

What's in Oracle Database 21c for Java Developers?

New enhancements to JDBC, UCP and OJVM for
Designing and Deploying Mission Critical and Cloud
Native Java Applications.

December 2020
Copyright © 2020, Oracle and/or its affiliates
Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates. This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Table of contents

Disclaimer	2
Introduction	4
Support for Popular Java Frameworks and IDEs	4
Eclipse and IntelliJ* Plugins	4
Support for Popular Frameworks and Java EE App Servers	5
UCP as JBoss DataSource	6
Sample Servlet Retrieving the datasource through JNDI.	7
Here is a sample Servlet using the datasource.	7
Enhancements to support Cloud Native Applications	8
JDBC and UCP Configuration with MicroServices Frameworks	8
Loading Wallet From Non-File bases System	8
Security Enhancements in the Previous Release	9
JDBC Support for Native JSON Datatype	9
Diagnosability and Tracing Enhancements	11
Connection Identifier for Improved Diagnosability	11
DMS Metrics and ClientInfo	11
Performance and Scalability for Mission Critical Deployment	11
Driver Support for Virtual Threads	11
The Reactive Streams Ingestion Library	12
JDBC Reactive Extensions	13
Driver Configuration for GraalVM Native Image	13
Sharding Datasource	14
Zero Downtime for Mission Critical Deployment	15
Application Continuity (AC) & Transparent AC	15
JVM in the Database - High-Availability and Security Enhancements	15
RAC Rolling OJVM Patching	15
Security Enhancements	16
Conclusions	16
References	16

Note: A Table of Contents (TOC) is recommended for documents that are more than 10 pages in length. If your statement of direction is shorter, you may remove the TOC page. To remove both the TOC and the page that it appears on, first display hidden characters by clicking on the Paragraph symbol on the Home toolbar. It is a small, square icon that appears to the left of the Quick Style Gallery. Notice the page break displayed as a line at the bottom of this page. Next, highlight all the text on this page and press the Delete key once to remove all the text. Then, highlight the page break and press the Delete key until the page disappears and your cursor is on the first page of body text. The TOC will update semi-automatically with the Heading 1 and Heading 2 styles when you 1) right click on the TOC and then 2) click Update Field in the contextual menu. You may also manually edit the TOC by placing your cursor within the text. Delete this note before publishing.

Introduction

As a Java developer or architect, why would you consider the Oracle database 21c release?

The new enhancements in the RDBMS, the embedded JVM (a.k.a. OJVM), the JDBC drivers, the Java connection pool, and the Oracle Cloud Infrastructure aim at: simplifying the onboarding of those new to Oracle database; improving the experience of those familiar to it; and simplifying the development and deployment of mission critical and Cloud native Java applications.

This technical brief discusses how the new capabilities fulfill the goals outlined above; specifically:

1. The enhancements to Cloud database connectivity
2. The support for popular Java frameworks and IDEs for simplifying Java developers experience or onboarding with the Oracle database
3. The support for Cloud native applications (Data-driven MicroServices and Serverless functions)
4. The enhancements to diagnosability and tracing for improving developers experience and mission critical deployments
5. The performance and scalability enhancements for mission critical deployments including: the async/reactive extensions, support for virtual threads (project Loom), the Reactive Streams Ingestion library, support for GraalVM Native Image, and a new Sharding datasource
6. The Zero Downtime enhancements for mission critical deployments including: Transparent Application Continuity, Zero Brownout OJVM Rolling patching

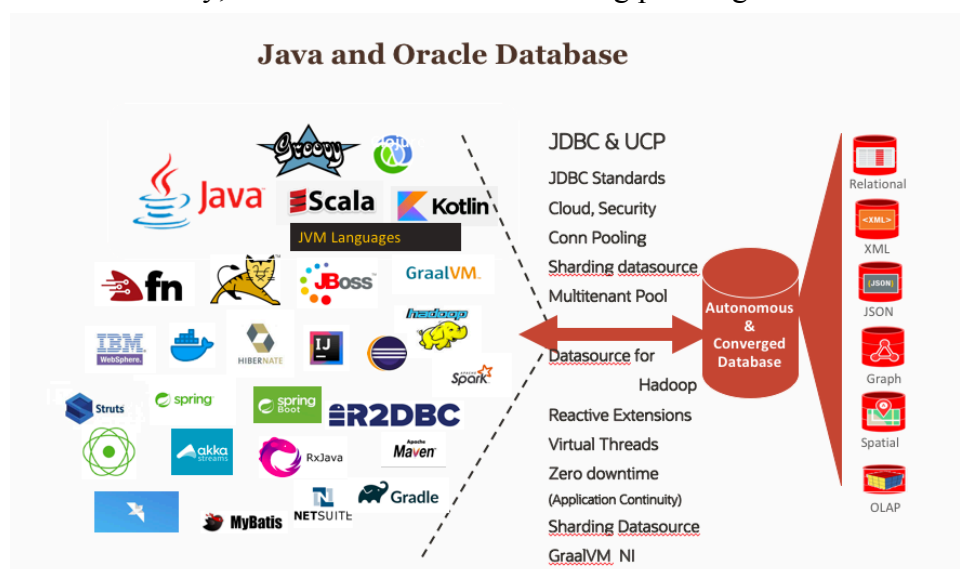


Figure 1 Global view of Java and the Oracle database

Support for Popular Java Frameworks and IDEs

There are tons of Java framework; we looked at the most popular and how to configure these to work efficiently with the Oracle JDBC driver and connection pool (UCP).

Eclipse and IntelliJ* Plugins

The Eclipse plugin is available @ <https://github.com/oracle/oci-toolkit-eclipse>. See the [related blog post](#) for more details.

* The IntelliJ plugin is in the process of being released, as of this write up.

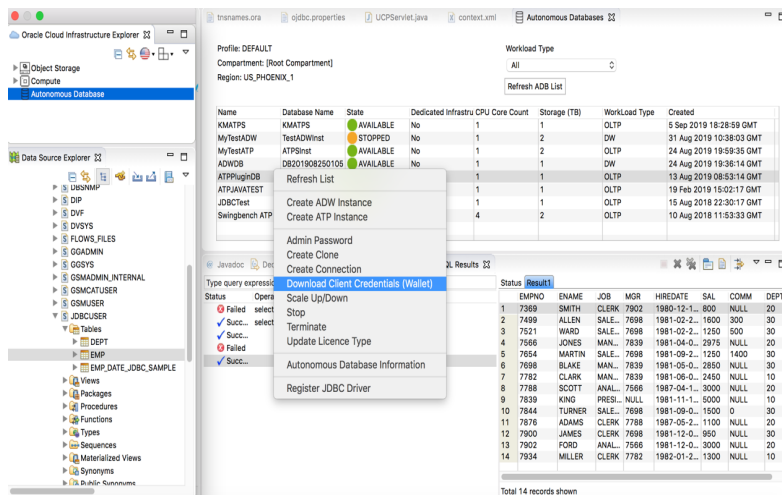


Figure 2 Oracle Cloud Infrastructure Toolkit for Eclipse

The features set for both plugins includes: provisioning new autonomous databases, start/stop/clone, scale up/down OCPU and Storage, change the ADMIN password; downloading the Cloud credentials; testing the connection to the Cloud database; browsing the database schema; and performing other database operations.

Support for Popular Frameworks and Java EE App Servers

1. Maven and Gradle are the more popular build automation frameworks used by Java developers. We now have Oracle JDBC on Maven Central @ <https://bit.ly/33bpLVJ>; the developers guide @ <https://bit.ly/2IBDXjJ> gives more details. Starting with the 19.8 release, we've also added the pre-established dependencies (a.k.a. "flavor POMs") to our BOM file so as to allow Sprint projects (i.e., Sprint Initializr) to pull these.
2. Watch [this video](#) for Apache Tomcat Servlet engine connectivity to the Autonomous database ATP.
3. Apache Hadoop: the [Oracle Datasource for Hadoop](#) turns Oracle database tables into Hadoop external tables.
4. Hibernate: see how to configure Hibernate and with UCP @ <https://bit.ly/32WrYUN>
5. MyBatis: see how to configure MyBatis and with @ <https://bit.ly/2IRgpGY>
6. Spring: the Java connection pool (UCP) can now be configured as a Spring datasource with no extra code. Spring retrieves the configuration from the application's properties file and autowires (injects) the values to the datasource

```
spring.datasource.url=jdbc:oracle:thin:@host:1521/mybservice
spring.datasource.driver-class-name=oracle.jdbc.OracleDriver
spring.datasource.type=oracle.ucp.jdbc.UCPDataSource
spring.datasource.ucp.connection-factory-class-name=oracle.jdbc.replay.OracleDataSourceImpl
spring.datasource.ucp.sql-for-validate-connection=select * from dual
spring.datasource.ucp.connection-pool-name=connectionPoolName1
spring.datasource.ucp.initial-pool-size=15
spring.datasource.ucp.min-pool-size=10
spring.datasource.ucp.max-pool-size=30
```

See also how to configure UCP with SpringBoot @ <https://bit.ly/3nDr6Mr>.

7. R2DBC: an open source Oracle-R2DBC driver will be available in the next couple of months.

8. GraalVM: the JDBC driver has been instrumented for GraalVM native image; see more details later in this write-up.

UCP as JBoss DataSource

In this release, the Java Universal Connection Pool (UCP) furnishes a class for its integration with JBoss for use by Java EE components (Servlets, JSP, JMS, EJB and so on).

- The class implements the `ServletContextListener` interface and supports method override `@Override public void contextInitialized(ServletContextEvent contextEvent)`.
- The class can either be annotated with `@WebListener` or added manually in `web.xml` using the `<listener>` and `<listener-class>` tags.
- The class is executed when the application server is started. The object reads a configuration/description file, creates a datasource with those values and binds it to a JNDI address and to an application object injected with CDI. The JNDI address and application object can be used by components to retrieve the datasource and connections. The parameters are retrieved either from the web descriptor (`web.xml`) as `<context-param>` values or using UCP XML configuration file.

Here is a working sample of the `ServletContextListener` implementation, a snippet of the web descriptor and its use from a Servlet.

Web descriptor snippet

```
<context-param>
    <param-name>ucp.jndiName</param-name>
    <param-value>java:/datasources/mypool_usingwl</param-value>
</context-param>
<context-param>
    <param-name>ucp.url</param-name>
    <param-value>jdbc:oracle:thin:@myhost:5521/myservice</param-value>
</context-param>
<context-param>
    <param-name>ucp.connectionFactoryClassName</param-name>
    <param-value>oracle.jdbc.replay.OracleDataSourceImpl</param-value>
</context-param>
<context-param>
    <param-name>ucp.dataSourceName</param-name>
    <param-value>myDataSource</param-value>
</context-param>
<context-param>
    <param-name>ucp.user</param-name>
    <param-value>scott</param-value>
</context-param>
<context-param>
    <param-name>ucp.password</param-name>
    <param-value>*****</param-value>
</context-param>
. . .
</web-app>
```

Using UCP XML Configuration

To use the UCP XML Configuration file, set the `oracle.ucp.jdbc.xmlConfigFile` JVM system property as follows:

```
-Doracle.ucp.jdbc.xmlConfigFile=file:/Users/scott/conf/ucp_config.xml
```

Then set the `ucp.dataSourceName` parameter in the web descriptor (`web.xml`); it must match the `data-source-attribute` in `/ucp-properties/connection-pool/data-source/data-source-name`.

Web.xml

```
<context-param>
    <param-name>ucp.dataSourceName</param-name>
    <param-value>myDataSource</param-value>
</context-param>
```

ucp_config.xml

```
<ucp-properties>
    <connection-pool
        connection-factory-class-name="oracle.jdbc.replay.OracleDataSourceImpl"
        connection-pool-name="pool1"
        initial-pool-size="10"
        max-connections-per-service="15"
        max-pool-size="30"
        min-pool-size="2"
        password="*****"
        url="jdbc:oracle:thin:@myhost:5521/myService"
        user="scott">
        <connection-property name="autoCommit" value="false"></connection-property>
        <connection-property name="oracle.net.OUTBOUND_CONNECT_TIMEOUT" value="2000">
        </connection-property>
        <data-source data-source-name="myDataSource" description="pdb1" service="ac">
        </data-source>
    </connection-pool>
</ucp-properties>
```

Sample Servlet Retrieving the datasource through CDI.

```
@Inject
@UCPResource
private DataSource ds;
```

Sample Servlet Retrieving the datasource through JNDI.

```
@WebServlet("/OracleUcp")
public class OracleUcp extends HttpServlet {
    private DataSource ds = null;

    // Retrieve Datasource reference using JNDI
    @Override
    public void init() throws ServletException {
        Context initContext;
        try {
            initContext = new InitialContext();
            ds = (DataSource)
                initContext.lookup("java:/datasources/mypool_usingwl");
        } catch (NamingException e) {
            // ...
        }
    }
}
```

Here is a sample Servlet using the datasource.


```
protected void doGet(HttpServletRequest request, HttpServletResponse response)
    throws ServletException, IOException {

    // Retrieve connection from the pool
    try (Connection conn = ds.getConnection();
        Statement st = conn.createStatement()) {
        ResultSet rs = null;
        rs = st.executeQuery("select empno, ename, job from emp");
    }
}
```

Known limitation: the management and monitoring tools provided from the container cannot be applied to the pool

Enhancements to support Cloud Native Applications

JDBC and UCP Configuration with MicroServices Frameworks

- **SpringBoot:** See how to configure the Oracle JDBC drivers in a SpringBoot project @ <https://bit.ly/3m4zQuH> and how to configure UCP with Spring @ <https://bit.ly/3nDr6Mr>
- **Helidon:** both Helidon MP and SE provide support for Oracle UCP datasources thus inheriting all of the features it provides. Helidon provides various integration and convenience features from automatically injecting UCP datasources into microservices to simplified use in Kubernetes and ATP environments to use of Oracle DB features (including those directly related to micro services architectures).
See how to configure the Oracle JDBC drivers in a Helidon project @ <https://bit.ly/2IP18a5>
- **Micronaut:** See how to configure the Oracle JDBC drivers in a Micronaut project @ <https://bit.ly/2KEw4KX>
- **Quarkus:** See how to configure the Oracle JDBC drivers in a Quarkus project @ <https://bit.ly/2J6dYR9>

Loading Wallet From Non-File bases System

In this release, a new API for setting an SSLContext on an OracleDataSource gives developers full control over how to load wallets i.e., from memory or non-file based system.
The following code snippets illustrates how to use such API.


```

static SSLContext createSSLContext()
    throws GeneralSecurityException, IOException {
    TrustManagerFactory trustManagerFactory =
        TrustManagerFactory.getInstance("PKIX");
    KeyManagerFactory keyManagerFactory =
        KeyManagerFactory.getInstance("PKIX");

    trustManagerFactory.init(loadKeyStore());
    keyManagerFactory.init(loadKeyStore(), null);

    SSLContext sslContext = SSLContext.getInstance("SSL");
    sslContext.init(
        keyManagerFactory.getKeyManagers(),
        trustManagerFactory.getTrustManagers(),
        null);
    return sslContext;
}

static KeyStore loadKeyStore()
    throws IOException, GeneralSecurityException {

    // This example is using a standard file-based source however, I can be reused
    // to load the keystore using an InputStream from a non-file based source.
    try (InputStream keyStoreStream =
        Files.newInputStream(Paths.get("cwallet.sso"))) {
        KeyStore keyStore = KeyStore.getInstance("SSO", new OraclePKIProvider());
        keyStore.load(keyStoreStream, null);
        // keyStore.load(KeyStore.LoadStoreParameter) could be used here as well.

        return keyStore;
    }
}

```

Security Enhancements in the Previous Release

The following security enhancements were made in the previous release.

- Support for HTTPS Proxy Configuration
- Automatic Provider Resolution (OraclePKIProvider)
- Setting server's domain name (*oracle.net.ssl_server_cert_dn*) in the URL
- Support of new wallet property (*my_wallet_directory*) in the URL
- Support for Key Store Service (KSS)

JDBC Support for Native JSON Datatype

The Oracle database 21c release furnishes a native JSON datatype and an [Autonomous JSON Cloud service](#). The `oracle.sql.json` package furnishes Java APIs for: accessing the native JSON type values and their storage in binary format; creating, querying and modifying JSON type values; encoding or decoding JSON type values in the same binary JSON format used by the database; converting JSON type values to and from JSON text; binding and accessing JSON type values using the JSON-P interfaces e.g., `javax.json.*`.

Here is a JDBC code sample illustrating the use of JSON Native datatype.

```

import java.sql.PreparedStatement;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;

import oracle.jdbc.OracleConnection;
import oracle.jdbc.pool.OracleDataSource;
import oracle.sql.json.OracleJsonFactory;
import oracle.sql.json.OracleJsonObject;

public class JsonExample {

    public static void main(String[] args) throws SQLException {
        OracleDataSource ds = new OracleDataSource();
        ds.setURL("jdbc:oracle:thin:@myhost:1521:orcl");
        ds.setUser(<user>);
        ds.setPassword(<password>);
        OracleConnection con = (OracleConnection) ds.getConnection();

        // create a table with a JSON column and insert one value
        Statement stmt = con.createStatement();
        stmt.executeUpdate("CREATE TABLE fruit (data JSON)");
        stmt.executeUpdate("INSERT INTO fruit VALUES ('{\"name\":\"pear\",\"count\":10}')");

        // create another JSON object
        OracleJsonFactory factory = new OracleJsonFactory();
        OracleJsonObject orange = factory.createObject();
        orange.put("name", "orange");
        orange.put("count", 12);

        // insert the orange object
        PreparedStatement pstmt = con.prepareStatement("INSERT INTO fruit VALUES (:1)");
        pstmt.setObject(1, orange, OracleType.JSON);
        pstmt.executeUpdate();
        pstmt.close();

        // retrieve the pear object
        ResultSet rs = stmt.executeQuery("SELECT data FROM fruit f WHERE f.data.name =
'pear'");
        rs.next();
        OracleJsonObject pear = rs.getObject(1, OracleJsonObject.class);
        int count = pear.getInt("count");

        // create a modifiable copy of the pear object
        pear = factory.createObject(pear);
        pear.put("count", count + 1);
        pear.put("color", "green");

        // update the pear object
        pstmt = con.prepareStatement("UPDATE fruit f SET data = :1 WHERE f.data.name =
'pear'");
        pstmt.setObject(1, pear, OracleType.JSON);
        pstmt.executeUpdate();
        pstmt.close();

        rs.close();
        stmt.close();
        con.close();
    }
}

```

Diagnosability and Tracing Enhancements

The Connection Identifier is a new diagnosability enhancements in this release.

Connection Identifier for Improved Diagnosability

The RDBMS allows assigning a connection id to each connection or group of connections.

1. As a name/value pair in the JDBC connect string (CONNECTION_ID =<value>)

```
(description =(address =(protocol=)( port =)(host=))
  (connect_data=(service_name=)(..)(CONNECTION_ID=<value>))
)
```

2. The connection identifier (CONNECTION_ID) will appear in logs

```
“Exception in thread "main" java.sql.SQLRecoverableException: ORA-
12506, TNS:listener rejected connection based on service ACL filtering
(CONNECTION_ID=ovDT/bCGfJngU602xAph8g==)”
```

DMS Metrics and ClientInfo

When present in the classpath, the Dynamic Monitoring System (DMS) jars (ojdbc8dms.jar, ojdbc11dms.jar) furnish diagnosability metrics and limited support for java.util.logging.

JDBC furnishes setClientInfo() and getClientInfo() methods for tagging applications with the end-to-end metrics including: Actions, ClientId, ExecutionContextId, Module and State.

Performance and Scalability for Mission Critical Deployment

The following new capabilities will significantly improve the performance and scalability of mission critical Java applications: the JDBC reactive extensions, the support for virtual threads in the driver, the reactive streams ingestion library, driver configuration for GraalVM native image, and the Sharding datasource.

Driver Support for Virtual Threads

[Project Loom](#) aims at “reducing the complexity of creating and maintaining the high-throughput concurrent applications” using inexpensive “virtual threads”, delimited continuations and tail-call elimination. The promise of virtual threads is that synchronous calls can perform as well as asynchronous calls.

DB21c JDBC has been instrumented (native method calls, intrinsic locking with synchronized) to support synchronous database access with inexpensive threads i.e., standard JDBC calls using virtual threads. In other words, no need to refactor simple blocking JDBC calls into complex reactive streams calls.

Here is a code snippet illustrating the use of Virtual Threads with JDBC

```
// Set the task executor to use Virtual Threads, or to use kernel threads.
final Executor taskExecutor =
    useVirtualThreads
    ? newVirtualThreadExecutor(kernelThreadExecutor)
    : kernelThreadExecutor;
...

// Execute a number of JDBC calls with the executor
for (int i = 0; i < taskCount; i++) {
    taskExecutor.execute(() -> {
        executeSql("SELECT * FROM emp");
        completionLatch.countDown();
    });
}
...

// Periodically prints a message until all tasks are complete
awaitCompletion(useVirtualThreads, taskCount, completionLatch);
```

The Reactive Streams Ingestion Library

The DB21c JDBC driver introduces a new Java library for high speed ingestion of massive volumes of data (sensors, Call Detail Records, logs) from large numbers of concurrent sources into the Oracle Database.

The library groups and buffers incoming data streams then gives control back to the client threads while using a separate thread pool for database I/O, using the Direct Path insert into database blocks, i.e., bypassing the SQL layer.

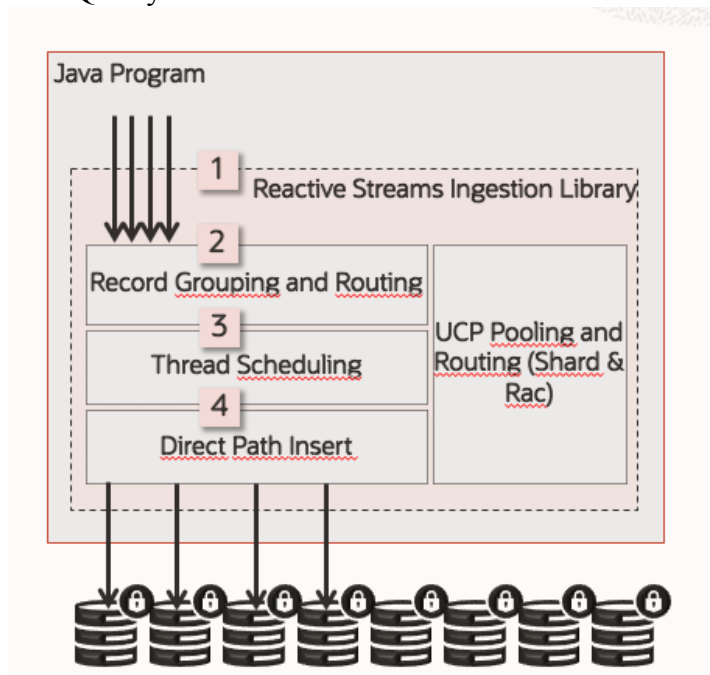


Figure 3 Reactive Streams Ingestion

The library (`rsi.jar`) furnishes the following simple APIs:

- Push Publisher: simplest usage, no backpressure.
- Flow Publisher: implementing Java Flow API (Publishers, Subscribers and Subscriptions).
- Record API for Object Relational Mapping

The following [short video](#) illustrates RSI in action. See the [Oracle JDBC doc](#) for more details and code fragments.

JDBC Reactive Extensions

The DB21c JDBC driver has been extended to support asynchronous database access with non-blocking network I/O. The extensions expose standard Java Flow Subscriber and Publisher types and APIs that can interoperate with the R2DBC API¹ and the reactive streams libraries including: Reactor, RxJava, Akka, Vert.x, Spring, jOOQ, Querydsl, Kotysa and so on which furnish operators (map, reduce, filters), concurrency modeling, monitoring, and tracing.

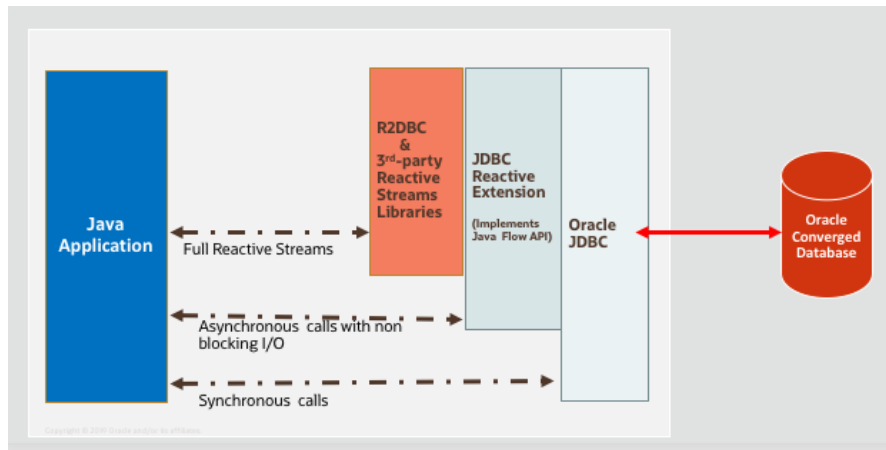


Figure 4 Database Access with Oracle JDBC 21c

The following code snippet shows how to open a connection, asynchronously.

```
/**
 * Asynchronously opens a new connection
 * @param dataSource DataSource configured with a URL, User, and
 * Password
 * @return A Publisher that emits a single connection
 * @throws SQLException If a database access error occurs before the
 * connection can be opened
 */
Flow.Publisher<OracleConnection> openConnection(DataSource dataSource)
    throws SQLException {
    return dataSource.unwrap(OracleDataSource.class)
        .createConnectionBuilder()
        .buildConnectionPublisherOracle();
}
```

See [the Oracle JDBC doc](#) for more details.

Driver Configuration for GraalVM Native Image

[GraalVM](#) speeds up mission critical applications (faster startup, low runtime memory, and so on). The DB21c JDBC driver have been instrumented with GraalVM configuration in META-INF/native-image. Additional code in the driver relates to the orai18n and xmlparserv2 companion jars used respectively for NLS or Internalization support, and XML parsing.

The following steps allows running a plain JDBC code [DataSourceSample.java](#) as a native image.

¹ The Oracle R2DBC driver will be available soon on Github.

- 1) Install Native Image with the GraalVM Updater tool

```
gu install native-image
```

- 2) Modify the DB_URL at line 40 to point to your database then compile the java file with the proper classpath:

```
javac -cp ../ojdbc11.jar DataSourceSample.java
```

- 3) Run the Native Image builder

```
native-image -cp ../ojdbc11.jar DataSourceSample
```

- 4) Run the image (the JVM is no longer needed at this stage):

```
./datasourcesample
```

In addition to instrumenting the driver, the JDBC dev team has been actively working with the Helidon, Micronaut, and Quarkus teams to integrate the driver and companion jars with their GraalVM native images.

Sharding Datasource

[Java SE 9](#) JDBC 4.3 introduces the `ShardingKey` and the `ShardingKeyBuilder` interfaces for creating a `ShardingKey` however, existing applications must be modified to use these new APIs. [This is a problem for many Java developers.

The DB21c JDBC driver introduces a new JDBC datasource for simplifying Java connectivity to Sharded databases.

It removes the requirement to explicitly set the Sharding key before requesting a connection. The SQL statements intended for a single shard must contain the Sharding key in the `WHERE` clause.

Example: `select id, name from customer where id = ?`

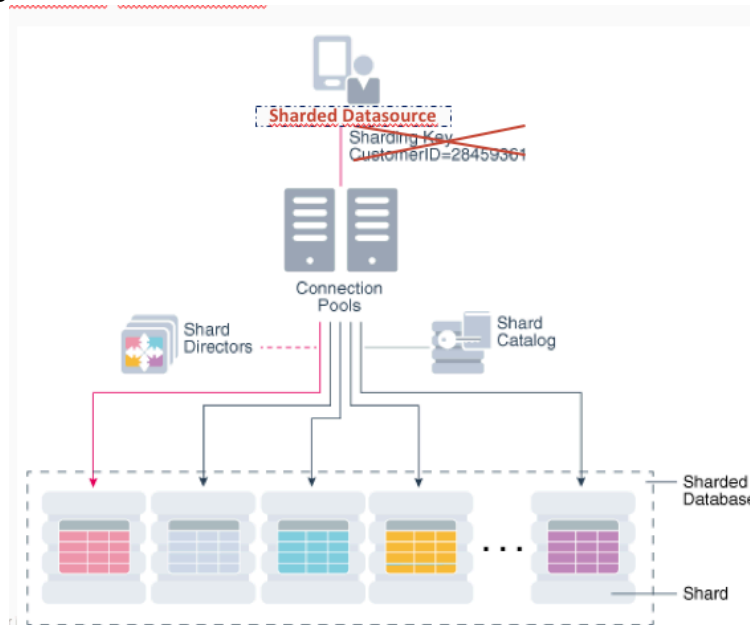


Figure 5 Sharding Datasource

Here are the requirements and best practices for using the Sharding datasource:

- 1) Java developers must set the `oracle.jdbc.useShardingDriverConnection` system property to true (in code or at the JDK level)

```
Properties prop = new Properties();  
prop.setProperty("oracle.jdbc.useShardingDriverConnection", "true");
```
- 2) Java developers must use the Shard director (a.k.a. GSM) connect string; it routes connections to the appropriate Shard

```
final static String gsmURL = "jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (HOST =
```



```
...)(PORT = ...)(PROTOCOL = tcp))(CONNECT_DATA = (SERVICE_NAME = ...)))";
```

- 3) The Sharding datasource supports local transaction against a single shard when Auto-COMMIT is OFF however, the Java developer must set the datasource property `allowSingleShardTxn` to TRUE to indicate that local transactions always start and end on a single Shard.

The following are the known limitations as of DB21c

- The Sharding datasource supports only the type 4 JDBC Thin driver. It does not support the type 2 JDBC OCI driver or the RDBMS server-side type 2 driver (a.k.a. KPRB driver).
- The Sharding data source does not currently support transactions spanning multiple shards.
- The Sharding datasource does not support some of the Oracle JDBC extensions such as the Direct Path load, and JDBC Dynamic Monitoring Service (DMS) metrics.

The following video shows the Sharding datasource in action.

Zero Downtime for Mission Critical Deployment

Zero downtime is an essential *Quality of Service* for mission critical applications. Application Continuity (AC), Transparent AC (TAC) and Transaction Guard are major zero downtime mechanisms for client Java applications. Java code running in the database also benefit from the zero-brownout patching and security enhancements.

Application Continuity (AC) & Transparent AC

Application Continuity (AC) consists in recording, in driver memory, all database calls made during an application unit of work (typically a transaction) and replaying these calls—upon the failure of an instance, host or network—against another instance of the same database, after hiding the exception. Upon replay, if the outcome is identical (i.e., result set checksum) then AC is successful and the control is given back to the application to continue, as if nothing has happened; if not, the exception is re-cast and visible the Java code. I described AC and TAC from Java developer's perspective in the Zero-Downtime section of [this blog post](#).

In this release, a new attribute named `RESET_STATE` of the database service has been added – this is transparent to your Java code -- cleans states set/used by the application when the unit of work i.e., transaction, completes. Specifically, cursors are cancelled, PL/SQL globals are cleared, temporary session tables are truncated, and temporary session lobes are cleared.

JVM in the Database - High-Availability and Security Enhancements

The JVM in the database (a.k.a. OJVM) is used for: in-place data processing; calling out Web-Services, Hadoop servers, 3rd party databases and legacy systems; running 3rd party Java libraries; running Java-based languages (e/g., Jython, Groovy Kotlin, Clojure, Scala, JRuby). OJVM is also used by database components such as AQ – JMS, XDB, Spatial, Scheduler, Java XA and OLAP.

This release furnishes the long-awaited zero brownout and rolling patching of OJVM in clustered database (RAC) environments and additional security enhancements.

RAC Rolling OJVM Patching

This release furnishes zero brownout and RAC rolling OJVM patching. RAC Rolling patching of OJVM was possible in the previous releases (18c and 19c) however, there was a few seconds brown-out period during which the Java is not available in the entire cluster.

Security Enhancements

A *lockdown profile* is a mechanism for restricting certain operations or functionalities in pluggable database (PDB) and the container database (CDB). In the previous releases, the following lockdown profiles were implemented: disable OS file access and networking to/from Java (in DB 19c release) and disabling granting `java.lang.RuntimePermission` and `java.io.FilePermission`.

In this release, the following lockdown profiles are implemented: confining OS file accesses in paths within `PATH_PREFIX`; the ability to specify OS user identity when forking OS process.

Conclusions

This technical brief walked you through the new and recent enhancements in the Oracle database 21c that simplify the onboarding or the experience of Java developers and architects. These new enhancements also simplify the development and deployment of mission critical and Cloud native Java applications.

The reference section hereafter furnishes links to the developer guides and javadocs. Our landing page <https://www.oracle.com/jdbc/> furnishes up to date resources.

References

[JDBC Developer's Guide](#)

[Java Developer's Guide](#)

[Universal Connection Pool Developer's Guide](#)


[JDBC Java API Reference](#)


[Universal Connection Pool API Reference](#)


[RAC FAN Events Java API Reference](#)

Connect with us

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com). Outside North America, find your local office at: [oracle.com/contact](https://www.oracle.com/contact).

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2020, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail REVREC_US@oracle.com.