



REDPOINT
— G L O B A L

RedPoint Data Management Release Notes
Version 8.4.4

RedPoint Global | 36 Washington Street

Wellesley Hills, MA 02481

+1 781 725 0250 | www.redpointglobal.com

Contents

Prerequisites and requirements.....	1
End of support notice.....	2
Features.....	2
Improvements.....	2
Known issues.....	3

RedPoint Data Management

Version 8.4.4 Release Notes

Prerequisites and requirements

This release of RedPoint Data Management offers several features and improvements. Before installing it, you should be aware of some prerequisites and requirements:

- You must have the Java Runtime Environment (JRE) Version 8 (64-bit) installed on your server machines to run RedPoint Data Management Version 8. If you have an older version of the Java JRE installed, visit <http://www.oracle.com/technetwork/java/javase/downloads/jre8-downloads-2133155.html> to upgrade.
- If you are installing the Data Management client on a different machine than the Data Management services, you may need to open the firewall ports associated with those processes so that the client can access the services. Data Management Version 8 uses different port numbers than previous versions, so that you can continue to run multiple versions side-by-side.

Process	Default ports
Site Server	20410
Execution Server	20411-20416
Web Proxy Server	20417

If any ports in the range 20410-20417 are already in use, you can change these defaults during installation:

- ➔ Click **Show advanced site server settings**, and specify an available **Base Port**. (To review a list of active connections, press the Windows logo key + R and then type `cmd`. Type `netstat -a -n` at the command line.)
- If during a previous Data Management installation you specified a base port other than the default value of 20410, and you select the option “Upgrade or repair the existing installation” during installation, the installation will fail with an error. See “Workaround” on page 4.
- If you are upgrading from Data Management Version 7.x, you will be given the option to migrate the repository to Version 8. *Note that previously-defined connections to Hadoop clusters cannot be migrated, and must be defined again.*
- If you are upgrading from Data Management Version 7.x and choose to migrate the repository to Version 8, both versions will use the same folder for temporary disk space. This can cause problems if you run both versions at the same time. To prevent this, change the temp space path to a different folder:

1. In the Data Management client, browse to **Repository > Settings > Site**, and click the **Resources** tab.
2. Edit each entry under **Temp Spaces** to ensure that Version 8 of Data Management is using a different **Path** than Version 7.
3. Click **Commit**.
4. Restart the Data Management services.

End of support notice

RedPoint Data Management ended support of MySQL 5.1 and 5.5 with version 8.3, and now supports MySQL 8.0.

Features

This release of RedPoint Data Management offers the following enhancements:

- RDBMS connections to AWS Aurora MySQL databases now support IAM authentication.
- The RDBMS Output tool now includes support for JDBC SQL Server Bulk Copy inserts.

Improvements

This release of RedPoint Data Management offers the following fixes:

- **Rename** and **Disable** options have been returned to tool and macro Properties Execution tab.
- Automation FTP steps configured for SSH/SFTP can download files larger than 2GB.
- Disabling and enabling tools in dynamic processing contexts works correctly.
- The **Refresh** command (**Edit** menu or **F5**) refreshes all data connections and schemas.
- The timeout limit for the RedPoint Data Management client's logon request has been increased to accommodate security environments that are slow to validate users.
- Rare issue no longer causes projects and automations to become unresponsive, leading to disconnection and project abort by `rpdm_cmd.exe`.
- The configuration locator syntax can be used to locate and modify the disabled state of project tools and automation steps.
- **General replace** command in Advanced Repository Actions works correctly for automations.
- The BigQuery Input tool correctly handles DATETIME data types containing values that include subseconds.
- RDBMS Azure Data Warehouse Bulk load provider correctly handles DATETIME data types containing NULL values.

Known issues

Java classpath settings are deprecated

In previous versions of Data Management, you could influence the classpath of Data Management's embedded Java Virtual Machine (JVM) in one of two ways:

- Enable the **Use CLASSPATH environment variable** option.
- Add folders and JARs to the class path in `/Settings/Site` or `/Settings/Machines/machine`.

These settings were most commonly used to reference third-party JDBC drivers and custom tools created with the Data Management SDK. However, class-path collisions between the referenced JARs and other Java-based applications could occur, causing errors that could be difficult to resolve.

As of version 8.2.1, Data Management employs a modular class-loader strategy to avoid such errors. However, this requires that you manage your Java dependencies more carefully. Therefore, the two techniques described above are **deprecated**. Any previously-defined configuration settings using this option will be removed during upgrade installations.

Installing JDBC drivers

If you previously accessed third-party JDBC drivers using one of the deprecated options, you should install those drivers in the Data Management installation folder instead.


To install JDBC drivers

1. Copy the JDBC driver JARs to the computers that host your Data Management Execution Servers, placing them in the `\jdbc_drivers` folder in the Data Management installation folder.
2. Close and re-open any projects or automations that use JDBC drivers, and restart any web services.

Installing Data Management SDK tools

If you previously accessed custom SDK tools using one of the deprecated options, you should install the custom JARs in the Data Management installation folder instead.

To install custom Data Management SDK tools

1. Close the Data Management clients.
2. Stop the Data Management Site and Execution services.
3. Copy the SDK tool JARs to the computer that hosts your Data Management Site Server, placing them in the `\java_plugins` folder in the Data Management installation folder.
4. Restart the Data Management Site and Execution services.
5. When you start the Data Management client, the SDK tools should appear on the Palette. If not, click the Palette menu , and then click **Reset Palette**.

Data Management will find the new tools and assign each SDK tool JAR its own ClassLoader to prevent classpath dependency collisions. If your tool requires additional libraries, you must include those dependencies in your SDK tool JAR.

Data Management SDK tool development

If you are developing your own SDK tools that depend on third-party libraries, we no longer recommend that you add those third-party library JARs to the class path. Instead, you must include those dependencies in your SDK tool JAR using the Apache Maven Shade Plugin. Refer to the Data Management SDK Developer Guide for details and examples.

Previously-defined shared tool settings may be removed

If you have previously defined shared tool settings in the repository for Data Management's MQ, Salesforce, or Kafka tools, these settings may be removed during upgrade, and you must re-enter them after installing version 8.3.1 and above.

To define shared tool settings

1. Open the Tools folder under Settings in the repository.
2. Click the tab for the desired tool name, and then view the Properties pane.
3. Configure the default tool properties for your environment.

Upgrade fails if existing Data Management installation specifies non-default base port

If during a previous Data Management installation you specified a base port other than the default value of 20410, and you select the option "Upgrade or repair the existing installation" during installation, the installation will fail with an error.

Workaround

During an upgrade installation:

1. Select the option **Install and configure new components**, and then click **Next**.
2. Do not specify a different **RedPoint Data Management program files directory**. (If you need to define a new installation location, you must uninstall RedPoint Data Management and perform a clean installation.) Optionally, you can select **Use an alternate data location** to specify a location for large or writable data. Click **Next**.
3. Click **Show advanced site server settings** and specify an available **Base port**. (To review a list of active connections, press the Windows logo key + R and then type `cmd`. Type `netstat -a -n` at the command line.)
4. Complete the installation in the usual way.

When Advanced Security is enabled, logons with no domain get an error

If advanced security mode is enabled, users with logons that do not include a domain may be unable to open or create projects and automations, and may see an error message like `errno6`.

Workaround

Either specify a **Default domain** in Site Settings, or specify a fully qualified **OS User** for each Data Management user.

Refresh command does not update changes to shared SDK tool settings

The **Refresh** command does not work for Data Management SDK-based tools that use shared settings. These include:

- Azure Service Bus tools
- BigQuery tools
- Cosmos DB tools
- Document tools
- Kafka tools
- MongoDB tools
- MQ tools
- Salesforce tools
- Snowflake tools
- SQS tools

Workaround

To refresh shared settings in a single tool:

1. Select the tool and click **Commit** on the Properties pane.
2. Click away from the tool, and then select it again.

To refresh all shared settings in a project that uses these tools, close and reopen the project.

Hadoop secure mode and Linux Integrated Security cannot be used together

When you configure a Data Management Hadoop connection to use Kerberos, Data Management creates a Kerberos configuration file (`krb5.conf`) in a non-default location, and specifies the Java System Property `java.security.krb5.conf` to point to that location. On Linux, Kerberos is used with Integrated Active Directory Authentication (or *Integrated Security*) via a configuration file created by the System Administrator, and also uses the `java.security.krb5.conf` Java System Property. This conflict can cause data connection failures in Data Management.

Workaround

None.

Indexes created with Cosmos DB Execute tool assigned default names

Due to a bug in Cosmos DB, index names you define using the Cosmos DB Execute tool's CREATE command may be assigned default names.

Workaround

None.

Using built-in SQL functions with RDBMS tool configured for Oracle JDBC causes error

Due to an issue with Oracle's JDBC driver, SQL statements that include Prepared Statement parameter replacement may fail with an error, because the JDBC driver is unable to replace parameters used as arguments to a SQL function in a Prepared Statement. For example, this statement:

```
SELECT * FROM oracle_tablename WHERE DUMMYDATE = TO_DATE(?, 'yyyy-mm-dd') ORDER BY ID, DUMMYDATE
```

will get the following error:

```
Error creating field specification list: Unsupported feature: checkValidIndex
```

Workaround

Do not use Prepared Statement parameter replacement with Oracle/JDBC.

Change to ML Trainer tool classification model output field names

In previous Data Management versions, Machine Learning Trainer tools configured to produce classification models included output field names with in-line parentheses in the models. This could cause errors in downstream Calculate tools. Projects that reference classification models created with previous versions of Data Management may show configuration errors when opened in the current version.

Workaround

Do not create new projects that reference classification models created with previous versions of Data Management. Re-run the ML Trainer project, and reference the new model in your project.

Configuring a Hadoop connection in edge node mode locks Execution Server ports

When the Data Management Hadoop configuration tool is run in "edge node" mode, it launches an external Java process. By default, child processes on Linux inherit all file handles from their parent process. If the parent process (in this case, the Execution Server) exits while the child process (the configuration tool) continues to run, the child process will take over its parent's listen port file handles. This will prevent the Execution server from starting.

This can happen when the Hadoop configuration tool attempts a connection to non-existent Resource Managers or Name Nodes. In this case, the configuration tool may wait for the cluster indefinitely. While the Execution server will report that configuration has failed after 5 minutes, the configuration tool continues running in the background—creating the scenario described above.

Workaround

Identify the process that is locking the Execution Server ports (typically 20411-20416) and kill it, then restart the Execution Server.

No support for Hive and HBase on ESP-enabled HDI 3.6

HDInsight 3.6 clusters configured with Enterprise Security Package (ESP) secure access to HBase and Hive with Apache Ranger, which Data Management does not yet support.

Workaround

None.

Azure tab missing for Azure Data Warehouse JDBC Bulk data connections

If you define a data connection and select the **Azure Data Warehouse JDBC Bulk** RDBMS provider, the Properties **Options** tab does not list **Azure JDBC loader** as an **Insert type**, and the **Azure** tab containing additional bulk-loading configuration options is not displayed.

Workaround

Configure a RDBMS Output tool to use the **Azure Data Warehouse JDBC bulk** provider, overriding the **Account** and **Secret key** properties on the tool's **Azure** tab.