



Red Hat JBoss Data Virtualization 6.3 Release Notes

Errata and late-breaking news for this release.

Documentation Team

Red Hat JBoss Data Virtualization 6.3 Release Notes

Errata and late-breaking news for this release.

Documentation Team
dlesage@redhat.com

Legal Notice

Copyright © 2017 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Table of Contents

1. BACK UP YOUR DATA	2
2. RED HAT DOCUMENTATION SITE	2
3. NEW FEATURES	2
3.1. JDV 6.3	2
4. KNOWN ISSUES	3
5. RESOLVED ISSUES	4
6. DEPRECATED FEATURES	6
7. WORKAROUNDS FOR ONGOING ISSUES	6

1. BACK UP YOUR DATA

Warning

Red Hat recommends that you back up your system settings and data before undertaking any of the configuration tasks mentioned in this book.

2. RED HAT DOCUMENTATION SITE

Red Hat's official documentation site is available at <https://access.redhat.com/site/documentation/>. There you will find the latest version of every book.

3. NEW FEATURES

3.1. JDV 6.3

- ✦ Supports JDK 6, 7 and 8
- ✦ In-memory technologies for Big Data processing
 - Apache Spark
 - JDG as materialization target for JDV
 - SAP HANA
 - HPE Vertica
- ✦ Development and Deployment productivity
 - OData V4
 - OAuth using RH-SSO
 - Teiid Designer usability improvements
 - VDB Builder CLI (technical preview only)
- ✦ Big Data and Cloud data sources
 - Apache Cassandra
 - Apache HBase
 - Amazon RedShift
 - Apache Accumulo
 - Geospatial Support
 - OSISoft PI (technical preview only)

4. KNOWN ISSUES

The following issues are known to exist in JBoss Data Virtualization 6.3.0:

Bugzilla Ticket 1168636 - When a user, who has not been assigned the appropriate role, tries to login to JBoss Data Virtualization Dashboard by entering the URL directly, they are redirected from the JBoss Data Virtualization Dashboard login page to a blank screen.

Bugzilla Ticket 1272955 - The Phoenix driver evaluates the clause `HAVING MIN/MAX/AVG(x) = null` to true. As a result, queries may return incorrect results. This is due to a bug in the driver.

Bugzilla Ticket 1272934 - There is a bug in the Phoenix driver whereby if you submit a query such as: `SELECT intkey FROM BQT1.SmallA AS A WHERE stringkey = (SELECT stringkey FROM BQT1.SmallA AS B WHERE b.intnum = a.intnum) ORDER BY intkey`, more than one row will be returned. To work around this issue, change the query to this: `WHERE b.intnum is not null and b.intnum = a.intnum` returns correct result.

Bugzilla Ticket 1272886 - There is a bug in the Phoenix driver that causes queries using `MAX(charvalue)` to be evaluated incorrectly and return null results.

Bugzilla Ticket 1272912 - The user may encounter an exception when reading the results of some queries from HBase using the Phoenix JDBC driver.

Bugzilla Ticket 1272862 - If a timestamp value is used in combination with the `ALL` predicate then an `IllegalArgumentException` is throw by the Phoenix driver. This occurs when there is a null timestamp value or when the subquery returns no rows. To work around this problem, avoid using the `ALL` predicate.

Bugzilla Ticket 1246153 - If a mongodb user has insufficient permissions to import metadata, the importation process will fail. The user needs adequate permissions to access the mongodb database. If there is a collection which cannot be accessed by the user, metadata import will fail. Currently, there is no way to exclude those tables from import.

Bugzilla Ticket 1361639 - If you define multiple resource adapters within the same resource-adapter ID you will be unable to use Teiid Designer's Teiid Connection import option to model data source. To work around this problem, ensure that each resource-adapter is defined within its own unique resource-adapter ID section. If you do so, the importer option can be used on each defined resource.

Bugzilla Ticket 1309067 - If you use the Designer plug-in to create a data source, remove that data source and then try to create a new data source of the same type, you will encounter an error. This is because Red Hat JBoss EAP expects a reload after the resource adapter is removed.

Bugzilla Ticket 1242938 - Null values are not sorted correctly when the Infinispan-dsl-cache translator is used.

Bugzilla Ticket 1278724 - The Informix translator adjusts time values according to the default server's timezone or the `DatabaseTimeZone` translator property (the translator property takes precedence). However date and timestamp values are adjusted only according to the `DatabaseTimeZone` translator property. Timestamp values are adjusted only if the `DatabaseTimeZone` property is set. Moreover, time values are adjusted in the "opposite" direction to date/timestamp values.

Bugzilla Ticket 1279473 - There is an issue with the Informix translator whereby neither `'<> ALL'` or `'NOT IN'` works in a query. You must avoid using these statements in your commands.

Bugzilla Ticket 1260621 - If a data source is configured to use the `RoleBasedCredentialMapIdentityLoginModule`, then the user will encounter an exception when launching the product. This is because the default username and password are null.

Bugzilla Ticket 1331742 - The MSSQL JDBC driver invalidates an active kerberos ticket on `Connection.close()`. As a result, if the user creates kerberos connection, the driver invalidates the ticket when the connection is closed and, therefore, the ticket cannot be re-used. The EAP team created a workaround for

this by adding the module option 'wrapGSSCredential=true' with the additional setting 'credentialLifetime=-1'. This works for static kerberos authentication. To make the PassthroughIdentityLoginModule (PTILM) work, you need to add an additional module option for PTILM 'wrapGSSCredential' and set it to 'true' (the default is 'false').

Bugzilla Ticket 1242999 - The Infinispan-dsl-cache translator will by default have support for NOT and NULL criteria set to false. Therefore, when NOT is used in the criteria, or Is Null or Is Not Null are used in the criteria, it will not be pushed down to JDG for processing, but will have the conditions applied by the Data Virtualization engine.

Bugzilla Ticket 1263571 - If a user wants to use a different data source for the dashboard builder, he or she must add a JDBC driver module name in the jboss-deployment-structure.xml configuration file in the teiid-dashboard-builder.war. Failure to do so will result in an exception.

Bugzilla Ticket 1345737 - The installer fails occasionally when the server is being started with the standalone-full-ha.xml profile. The step on which the installer fails is called the "Starting Process Welcome Page".

Bugzilla Ticket 1208263 - The EAP Management Console refers to the JBoss Enterprise Application Platform instead of Red Hat JBoss Data Virtualization under the navigation panel.

Bugzilla Ticket 1270336 - If you try to deploy a virtual database after a dynamic virtual database has already been deployed, it will fail. To work around this problem, avoid mixing VDB types.

Bugzilla Ticket 1098405 - It is not possible to update a VDB which has been imported as a source model using the Designer plug-in as the "Update (if existing model selected)" check box is not visible for imported virtual databases being used as source models.

Bugzilla Ticket 1323539 - The ODBC driver installation for RHEL does not provide iSQL to test the ODBC connection.

Bugzilla Ticket 1265407 - If the user tries to install the product in console mode or by using an automatic script, they may encounter an exception if they are logged in as the root user. To work around this issue, ensure you log in as a JBoss user.

Bugzilla Ticket 1270336 - If you try to deploy a virtual database after a dynamic virtual database with exactly the same name has already been deployed, it will fail. To work around this problem, avoid mixing VDB types.

Bugzilla Ticket 1342917 - A warning message informing the user that "The use of the Modeshape JCR APIs have been deprecated" appears on the installer's package selection screen even if the user has not selected Modeshape for installation.

Bugzilla Ticket 1351093 - The name of the table in Red Hat JBoss Data Virtualization must match the name of the class name of the Pojo. If not, when executing updates, the table will not be found. Ensure that the name of the Table and class name of the Pojo must match.

Bugzilla Ticket 1355777 - If you run reload several times on a server started with Java 6 or 7, you will encounter an OutOfMemoryError. To work around this issue, you can remove the Dashboard Builder component if you are not using it.

Bugzilla Ticket 1381229 - If the user issues a multi-value INSERT SQL statement, they will encounter a null-pointer exception.

Bugzilla Ticket 1401324 - An implicitly-started xa-transaction which uses multiple non-xa datasources will produce a warning that 'Adding multiple last resources is disallowed' and then throw an exception 'Unchecked throwable in managedConnectionReconnected'. To avoid this, define your datasources as xa-datasources.

5. RESOLVED ISSUES

The following issues have been resolved in JBoss Data Virtualization 6.3.0.

Bugzilla Ticket 1244149 - The README.md files for the quick starts contained outdated information in some cases. This data has now been updated so the files are now therefore accurate.

Bugzilla Ticket 1022650 - The 'drools-integration' quick start hadnot been included with JBoss Data Virtualization. It needed to be updated for the latest version of the rules management software before it can be included in the product. It has now been updated.

Bugzilla Ticket 1018996 - There was no indication that a failed install would be reported to Red Hat. Furthermore, the content in the install log couldnot be copied into the clipboard to email details manually. The data is now captured in a log file correctly.

Bugzilla Ticket 1237049 - In high load/high contention cases, the client had to change the ISPN configuration and increase the locking timeout. A fix to EAP has now resolved this issue and the user should no longer need to change the ISPN configuration.

Bugzilla Ticket 1226323 - If an insert statement was used with an object value on Postgres, it would fail. As a result, the user would encounter an exception. A code change has been applied so that this no longer causes a problem.

Bugzilla Ticket 1226824 - Teradata did not support Cast statements. If one was used, an exception was thrown. These statements are now supported so the user should no longer see these exceptions.

Bugzilla Ticket 1227214 - When the convert(... , float) and convert(..., double) functions were used in complex operations, the translator did not take into account the priority of the operators. This has now been fixed so that the priority is accounted for.

Bugzilla Ticket 1234270 - The get-vdb command's failure description was lacking a status. Therefore, it did not give the user all of the information they needed when a problem was encountered obtaining a VDB. This has been corrected so that it now provides the user with all of the required information.

Bugzilla Ticket 1240565 - Red Hat JBoss Data Virtualization shifted date/time/timestamp values returned from Teradata according to the user.timezone value. This has been corrected so that both simple SELECT and time-defined parameters in WHERE now work.

Bugzilla Ticket 1243466 - The Teiid ODBC driver handled errors inconsistently. As a result, log information was sometimes incomplete. This has now been corrected. As a result, the log should contain all of the expected error information.

Bugzilla Ticket 1245204 - If an INSERT statement contained an object value execution would fail and users would encounter an exception. This has now been corrected so that users can add object values to INSERT statements.

Bugzilla Ticket 1247976 - Red Hat JBoss Data Virtualization handled nulls incorrectly in the UCASE,LCASE and SUBSTRING functions when working with MongoDB. It preserved Mongo's results, which were different from what was expected from the SQL function. This has now been fixed so that the nulls are handled correctly and users can expect to see the correct results from the SQL.

Bugzilla Ticket 1252307 - When using the HBase translator, Teiid's INSERT and UPDATE statement had no effect. Both statements wee translated to UPSERT. This has now been corrected so that both INSERT and UPDATE statements work correctly.

Bugzilla Ticket 1253602 - When users tried to import Salesforce models with the Teiid Connection, they encountered log errors. This was due to a property not being renamed. This has now been corrected by a coding change so that users can import Salesforce models without error.

Bugzilla Ticket 1259592 - If a table did not contain an element with a required key, then the user would encounter a null-pointer error. This has now been fixed by a coding change so users will no longer encounter this error.

Bugzilla Ticket 1260621 - If a data source was configured to use the RoleBasedCredentialMapIdentityLoginModule, then the user would encounter an exception on launching the product. This was because the default username and password were null. This has now been corrected. Note that the RoleBasedCredentialMapIdentityLoginModule is now deprecated.

Bugzilla Ticket 1231769 - If a Solr Translator PreparedStatement.executeBatch() had more than one query the user would encounter an error. This has now been resolved so that users can supply more than one query and no error will occur.

Bugzilla Ticket 1231613 - The ODBC port number was not written to the log during shutdown. This has now been corrected so that the log now captures this information.

Bugzilla Ticket 1180144 - If the user attempted to join data from a MongoDB table that had the teiid_mongo:MERGE property enabled, they would encounter a java.util.EmptyStackException. This has now been corrected so that users can undertake joins without encountering an exception.

Bugzilla Ticket 1246153 - If a MongoDB user had insufficient permissions to import metadata, the importation process would fail. The user needed adequate permissions to access the MongoDB database. This has been corrected so that users no longer encounter a null-pointer exception.

6. DEPRECATED FEATURES

Please note that OData 2 is deprecated in this release.

For customers migrating from 5.3.1, please note that, if you are using custom UPDATE, INSERT or DELETE procedures in view transformations, that the following special variables are no longer valid: TRANSLATE CRITERIA, HAS CRITERIA, INPUT Variables, ROWS_UPDATED Variable and Update procedures defined by "CREATE PROCEDURE."

The ability to use Lucene to search a JDG cache from within the infinispan-cache translator has been deprecated. This is because the JDG DSL query language provides most of the features that match SQL querying. Therefore, there is no need to support two underlying querying options.

7. WORKAROUNDS FOR ONGOING ISSUES

The following are workarounds for ongoing issues known to exist in JBoss Data Virtualization 6.

BZ#1311213 - EDS 5.3 jdbc driver virtual procedure call against JDV 6.2.3

If you use a driver from a prior major release against a new major release you may see errors due to compatibility issues. Make sure you upgrade your drivers so that they are compatible with the current version of Red Hat JBoss Data Virtualization.

BZ#996550 - Sybase15 subquery using HAVING clause failing

Compatibility, Documentation, Integration, Teiid - #In SYBASE 15, if a query is submitted with a subquery that contains the HAVING clause, it will fail. You can work around this by using a custom translator to set supportsOnlyCorrelatedSubqueries to true, which will disable a non-correlated subquery pushdown such as this. However similar issues may still exist with other subqueries.

BZ#1035035 - Unable to create function from the reconciler

Tooling - #When building functions in Teiid Designer, users are not able to create recursive functions. The workaround is:

1. Cancel out of the reconciler dialog
2. Place your cursor in the SQL near the columns that need reconciling
3. Select the "Expression Builder" action on the Transformation Editor toolbar
4. Create/edit your complex/nested function definition

BZ#1035480 - Changing connection profile jar then reset CP on Model does not update the jar

Tooling - #When using Teiid Designer, users are finding that changes to the list of driver JARs for a connection profile are not being displayed when the user goes to "View Connection Info" for the model. To workaroud this issue, after making changes to the list of driver JARs, perform one of the following options:

1. Restart JBoss Developer Studio
2. Create a new connection profile

BZ#1312196 - If you configure a Red Hat JBoss Data Grid DSL resource adapter, you still need to use the marshaller and protobuf when using annotations to define a schema

To see how to use an annotated POJO, please refer to this section of the JDG documentation:
https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Data_Grid/6.6/html-single/Infinispan_Query_Guide/index.html#Defining_Protocol_Buffers_Schemas_With_Java_Annotati

Q: Where is the documentation?

A: You can find the Red Hat JBoss Data Virtualization documentation at <https://access.redhat.com/site/documentation/>

You can also find many articles about specific use cases at the Red Hat Customer Portal Knowledge Base: <https://access.redhat.com/knowledge>

Javadoc packages are available for download with the software at the Red Hat Customer Portal: <https://access.redhat.com/jbossnetwork>

Red Hat recommends reading the Getting Started Guide first as it contains instructions for installing and configuring the product.

Q: Where are the installation instructions?

A: Complete installation instructions for JBoss Data Virtualization can be found in the Installation and Configuration Guide at <https://access.redhat.com/site/documentation/>

Q: Which operating systems, Java Virtual Machines and database servers are supported?

A: For a complete list of the supported operating systems, Java Virtual Machines (JVMs) and database servers on which you can run JBoss Data Virtualization, refer to <https://access.redhat.com/support/configurations/jboss/>

Q: Why isn't the included Hypersonic database supported?

A: The default configuration includes the embedded Hypersonic database. This configuration is included only for evaluation and demonstration purposes. It is not supported in a production environment.

You can read about this at <https://access.redhat.com/kb/docs/DOC-41794>

The product now comes with H2, an in-memory reference database.

Q: What components are included in this product?

A: A link to the list of components can be found here:

<https://access.redhat.com/support/configurations/jboss/>

Q: What is a Technology Preview?

A: Technology Preview features are not supported, may not be functionally complete, and are not intended for production use. These features are included to provide customers with early access to upcoming product innovations, enabling them to test functionality and provide feedback during the development process.

Red Hat support will provide commercially reasonable efforts to resolve any reported issues that customers experience when using these features.

Q: Where can I find out more details about my support entitlements?

A: Details of support policies can be found at the following URLs:

http://access.redhat.com/support/policy/support_process.html

<https://access.redhat.com/support/offerings/production/>

<https://access.redhat.com/support/offerings/developer/>

https://access.redhat.com/support/policy/updates/jboss_notes/

http://www.redhat.com/licenses/jboss_eula.html

More information about support can be found at <https://access.redhat.com/support/>

Q: Where can I obtain the source code?

A: The source code for this and earlier releases can be downloaded from the Red Hat Customer Portal at

<https://access.redhat.com/jbossnetwork/>

Q: What is Red Hat JBoss Data Virtualization?

A: Red Hat JBoss Data Virtualization is a data integration solution. It includes the following features:

Open source data virtualization engine based on the Teiid community project.

JBoss Data Virtualization tooling is integrated with Red Hat JBoss Developer Studio.

Metadata repository based on the Modeshape community project.

Management through the JBoss EAP Management Console. Management and monitoring through the Red Hat JBoss Operations Network.

Updated data source support.

You can find out more about JBoss Data Virtualization at
<http://au.redhat.com/products/jbossenterprisemiddleware/>
