
        while row:
            print(row)
            row = res.fetch_one()
    else:
        print('Empty list of rows.')
else:
    print('No row result.)
```

Java Code

```
SqlResult res = mySession.sql("CALL my_proc()").execute();

if (res.hasData()){

    Row row = res.fetchOne();
    if (row != null){
        System.out.println("List of rows available for fetching.");
        do {
            for (int c = 0; c < res.getColumnCount(); c++) {
                System.out.println(row.getString(c));
            }
        } while ((row = res.fetchOne()) != null);
    }
    else{
```

```
        System.out.println("Empty list of rows.");
    }
}
else {
    System.out.println("No row result.");
}
```

C++ Code

```
SqlResult res = mysession.sql("CALL my_proc()").execute();

if (res.hasData())
{
    Row row = res.fetchOne();
    if (row)
    {
        cout << "List of rows available for fetching." << endl;
        do {
            cout << "next row: ";
            for (unsigned i=0 ; i < row.colCount(); ++i)
                cout << row[i] << ", ";
            cout << endl;
        } while ((row = res.fetchOne()));
    }
    else
    {
        cout << "Empty list of rows." << endl;
    }
}
else
{
    cout << "No row result." << endl;
}
```

It is an error to call either `fetchOne()` or `fetchAll()` when `hasData()` indicates that an `SqlResult` is not a data set.

MySQL Shell JavaScript Code

```
function print_result(res) {
    if (res.hasData()) {
        // SELECT
        var columns = res.getColumns();
        var record = res.fetchOne();

        while (record){
            for (index in columns){
                print (columns[index].getColumnName() + ":" + record[index] + "\n");
            }

            // Get the next record
            record = res.fetchOne();
        }

    } else {
        // INSERT, UPDATE, DELETE, ...
        print('Rows affected: ' + res.getAffectedItemCount());
    }
}

print_result(mysession.sql('DELETE FROM users WHERE age < 30').execute());
print_result(mysession.sql('SELECT * FROM users WHERE age = 40').execute());
```

MySQL Shell Python Code

```
def print_result(res):
```

```
if res.has_data():
    # SELECT
    columns = res.get_columns()
    record = res.fetch_one()

    while record:
        index = 0

        for column in columns:
            print("%s: %s \n" % (column.get_column_name(), record[index]))
            index = index + 1

        # Get the next record
        record = res.fetch_one()
else:
    #INSERT, UPDATE, DELETE, ...
    print('Rows affected: %s' % res.get_affected_items_count())

print_result(mySession.sql('DELETE FROM users WHERE age < 30').execute())
print_result(mySession.sql('SELECT * FROM users WHERE age = 40').execute())
```

Node.js JavaScript Code

```
function print_result(res) {
    if (res.hasData()) {
        // SELECT
        var columns = res.getColumns();
        var record = res.fetchOne();

        while (record) {
            for (index in columns) {
                console.log(columns[index].getColumnName() + ": " + record[index]);
            }

            // Get the next record
            record = res.fetchOne();
        }
    } else {
        // INSERT, UPDATE, DELETE, ...
        console.log('Rows affected: ' + res.getAffectedItemCount());
    }
}

mySession.sql(`DELETE FROM users WHERE age < 30`)
    .execute()
    .then(function (res) {
        print_result(res);
    });

mySession.sql(`SELECT * FROM users WHERE age = 40`)
    .execute()
    .then(function (res) {
        print_result(res);
    });
}
```

C# Code

```
private void print_result(SqlResult res)
{
    if (res.HasData)
    {
        // SELECT
    }
    else
    {
```

```

        // INSERT, UPDATE, DELETE, ...
        Console.WriteLine("Rows affected: " + res.RecordsAffected);
    }
}

print_result(Session.SQL("DELETE FROM users WHERE age < 30").Execute());
print_result(Session.SQL("SELECT COUNT(*) AS forty FROM users WHERE age = 40").Execute());

```

Python Code

```

# Connector/Python
def print_result(res):
    if res.has_data():
        # SELECT
        columns = res.get_columns()
        record = res.fetch_one()

        while record:
            index = 0

            for column in columns:
                print("%s: %s \n" % (column.get_column_name(), record[index]))
                index = index + 1

            # Get the next record
            record = res.fetch_one()

    else:
        #INSERT, UPDATE, DELETE, ...
        print('Rows affected: %s' % res.get_affected_items_count())

print_result(mySession.sql('DELETE FROM users WHERE age < 30').execute())
print_result(mySession.sql('SELECT * FROM users WHERE age = 40').execute())

```

Java Code

```

private void print_result(SqlResult res) {
    if (res.hasData()) {
        // SELECT
        Row row;
        while ((row = res.fetchOne()) != null){
            for (int c = 0; c < res.getColumnCount(); c++) {
                System.out.println(row.getString(c));
            }
        }
    } else {
        // INSERT, UPDATE, DELETE, ...
        System.out.println("Rows affected: " + res.getAffectedItemCount());
    }
}

print_result(mySession.sql("DELETE FROM users WHERE age < 30").execute());
print_result(mySession.sql("SELECT COUNT(*) AS forty FROM users WHERE age = 40").execute());

```

C++ Code

```

void print_result(SqlResult &&_res)
{
    // Note: We need to store the result somewhere to be able to process it.

    SqlResult res(std::move(_res));

    if (res.hasData())
    {

```

```
// SELECT
const Columns &columns = res.getColumns();
Row record = res.fetchOne();

while (record)
{
    for (unsigned index=0; index < res.getColumnCount(); ++index)
    {
        cout << columns[index].getColumnName() << ":" <<
            record[index] << endl;
    }

    // Get the next record
    record = res.fetchOne();
}

}
else
{
    // INSERT, UPDATE, DELETE, ...
    // Note: getAffectedItemsCount() not yet implemented in Connector/C++.
    cout << "No rows in the result" << endl;
}

print_result(mysession.sql("DELETE FROM users WHERE age < 30").execute());
print_result(mysession.sql("SELECT * FROM users WHERE age = 40").execute());
```

Calling a stored procedure might result in having to deal with multiple result sets as part of a single execution. As a result for the query execution an `SqlResult` object is returned, which encapsulates the first result set. After processing the result set you can call `nextResult()` to move forward to the next result, if there is any. Once you advanced to the next result set, it replaces the previously loaded result which then becomes unavailable.

MySQL Shell JavaScript Code

```
function print_result(res) {
    if (res.hasData()) {
        // SELECT
        var columns = res.getColumns();
        var record = res.fetchOne();

        while (record){
            for (index in columns){
                print (columns[index].getColumnName() + ":" + record[index] + "\n");
            }

            // Get the next record
            record = res.fetchOne();
        }

    } else {
        // INSERT, UPDATE, DELETE, ...
        print('Rows affected: ' + res.getAffectedItemCount());
    }
}

var res = mySession.sql('CALL my_proc()').execute();

// Prints each returned result
var more = true;
while (more){
    print_result(res);

    more = res.nextResult();
```

```
}
```

MySQL Shell Python Code

```
def print_result(res):
    if res.has_data():
        # SELECT
        columns = res.get_columns()
        record = res.fetch_one()

        while record:
            index = 0

            for column in columns:
                print("%s: %s \n" % (column.get_column_name(), record[index]))
                index = index + 1

            # Get the next record
            record = res.fetch_one()
    else:
        #INSERT, UPDATE, DELETE, ...
        print('Rows affected: %s' % res.get_affected_items_count())

    res = mySession.sql('CALL my_proc()').execute()

    # Prints each returned result
    more = True
    while more:
        print_result(res)
        more = res.next_result()
```

Node.js JavaScript Code

```
function print_result(res) {
    if (res.hasData()) {
        // SELECT
        var columns = res.getColumns();
        var record = res.fetchOne();

        while (record) {
            for (index in columns) {
                console.log(columns[index].getColumnName() + ": " + record[index]);
            }

            // Get the next record
            record = res.fetchOne();
        }
    } else {
        // INSERT, UPDATE, DELETE, ...
        console.log('Rows affected: ' + res.getAffectedItemCount());
    }
}

mySession.sql('CALL my_proc()')
    .execute()
    .then(function (res) {
        // Prints each returned result
        var more = true;

        while (more) {
            print_result(res);

            more = res.nextResult();
        }
    })
}
```

C# Code

```
var res = Session.SQL("CALL my_proc()").Execute();

if (res.HasData)
{
    do
    {
        Console.WriteLine("New resultset");
        while (res.Next())
        {
            Console.WriteLine(res.Current);
        }
    } while (res.NextResult());
}
```

Python Code

```
# Connector/Python
def print_result(res):
    if res.has_data():
        # SELECT
        columns = res.get_columns()
        record = res.fetch_one()

        while record:
            index = 0

            for column in columns:
                print("%s: %s \n" % (column.get_column_name(), record[index]))
                index = index + 1

            # Get the next record
            record = res.fetch_one()
    else:
        #INSERT, UPDATE, DELETE, ...
        print('Rows affected: %s' % res.get_affected_items_count())

res = mySession.sql('CALL my_proc()').execute()

# Prints each returned result
more = True
while more:
    print_result(res)

    more = res.next_result()
```

Java Code

```
SqlResult res = mySession.sql("CALL my_proc()").execute();
```

C++ Code

```
SqlResult res = mysession.sql("CALL my_proc()").execute();

while (true)
{
    if (res.hasData())
    {
        cout << "List of rows in the resultset." << endl;
        for (Row row; (row = res.fetchOne());)
        {
            cout << "next row: ";
            for (unsigned i = 0; i < row.colCount(); ++i)
                cout << row[i] << ", ";
```

```
        cout << endl;
    }
}
else
{
    cout << "No rows in the resultset." << endl;
}

if (!res.nextResult())
    break;

cout << "Next resultset." << endl;
}
```

When using Node.js, individual rows can be returned immediately using a callback, which has to be provided to the `execute()` method. To identify individual result sets you can provide a second callback, which is called for meta data that marks the beginning of a result set.

Node.js JavaScript Code

```
var resultcount = 0;
var res = session
    .sql('CALL my_proc()')
    .execute(
        function (row) {
            console.log(row);
        },
        function (meta) {
            console.log('Begin of result set number ', resultCount++);
        });

```

The number of result sets is not known immediately after the query execution. Query results can be streamed to the client or buffered at the client. In the streaming or partial buffering mode a client cannot tell whether a query emits more than one result set.

9.6 Working with Metadata

Results contain metadata related to the origin and types of results from relational queries. This metadata can be used by applications that need to deal with dynamic query results or format results for transformation or display. Result metadata is accessible via instances of `Column`. An array of columns can be obtained from any `RowResult` using the `getColumns()` method.

For example, the following metadata is returned in response to the query `SELECT 1+1 AS a, b FROM mydb.some_table_with_b AS b_table.`

```
Column[0].databaseName = NULL
Column[0].tableName = NULL
Column[0].tableLabel = NULL
Column[0].columnName = NULL
Column[0].columnLabel = "a"
Column[0].type = BIGINT
Column[0].length = 3
Column[0].fractionalDigits = 0
Column[0].numberSigned = TRUE
Column[0].collationName = "binary"
Column[0].characterSetName = "binary"
Column[0].padded = FALSE

Column[1].databaseName = "mydb"
Column[1].tableName = "some_table_with_b"
Column[1].tableLabel = "b_table"
Column[1].columnName = "b"
Column[1].columnLabel = "b"
```

```
Column[1].type = STRING
Column[1].length = 20 (e.g.)
Column[1].fractionalDigits = 0
Column[1].numberSigned = TRUE
Column[1].collationName = "utf8mb4_general_ci"
Column[1].characterSetName = "utf8mb4"
Column[1].padded = FALSE
```

9.7 Support for Language Native Iterators

All implementations of the DevAPI feature the methods shown in the UML diagram at the beginning of this chapter. All implementations allow result set iteration using `fetchOne()`, `fetchAll()` and `nextResult()`. In addition to the unified API drivers should implement language native iteration patterns. This applies to any type of data set (`DocResult`, `RowResult`, `SqlResult`) and to the list of items returned by `fetchAll()`. You can choose whether you want your X DevAPI-based application code to offer the same look and feel in all programming languages used or go for the natural style of a programming language.

Chapter 10 Building Expressions

Table of Contents

10.1 Expression Strings	125
10.1.1 Boolean Expression Strings	125
10.1.2 Value Expression Strings	125

This section explains how to build expressions using X DevAPI.

When working with MySQL expressions used in CRUD, statements can be specified in two ways. The first is to use strings to formulate the expressions which should be familiar if you have developed code with SQL before. The other method is to use Expression Builder functionality.

10.1 Expression Strings

Defining string expressions is straightforward as these are easy to read and write. The disadvantage is that they need to be parsed before they can be transferred to the MySQL server. In addition, type checking can only be done at runtime. All implementations can use the syntax illustrated here, which is shown as MySQL Shell JavaScript code.

```
// Using a string expression to get all documents that
// have the name field starting with 'S'
var myDocs = myColl.find('name like :name').bind('name', 'S%').execute();
```

10.1.1 Boolean Expression Strings

Boolean expression strings can be used when filtering collections or tables using operations, such as `find()` and `remove()`. The expression is evaluated once for each document or row.

The following example of a boolean expression string uses `find()` to search for all documents with a “red” color attribute from the collection “apples”:

```
apples.find('color = "red"').execute()
```

Similarly, to delete all red apples:

```
apples.remove('color = "red"').execute()
```

10.1.2 Value Expression Strings

Value expression strings are used to compute a value which can then be assigned to a given field or column. This is necessary for both `modify()` and `update()`, as well as computing values in documents at insertion time.

An example use of a value expression string would be to increment a counter. The `expr()` function is used to wrap strings where they would otherwise be interpreted literally. For example, to increment a counter:

```
// the expression is evaluated on the server
collection.modify('true').set("counter", expr("counter + 1")).execute()
```

If you do not wrap the string with `expr()`, it would be assigning the literal string “counter + 1” to the “counter” member:

```
// equivalent to directly assigning a string: counter = "counter + 1"  
collection.modify('true').set("counter", "counter + 1").execute()
```

Chapter 11 CRUD EBNF Definitions

Table of Contents

11.1 Session Objects and Functions	127
11.2 Schema Objects and Functions	129
11.3 Collection CRUD Functions	132
11.4 Collection Index Management Functions	134
11.5 Table CRUD Functions	134
11.6 Result Functions	136
11.7 Other EBNF Definitions	139


This chapter provides a visual reference guide to the objects and functions available in the X DevAPI.

11.1 Session Objects and Functions

Session

The syntax for this object shown in EBNF is:

```
Session
 ::= '.getSchema(' StringLiteral ')'
   | '.getSchemas()'
   | '.createSchema(' StringLiteral ')'
   | '.dropSchema(' StringLiteral ')'
   | '.getDefaultSchema()'
   | '.startTransaction()'
   | '.commit()'
   | '.rollback()'
   | '.setSavepoint()'
   | '.setSavepoint(' StringLiteral ')'
   | '.releaseSavePoint(' StringLiteral ')'
   | '.rollbackTo(' StringLiteral ')'
   | '.close()'
   | SqlExecute
```


Figure 11.1 Session

SqlExecute

The syntax for this function shown in EBNF is:

```

SqlExecute
 ::= '.sql(' SqlStatementStr ')'
   ( '.bind(' Literal (',' Literal)* ')')*
   ( '.execute()' )?
  
```

Figure 11.2 SqlExecute

SQLPlaceholderValues

The syntax for this function shown in EBNF is:

```
SQLPlaceholderValues
 ::= '{' SQLPlaceholderName ':' ( SQLLiteral ) '}'
```

Figure 11.3 SQLPlaceholderValues

SQLPlaceholderName

The syntax for this function shown in EBNF is:


```
SQLPlaceholderName
 ::= '?'
```

Figure 11.4 SQLPlaceholderName

SQLLiteral

The syntax for this function shown in EBNF is:

```
SQLLiteral
 ::= ''' StringLiteral ''' | Number | Document
```


Figure 11.5 SQLLiteral

11.2 Schema Objects and Functions

Schema

The syntax for this function shown in EBNF is:

```
Schema
 ::= '.getName()'
 | '.existsInDatabase()'
 | '.getSession()'
 | '.getCollection(' StringLiteral ')'
 | '.getCollections()'
 | '.getCollectionAsTable(' StringLiteral ')'
 | '.dropCollection(' StringLiteral ')'
 | '.getTable(' StringLiteral ')'
 | '.getTables()'
 | '.createCollection(' StringLiteral ')'
```


Figure 11.6 Schema

Collection

The syntax for this function shown in EBNF is:

```

Collection
 ::= '.getSchema()'
 | '.getName()'
 | '.getSession()'
 | '.existsInDatabase()'
 | '.replaceOne(' DocumentId ',' DocumentOrJSON '' )
 | '.addOrReplaceOne(' DocumentId ',' DocumentOrJSON '' )
 | '.getOne(' DocumentId '' )
 | '.removeOne(' DocumentId '' )
 | CollectionFindFunction
 | CollectionModifyFunction
 | CollectionAddFunction
 | CollectionRemoveFunction
 | CollectionCreateIndex
 | CollectionDropIndex
  
```


Figure 11.7 Collection

Table

The syntax for this function shown in EBNF is:

```

Table
 ::= '.getSchema()'
 | '.getName()'
 | '.getSession()'
 | '.existsInDatabase()'
 | '.isView()'
 | TableSelectFunction
 | TableUpdateFunction
 | TableInsertFunction
 | TableDeleteFunction
  
```

Figure 11.8 Table

11.3 Collection CRUD Functions


CollectionFindFunction

The syntax for this function in EBNF is:

```

CollectionFindFunction
 ::= '.find(' SearchConditionStr? ')'
   ( '.fields(' ProjectedDocumentExprStr ')')?
   ( '.groupBy(' SearchExprStrList ')')?
   ( '.having(' SearchConditionStr ')')?
   ( '.sort(' SortExprStrList ')')?
   ( '.limit(' NumberOfRows ')'
     ( '.offset(' NumberOfRows ')')?)?
   ( '.lockExclusive(' LockContention ')'
     | '.lockShared(' LockContention ')')?)?
   ( '.bind(' PlaceholderValues ')')*
   ( '.execute()' )?

```


Figure 11.9 CollectionFindFunction

CollectionModifyFunction

The syntax for this function shown in EBNF is:

```
CollectionModifyFunction
 ::= '.modify(' SearchConditionStr ')'
   ( '.set(' DocPath ',' ExprOrLiteral ')' | 
     '.unset(' DocPath ( ',' DocPath )* ')' | 
     '.arrayInsert(' DocPath ',' ExprOrLiteral ')' | 
     '.arrayAppend(' DocPath ',' ExprOrLiteral ')' | 
     '.patch(' DocumentOrJSON ')'
   )+
   ( '.sort(' SortExprStrList ')' )? ( '.limit(' NumberOfRows ')' )?
   ( '.bind(' PlaceholderValues ')' )*
   ( '.execute()' )?
```

Figure 11.10 CollectionModifyFunction




CollectionAddFunction

The syntax for this function shown in EBNF is:

```
CollectionAddFunction
 ::= ( '.add(' ( DocumentOrJSON | '[]' DocumentOrJSON ( ',' DocumentOrJSON )* ']' )? ' )' )+
   ( '.execute()' )?
```

Figure 11.11 CollectionAddFunction




CollectionRemoveFunction

The syntax for this function shown in EBNF is:

```
CollectionRemoveFunction
 ::= '.remove(' SearchConditionStr ')'
   ( '.sort(' SortExprStrList ')' )?
   ( '.limit(' NumberOfRows ')' )?
   ( '.bind(' PlaceholderValues ')' )*
   ( '.execute()' )?
```

Figure 11.12 CollectionRemoveFunction



11.4 Collection Index Management Functions

Collection.createIndex() Function

The syntax for this function shown in EBNF is:

```
CollectionCreateIndex
 ::= '.createIndex(' StringLiteral ',' DocumentOrJSON ')'
```

Figure 11.13 CollectionCreateIndexFunction



CollectionDropIndex

The syntax for this function shown in EBNF is:

```
CollectionDropIndex
 ::= '.dropIndex(' StringLiteral ')'
```

Figure 11.14 CollectionDropIndex



11.5 Table CRUD Functions

TableSelectFunction

Table.select() and collection.find() use different methods for sorting results.
Table.select() follows the SQL language naming and calls the sort method orderBy().


TableInsertFunction

`Collection.find()` does not. Use the method `sort()` to sort the results returned by `Collection.find()`. Proximity with the SQL standard is considered more important than API uniformity here.

The syntax for this function shown in EBNF is:

```
TableSelectFunction
 ::= '.select(' ProjectedSearchExprStrList? ')'
   ( '.where(' SearchConditionStr ')')?
   ( '.groupBy(' SearchExprStrList ')')?
   ( '.having(' SearchConditionStr ')')?
   ( '.orderBy(' SortExprStrList ')')?
   ( '.limit(' NumberOfRows ')')?
   ( '.offset(' NumberOfRows ')')?
   ( '.lockExclusive(' LockContention ')')|'.lockShared(' LockContention ')')?
   ( '.bind(' ( PlaceholderValues ) ')')*
   ( '.execute()')?
```

Figure 11.15 TableSelectFunction




TableInsertFunction

The syntax for this function shown in EBNF is:

```
TableInsertFunction
 ::= '.insert(' ( TableFields )? ')'
   ( '.values(' ExprOrLiteral (', ExprOrLiteral)* ')')+
   ( '.execute()')?
```

Figure 11.16 TableInsertFunction




TableUpdateFunction

The syntax for this function shown in EBNF is:

TableDeleteFunction

```
TableUpdateFunction
 ::= '.update()'
   ( '.set(' TableField ',' ExprOrLiteral ')' )+ '.where(' SearchConditionStr ')
   ( '.orderBy(' SortExprStrList ')' )? ( '.limit(' NumberOfRows ')' )?
   ( '.bind(' ( PlaceholderValues ) ')' )*
   ( '.execute()' )?
```

Figure 11.17 TableUpdateFunction




TableDeleteFunction

The syntax for this function shown in EBNF is:

```
TableDeleteFunction
 ::= '.delete()' '.where(' SearchConditionStr ')
   ( '.orderBy(' SortExprStrList ')' )? ( '.limit(' NumberOfRows ')' )?
   ( '.bind(' ( PlaceholderValues ) ')' )*
   ( '.execute()' )?
```

Figure 11.18 TableDeleteFunction




11.6 Result Functions

Result

The syntax for this function shown in EBNF is:

```
Result
 ::= '.getAffectedItemCount()'
   | '.getAutoIncrementValue()'
   | '.getGeneratedIds()'
   | '.getWarningsCount()'
   | '.getWarnings()'
```

Figure 11.19 Result

DocResult

The syntax for this function shown in EBNF is:

```

DocResult
 ::= '.getWarningsCount()'
 | '.getWarnings()'
 | '.fetchAll()'
 | '.fetchOne()'

```

Figure 11.20 DocResult


RowResult

The syntax for this function shown in EBNF is:

```

RowResult
 ::= '.getWarningsCount()'
 | '.getWarnings()'
 | '.fetchAll()'
 | '.fetchOne()'
 | '.getColumns()'

```


Figure 11.21 RowResult

Column

The syntax for this function shown in EBNF is:

```
Column
 ::= '.getSchemaName()'
 | '.getTableName()'
 | '.getTableLabel()'
 | '.getColumnName()'
 | '.getColumnLabel()'
 | '.getType()'
 | '.getLength()'
 | '.getFractionalDigits()'
 | '.isNumberSigned()'
 | '.getCollationName()'
 | '.getCharacterSetName()'
 | '.isPadded()'
```

Figure 11.22 Column




SqlResult

The syntax for this function shown in EBNF is:

```
SqlResult
 ::= '.getWarningsCount()'
 | '.getWarnings()'
 | '.fetchAll()'
 | '.fetchOne()'
 | '.getColumns()'
 | '.getAutoIncrementValue()'
 | '.hasData()'
```

```
| '.nextResult()'
```

Figure 11.23 SqlResult

11.7 Other EBNF Definitions

SearchConditionStr

The syntax for this function shown in EBNF is:

```
SearchConditionStr
 ::= ''' Expression '''
```

Figure 11.24 SearchConditionStr

SearchExprStrList

The syntax for this function shown in EBNF is:

```
SearchExprStrList
 ::= '[' ''' Expression ''' ( ',' ''' Expression ''' )* ']'
```

Figure 11.25 SearchExprStrList

ProjectedDocumentExprStr

The syntax for this function shown in EBNF is:

```
ProjectedDocumentExprStr
```


```
 ::= ProjectedSearchExprStrList | 'expr("' JSONDocumentExpression '')'
```

Figure 11.26 ProjectedDocumentExprStr

ProjectedSearchExprStrList

The syntax for this function shown in EBNF is:


```
ProjectedSearchExprStrList
 ::= '[' ''' Expression ( 'AS' Alias )? ''' ( ',' ''' Expression ( 'AS' Alias )? ''' )* ']'
```

Figure 11.27 ProjectedSearchExprStrList

SortExprStrList

The syntax for this function shown in EBNF is:

```
SortExprStrList
 ::= '[' ''' Expression ( 'ASC' | 'DESC' )? ''' ( ',' ''' Expression ( 'ASC' | 'DESC' )? ''' )* ']'
```

Figure 11.28 SortExprStrList

ExprOrLiteral

The syntax for this function shown in EBNF is:

```
ExprOrLiteral
 ::= 'expr("' Expression '"')' | Literal
```

Figure 11.29 ExprOrLiteral

ExprOrLiterals

The syntax for this function shown in EBNF is:

```
ExprOrLiterals
 ::= ExprOrLiteral ( ',' ExprOrLiteral )*
```

Figure 11.30 ExprOrLiterals




ExprOrLiteralOrOperand

The syntax for this function shown in EBNF is:

```
ExprOrLiteralOrOperand
 ::= ExprOrLiteral
```

Figure 11.31 ExprOrLiteralOrOperand




PlaceholderValues

The syntax for this function shown in EBNF is:

```
PlaceholderValues
 ::= '{' PlaceholderName ':' ( ExprOrLiteral ) '}'
```

Figure 11.32 PlaceholderValues




PlaceholderName

The syntax for this function shown in EBNF is:

```
PlaceholderName
 ::= NamedPlaceholderNotQuestionmarkNotNumbered
```


Figure 11.33 PlaceholderName



DocPath

The syntax for this function shown in EBNF is:

```
DocPath
 ::= ( '[*]' | ( '[' Index ']' ) | '.*' | ( '.', StringLiteral ) | '**' )+
```


Figure 11.34 DocPath

Literal

The syntax for this function shown in EBNF is:

```


Literal
 ::= ''' StringLiteral ''' | Number | true | false | Document
  
```

Figure 11.35 Literal

Expression

```

Expression
 ::= Literal
 | DocPath
 | TableField
 | FunctionName '(' Expression ( ',' Expression )* ')'
 | ':' PlaceholderName
 | Expression Operator Expression
 | JSONExpression
  
```

Figure 11.36 Expression

Document

An API call expecting a JSON document allows the use of many data types to describe the document. Depending on the X DevAPI implementation and language any of the following data types can be used:


- String
- Native JSON
- JSON equivalent syntax
- DbDoc
- Generated Doc Classes

All implementations of X DevAPI allow expressing a document by the special DbDoc type and as a string.

The syntax for this function shown in EBNF is:

```


Document
 ::= JSONDocument | JSONEquivalentDocument | DbDoc | GeneratedDocumentClasses
    
```

Figure 11.37 Document

JSONExpression

The syntax for this function shown in EBNF is:


```
JSONExpression
 ::= JSONDocumentExpression | '[' Expression ( ',' Expression )* ']'
```

Figure 11.38 JSONExpression

JSONDocumentExpression

The syntax for this function shown in EBNF is:


```
JSONDocumentExpression
 ::= '{' StringLiteral ':' JSONExpression (',' StringLiteral ':' JSONExpression)* '}'
```

Figure 11.39 JSONDocumentExpression

FunctionName

The syntax for this function shown in EBNF is:


```
FunctionName
 ::= StringLiteral | StringLiteral '.' StringLiteral
```

Figure 11.40 FunctionName

DocumentOrJSON

The syntax for this function shown in EBNF is:

```
DocumentOrJSON
 ::= Document | 'expr("' JSONDocumentExpression '"')'
```


Figure 11.41 DocumentOrJSON

TableField

The syntax for this function shown in EBNF is:

```
TableField
 ::= ( StringLiteral '.' )? ( StringLiteral '.' )? StringLiteral ( '@' DocPath )?
```

Figure 11.42 TableField




TableFields

The syntax for this function shown in EBNF is:

```
TableFields
 ::= [ TableField ( ',' TableField )* ]'
```

Figure 11.43 TableFields



Chapter 12 Expressions EBNF Definitions

This section provides a visual reference guide to the grammar for the expression language used in X DevAPI.

ident

```
ident
 ::= ID | QUOTED_ID
```


Figure 12.1 ident



schemaQualifiedIdent

```
schemaQualifiedIdent
 ::= ( ident_schema '.' )? ident
```


Figure 12.2 schemaQualifiedIdent



columnIdent


Figure 12.3 columnIdent

```
columnIdent
 ::= ( ident '.' ( ident '.' )? )? ident ( ( '->' | '->>' ) "" '$' documentPath ''')?
```



documentPathLastItem

```
documentPathLastItem
 ::= '[' '*' '
 | '[' INT ']'
 | '..*'
 | '..' documentPathMember
```

Figure 12.4 documentPathLastItem

documentPathItem

```


documentPathItem
  ::= documentPathLastItem
  | ...
  
```

Figure 12.5 documentPathItem

documentPath

```

documentPath
  ::= documentPathItem* documentPathLastItem
  
```


Figure 12.6 documentPath

documentField

Figure 12.7 documentField

```


documentField
  ::= fieldId documentPath*
  | '$' documentPath
  
```



argsList

```

argsList ::= expr ( ',' expr )*
  
```

Figure 12.8 argsList

lengthSpec

```
lengthSpec ::= '(' INT ')'
```


Figure 12.9 lengthSpec

castType

```
castType ::= 'SIGNED' 'INTEGER'*  
          | 'UNSIGNED' 'INTEGER'*  
          | 'CHAR' lengthSpec*  
          | 'BINARY' lengthSpec*  
          | 'DECIMAL' ( lengthSpec | '(' INT ',' INT ')' )?  
          | 'TIME'  
          | 'DATE'  
          | 'DATETIME'  
          | 'JSON'
```

functionCall


Figure 12.10 castType



functionCall

```
functionCall ::= schemaQualifiedIdent '(' argsList? ')'
```


Figure 12.11 functionCall



placeholder

```
placeholder ::= ':' ID
```

Figure 12.12 placeholder




groupedExpr

```
groupedExpr ::= '(' expr ')'
```

Figure 12.13 groupedExpr


unaryOp

```
unaryOp ::= ( '!' | '~' | '+' | '-' ) atomicExpr
```

Figure 12.14 unaryOp


literal

```
literal ::= INT
          | FLOAT
          | STRING_SQ
          | STRING_DQ
          | 'NULL'
          | 'FALSE'
          | 'TRUE'
```

Figure 12.15 literal

jsonKeyValue


```
jsonKeyValue ::= STRING_DQ ':' expr
```

Figure 12.16 jsonKeyValue

jsonDoc

```
jsonDoc ::= '{' ( jsonKeyValue ( ',' jsonKeyValue )* )* '}'
```


Figure 12.17 jsonDoc



jsonarray

```
jsonArray ::= '[' ( expr ( ',' expr )* )* ']'
```


Figure 12.18 jsonArray



atomicExpr


```
atomicExpr
 ::= placeholder
 | columnOrPath
 | functionCall
 | groupedExpr
 | unaryOp
 | castOp
```

Figure 12.19 atomicExpr



intervalUnit

```
INTERVAL_UNIT
  ::= 'MICROSECOND'
    | 'SECOND'
    | 'MINUTE'
    | 'HOUR'
    | 'DAY'
    | 'WEEK'
    | 'MONTH'
    | 'QUARTER'
    | 'YEAR'
    | 'SECOND_MICROSECOND'
    | 'MINUTE_MICROSECOND'
    | 'MINUTE_SECOND'
    | 'HOUR_MICROSECOND'
    | 'HOUR_SECOND'
    | 'HOUR_MINUTE'
    | 'DAY_MICROSECOND'
    | 'DAY_SECOND'
    | 'DAY_MINUTE'
    | 'DAY_HOUR'
    | 'YEAR_MONTH'
```

Figure 12.20 INTERVAL_UNIT**interval**


```
interval ::= 'INTERVAL' expr INTERVAL_UNIT
```

Figure 12.21 interval

intervalExpr

```
intervalExpr
  ::= atomicExpr ( ( '+' | '-' ) interval )*
```


Figure 12.22 intervalExpr



mulDivExpr

```
mulDivExpr
  ::= intervalExpr ( ( '*' | '/' | '%' ) intervalExpr )*
```


Figure 12.23 mulDivExpr



addSubExpr

```
addSubExpr
  ::= mulDivExpr ( ( '+' | '-' ) mulDivExpr )*
```


Figure 12.24 addSubExpr



shiftExpr

```
shiftExpr
  ::= addSubExpr ( ( '<<' | '>>' ) addSubExpr )*
```


Figure 12.25 shiftExpr



bitExpr

```
bitExpr ::= shiftExpr ( ( '&' | '!' | '^' ) shiftExpr )*
```


Figure 12.26 bitExpr



compExpr


```
compExpr ::= bitExpr ( ( '>=' | '>' | '<=' | '<' | '=' | '<>' | '!=') bitExpr )*
```

Figure 12.27 compExpr




ilriExpr

```
ilriExpr ::= compExpr 'IS' 'NOT'* ( 'NULL' | 'TRUE' | 'FALSE' )
           | compExpr 'NOT'* 'IN' '(' argsList* ')'
           | compExpr 'NOT'* 'IN' compExpr
           | compExpr 'NOT'* 'LIKE' compExpr ( 'ESCAPE' compExpr )*
           | compExpr 'NOT'* 'BETWEEN' compExpr 'AND' compExpr
           | compExpr 'NOT'* 'REGEXP' compExpr
           | compExpr
```

Figure 12.28 ilriExpr

andExpr

```
andExpr ::= ilriExpr ( ( '&&' | 'AND' ) ilriExpr )*
```


Figure 12.29 andExpr

orExpr

```
orExpr ::= andExpr ( ( ' | ' | 'OR' ) andExpr )*
```

expr

Figure 12.30 orExpr



expr

```
expr      ::= orExpr
```


Figure 12.31 expr



DIGIT

```
DIGIT     ::= '0' - '9'
```


Figure 12.32 DIGIT



FLOAT

```
FLOAT      ::= DIGIT* '.' DIGIT+ ( 'E' ( '+' | '-' )* DIGIT+ )*
              | DIGIT+ 'E' ( '+' | '-' )* DIGIT+
```

Figure 12.33 FLOAT



INT

```
INT      ::= DIGIT+
```

QUOTED_ID


Figure 12.34 INT



QUOTED_ID

```
QUOTED_ID ::= `` ID ``  
           | `` ( ~ ``* | ``* )* ``
```


Figure 12.35 QUOTED_ID



ID


```
ID ::= ( 'a' - 'z' | 'A' - 'Z' | '_' ) ( 'a' - 'z' | 'A' - 'Z' | '0' - '9' | '_' )*
```

Figure 12.36 ID



WS

```
WS ::= [ \t\r\n]+
```


Figure 12.37 WS

SCHAR

```
SCHAR      ::= [ \u0020\u0021\u0023\u0024\u0025\u0026\u0028-\u005B\u005D-\u007E ]
```

SCHAR


Figure 0-38 SCHAR



STRING_DQ

```
STRING_DQ
 ::= " " ( SCHAR | " " | ESCAPED_DQ )* " "
```


Figure 12.39 STRING_DQ



STRING_SQ

```
STRING_SQ
 ::= " " ( SCHAR | " " | ESCAPED_SQ )* " "
```

Figure 12.40 STRING_SQ



Chapter 13 Implementation Notes

Table of Contents

13.1 MySQL Connector Notes	163
13.2 MySQL Shell X DevAPI extensions	163
13.3 MySQL Connector/Node.js Notes	163
13.4 MySQL Connector/J Notes	163

This section provides notes on the different language-specific implementations of X DevAPI.

13.1 MySQL Connector Notes

Each driver implementation of X DevAPI may deviate from the description in marginal details to align the implementation to the common pattern and styles of the host language. All class names are identical among drivers and all drivers support the same core concepts such as `find()` or the chaining supported for `find()` to ensure developers experience similar APIs in all implementations.

The following implementation differences are possible:

- Function names can be postfixed to add specialisation. For example, implementations can choose between 'execute([<flag_async>])' and/or 'executeAsync()'.
- Functions can have prefixes such as 'get'.
- Connectors may offer native language result set iteration patterns in addition to a basic `while()` loop shown in many examples. For example, drivers may define iterator interfaces or classes.

13.2 MySQL Shell X DevAPI extensions

MySQL Shell deviates from the Connector implementations of X DevAPI in certain places. A Connector can connect to MySQL Server instances running X Plugin only by means of X Protocol. MySQL Shell contains an extension of X DevAPI to access MySQL Server instances through X Protocol. An additional `ClassicSession` class is available to establish a connection to a single MySQL Server instance using classic MySQL protocol. The functionality of the `ClassicSession` is limited to basic schema browsing and SQL execution.

See [MySQL Shell 8.4](#), for more information.

13.3 MySQL Connector/Node.js Notes

MySQL Connector/Node.js is built with ECMAScript 6 Promise objects to provide an asynchronous API. All network operations return a Promise, which resolves when the server responds. Please refer to the [information](#) on the ES6 Promise implementation.

13.4 MySQL Connector/J Notes

The following are some features unique to X DevAPI for Connector/J:

- The creation of new `Session` objects with `SessionFactory`.
- The creation of new `client` objects with `ClientFactory`.

- The `JsonValue` interface and its implementations, including `JsonObject`, `JsonString`, `JsonArray`, and `JsonNumber`.
- X Protocol connection properties are prefixed by `xdevapi:`.
- An extended connection string syntax; see [Connection URL Syntax](#) for details.