



First Look: JDBC and InterSystems IRIS

Version 2018.1
2018-01-31

First Look: JDBC and InterSystems IRIS
InterSystems Version 2018.1 2018-01-31
Copyright © 2018 InterSystems Corporation
All rights reserved.



InterSystems, InterSystems Caché, InterSystems Ensemble, InterSystems HealthShare, HealthShare, InterSystems TrakCare, TrakCare, InterSystems DeepSee, and DeepSee are registered trademarks of InterSystems Corporation.



InterSystems IRIS Data Platform, InterSystems iKnow, Zen, and Caché Server Pages are trademarks of InterSystems Corporation.

All other brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

This document contains trade secret and confidential information which is the property of InterSystems Corporation, One Memorial Drive, Cambridge, MA 02142, or its affiliates, and is furnished for the sole purpose of the operation and maintenance of the products of InterSystems Corporation. No part of this publication is to be used for any other purpose, and this publication is not to be reproduced, copied, disclosed, transmitted, stored in a retrieval system or translated into any human or computer language, in any form, by any means, in whole or in part, without the express prior written consent of InterSystems Corporation.

The copying, use and disposition of this document and the software programs described herein is prohibited except to the limited extent set forth in the standard software license agreement(s) of InterSystems Corporation covering such programs and related documentation. InterSystems Corporation makes no representations and warranties concerning such software programs other than those set forth in such standard software license agreement(s). In addition, the liability of InterSystems Corporation for any losses or damages relating to or arising out of the use of such software programs is limited in the manner set forth in such standard software license agreement(s).

THE FOREGOING IS A GENERAL SUMMARY OF THE RESTRICTIONS AND LIMITATIONS IMPOSED BY INTERSYSTEMS CORPORATION ON THE USE OF, AND LIABILITY ARISING FROM, ITS COMPUTER SOFTWARE. FOR COMPLETE INFORMATION REFERENCE SHOULD BE MADE TO THE STANDARD SOFTWARE LICENSE AGREEMENT(S) OF INTERSYSTEMS CORPORATION, COPIES OF WHICH WILL BE MADE AVAILABLE UPON REQUEST.

InterSystems Corporation disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

For Support questions about any InterSystems products, contact:

InterSystems Worldwide Response Center (WRC)
Tel: +1-617-621-0700
Tel: +44 (0) 844 854 2917
Email: support@InterSystems.com

Table of Contents

First Look: JDBC and InterSystems IRIS	1
1 JDBC: How to Use It with InterSystems IRIS	3
2 JDBC: Why Is It Important?	1
3 JDBC: Part of InterSystems IRIS Java Connectivity Options	2
3.1 JDBC: What's Unique about Shared Memory Connections	2
4 JDBC: Exploring It	2
5 JDBC: For More Information	3

First Look: JDBC and InterSystems IRIS

If you want to use Java with InterSystems IRIS™, this document provides an introduction to how to use JDBC.

1 JDBC: How to Use It with InterSystems IRIS

InterSystems provides a fully compliant (JDBC 4.2), pure Java, type 4 JDBC driver, which is a single standalone JAR file with no dependencies.

If you are already familiar with JDBC, and have a JDK 1.7 or 1.8 installed, all you need to do is to add `iris-jdbc-3.0.0.jar` to your `CLASSPATH`. The driver name is `com.intersys.jdbc.CacheDriver`, and it's typically installed under

```
<install-dir>\dev\java\lib\JDK1x
```

where `JDK1x` is either `JDK17` or `JDK18` (which are the two currently JDK supported versions).

The URL (connection string) is:

```
jdbc:Cache://ipAddress:superserverPort/namespace
```

If you want to try an even faster way to connect and your application is on the same machine as your InterSystems IRIS instance, try a shared memory connection with the connection string:

```
jdbc:Cache://SHM||||:superserverPort/namespace
```

For more information about shared memory connections, see “[JDBC: What’s Unique about Shared Memory Connections](#)”.

The point of this document is to give you a taste of using JDBC with InterSystems IRIS without bogging you down in details, so we’ve kept this exploration simple. When you bring InterSystems IRIS to your production systems, though, there are many things you will need to do differently, such as in regard to (but not limited to) security. So be sure not to confuse this exploration of InterSystems IRIS with the real thing! The sources provided at the end of this document will give you a good idea of what’s involved in using JDBC with InterSystems IRIS in production.

2 JDBC: Why Is It Important?

If you don’t know all the fine points of JDBC: simply put, it is the industry standard for platform-agnostic database connectivity. Specifically, JDBC is a database technology from Oracle that provides access to data sources. It is used to establish database connectivity and provides a standard, reliable way for Java programmers to connect to many types of data sources or perform operations on them with SQL.

Connecting to InterSystems IRIS via JDBC is simple, especially if you’ve used JDBC before. Establishing a JDBC connection to InterSystems IRIS from a Java application allows you to run SQL commands against InterSystems IRIS databases from your Java application.

And the InterSystems IRIS implementation of JDBC has been tuned to provide the fastest, most efficient access to your data.

3 JDBC: Part of InterSystems IRIS Java Connectivity Options

The InterSystems IRIS JDBC driver is the core InterSystems IRIS Java component, and supports traditional relational (SQL) access. For tighter integration, InterSystems IRIS also provides a separate feature — the [InterSystems IRIS XEP component](#).

Taken all together, InterSystems IRIS provides a unique set of capabilities to use the same physical connection and transaction context to manipulate data using relational and object-oriented paradigms. For more complex applications, InterSystems fully supports Hibernate. Enabling all these forms of connectivity — InterSystems IRIS XEP, Hibernate, and also the InterSystems IRIS Spark Connector — is the InterSystems IRIS JDBC driver.

3.1 JDBC: What's Unique about Shared Memory Connections

As with other database platforms, a JDBC connection to InterSystems IRIS is over TCP/IP. To maximize performance, InterSystems IRIS also offers a Java *shared memory connection*. Shared memory connections are available to Java applications running on the same Windows machine as an InterSystems IRIS instance.

A shared memory connection is a temporary device, backed either by a file or virtual memory, that is shared by a JDBC client and an instance of InterSystems IRIS running on the same physical machine. Further, these connections do not require potentially expensive calls into the kernel network stack. By using a channel directly from the JDBC client to InterSystems IRIS, they provide the ultimate in low latency and high throughput for JDBC operations.

For now, shared memory connections are available for JDBC only on Microsoft Windows. In future releases of InterSystems IRIS, you will be able to use shared memory connections on other platforms, including Linux.

4 JDBC: Exploring It

We've developed a demo that shows you how to work with JDBC and InterSystems IRIS — and how straightforward that is.

Please note that this demo does not show you the power of the [InterSystems Java shared memory connection](#), because it does not deal with the large volumes of data that the shared memory connection can handle so efficiently. On the other hand, it does allow you to see how easy it is to use the shared memory connection — by simply changing part of one line of code.

For instructions on how to install and license a development instance of InterSystems IRIS, see [Quick Start: InterSystems IRIS Installation](#).

Here is the sample code:

```
import java.sql.*;

public class JDBCSTest {
    public static void main(String[] str) throws Exception {
        String url = "jdbc:Cache://127.0.0.1:1972/USER";
        // To use shared memory, the connection string is:
        // String url = "jdbc:Cache://SHM|||||:1972/USER";

        Class.forName("com.intersys.jdbc.CacheDriver");
        Connection connection = DriverManager.getConnection(url, "_SYSTEM", "SYS");
        // Replace _SYSTEM and SYS with a username and password on your system

        String createTable = "CREATE TABLE People(ID int, FirstName varchar(255), LastName varchar(255))";

        String insert1 = "INSERT INTO People VALUES (1, 'John', 'Smith')";
```

```

String insert2 = "INSERT INTO People VALUES (2, 'Jane', 'Doe')";
String query = "SELECT * FROM People";

Statement statement = connection.createStatement();
statement.executeUpdate(createTable);
statement.executeUpdate(insert1);
statement.executeUpdate(insert2);
ResultSet resultSet = statement.executeQuery(query);
System.out.println("Printing out contents of SELECT query: ");
while (resultSet.next()) {
    System.out.println(resultSet.getString(1) + ", " + resultSet.getString(2) + ", " +
resultSet.getString(3));
}
resultSet.close();
statement.close();
connection.close();
}
}

```

If the connection and queries have completed successfully, you should see a console window containing the results of the SELECT query.

A Note on the Sample

The connection string syntax is:

```
jdbc:Cache://ipAddress:superserverPort/namespace
```

where:

- *ipAddress* — The IP address of the InterSystems IRIS instance. If you are connecting locally, use 127.0.0.1 instead of localhost.
- *superserverPort* — The superserver port number for the IRIS instance, which is not the same as the webserver port number. To find the superserver port number, in the Management Portal, go to **System Administration > Configuration > System Configuration > Memory and Startup**.
- *namespace* — An existing namespace in the InterSystems IRIS instance. In this demo, we connect to the **USER** namespace.

5 JDBC: For More Information

For more information on JDBC, SQL and InterSystems IRIS, and other Java interoperability technologies, see:

- [Using JDBC with InterSystems IRIS](#) — InterSystems documentation: step-by-step instructions for using JDBC
- [Using Java with InterSystems IRIS XEP](#) — InterSystems documentation: step-by-step instructions for using XEP
- [Using the InterSystems Hibernate Dialect](#) — InterSystems documentation: connecting to InterSystems IRIS using Hibernate
- [Using the InterSystems Spark Connector](#) — InterSystems documentation: using InterSystems IRIS as an Apache data source
- [Using InterSystems SQL](#) — InterSystems documentation: step-by-step instructions for using JDBC
- [InterSystems SQL Reference](#) — InterSystems documentation: SQL reference
- [SQL Optimization Guide](#) — InterSystems documentation: guide to optimizing InterSystems SQL performance
- [First Look: XEP and InterSystems IRIS](#) — InterSystems documentation: Java XEP First Look
- [Java Overview](#) — InterSystems online learning: introductory video
- [Hibernate releases](#) — Hibernate website's release information

- [Hibernate and JDBC compared](#) — Stack Overflow article