



# Redpoint Data Management

Version 9.0.4 Release Notes

July 30, 2020

---

Redpoint Global Inc., 36 Washington Street, Suite 120 Wellesley Hills, MA 02481 USA  
T: 1 781-725-0250 F: 1 781-583-0464 [www.redpointglobal.com](http://www.redpointglobal.com)

# Contents

- Prerequisites and requirements.....1
- Redpoint Interaction and Redpoint Data Management.....1
- End of support notices .....2
- End of support advance notice.....2
- Features.....2
  - Security .....3
  - Usability .....3
  - Manageability .....4
  - Platforms and data environments .....4
  - Miscellaneous .....4
- Improvements .....4
- Known issues .....7

# Redpoint Data Management

## Version 9.0.4 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:

- If you are upgrading from Data Management Version 8 you will be given the option to migrate the repository to Version 9. *Note that previously-defined connections to Hadoop clusters cannot be migrated, and must be defined again.*
- If you are upgrading from Data Management Version 8 and choose to migrate the repository to Version 9, 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 avoid this, change the Version 9 temp space path to a different folder:
  1. In the Version 9 Data Management client, browse to **Repository** > **Settings** > **Site**, and click the **Resources** tab.
  2. Edit each entry under **Temp Spaces** to ensure that Version 9 of Data Management is using a different **Path** than the previous version.
  3. Click **Commit**.
  4. Restart the Data Management services.

### Redpoint Interaction and Redpoint Data Management

With Version 9.0 of Redpoint Data Management, Redpoint Interaction (RPI) functionality that relies on Data Management (for example, Data Intake, Data Process Projects, and SFMC Channel Synch) is implemented using the Data Management Operational API (OAPI). If you use RPI, you must ensure that the OAPI is deployed and is accessible to both the Data Management and the RPI environments, and that RPI is configured to use the OAPI.

The following system configuration parameters must be defined:

System parameter	Value
DataManagementIntegrationType	Set this to 1
DataManagementOAPIWebServiceAddress	Set to the OAPI Web Service URL

For more information about configuring the integration between RPI and Redpoint Data Management, please refer to the section “RPI to RPDM Integration” in the Redpoint Interaction System Administration Guide (<https://cdn.redpointglobal.com/redpoint-interaction/RPI%20v5.2%20Admin%20Guide.pdf>) .

For more information about the Redpoint Data Management Operational API, please refer to <https://support.redpointglobal.com/hc/en-us/sections/115003520608-Redpoint-Data-Management-Operational-API>.

## End of support notices

- Redpoint Data Management Version 7.4 is no longer supported.
- Redpoint Data Management ended support of MapR 5.2 with version 8.4, and now supports MapR 6.1.
- Redpoint Data Management ended support of CentOS 6 with version 8.4, and now supports CentOS 7.
- ActiveReports has been completely removed and existing reports can no longer be run with Data Management.
- Data Management RDBMS tools now implement OLE DB connections via the MSOLEDBSQL driver introduced by Microsoft in April 2018. This driver replaces the SQLNCLI, SQLOLEDB, and SNAC drivers, which will not work with our OLE DB data provider.
- Redpoint Data Management ends support of Netezza 7.1.x with version 9.0, and now supports Netezza 7.2.
- Redpoint Data Management now includes native RDBMS support for Snowflake, and support for the standalone Snowflake Input and Output tools will end in Version 9.11. You should migrate projects to use the new capabilities.

## End of support advance notice

- Support for Windows Server 12 will end with Redpoint Data Management Version 9.1.
- Support for the ODBC Teradata RDBMS connector will end with Redpoint Data Management Version 9.1. You should migrate Teradata connections to the JDBC provider.
- Support for Cloudera's CDH 5.x distribution will end with Redpoint Data Management Version 9.1.

## Features

Version 9.0.4 of Redpoint Data Management includes the following enhancement:

- Improvements to input file analysis.

Version 9.0.3 of Redpoint Data Management included the following enhancement:

- JDBC support for Teradata 16.20.

Version 9.0.2 of Redpoint Data Management included the following enhancements:

- Support for Microsoft's newly-released Azure Data Lake Gen2.
- Support for FIFO queue types in Amazon Simple Queue Service (SQS) tools.
- Support for Greenplum Database 4.3.

Version 9.0.1 of Redpoint Data Management included the following enhancements:

## Security

Data Management Version 9.0 introduced new password storage and encryption technologies with enhanced encryption algorithms, keys, and access control. These changes support regulatory compliance and strong encryption for security-conscious customers. Some highlights:

- FIPS 140-2 compliant security modules.
- Best-practice encryption algorithms: 2048-bit RSA keys, 256-bit AES ciphers, and password-protected key stores in PKCS#12 format.
- Key access provided on an as-needed basis.
- Digital authentication credentials encrypted using a hybrid RSA/AES-GCM algorithm and encoded into self-describing ASN.1 envelopes.
- Access to private keys for decryption and signing controlled via password-protected key stores.
- Access to public keys for encryption and signature validation available to all users and clients.
- Signed timestamp strategy with dedicated RSA key pairs for inter-service authentication ensure that RPDM services cannot be spoofed by bad actors on the network.
- JMS tools support SSL connections, and Data Management uses the latest version of OpenSSL (version 1.0.2q-1).
- Support for SSL certificate chains.
- Enhanced audit logging for session start and log on.
- Improved logging information for user auditing of live automations.
- Ability to write automation logs to the file system.
- Signed publisher for the Windows installer.

## Usability

- A complex project can contain many tools, making the project challenging to view. You can use **tool containers** to organize collections of tools on the canvas. These tool containers can be minimized to reduce visual complexity on the canvas, displaying a text label if one has been defined.
- The **Find in Project** window is now dockable and can be toggled from a toolbar icon, the **View** menu, or a hotkey. It displays a list of all the tools in the current project. For each tool, you can view tool name, status, percentage of completion, memory and temp disk space usage, and any notes defined on the tool's **Execution** tab. You can navigate the project canvas by selecting tools in the grid. Use the **Search** box at the top of the grid to search for text in the tool name, tool notes, or both. You can filter the tools displayed in the grid by status. And you can double-click a tool name or note to edit its value directly, rather than defining it on the **Execution** tab.
- All control types in project parameters now support multi-line tooltips.
- You can now define tool **Name** and **Notes** on the Execution tab of an Automation step's Properties.
- You can now search the database **Providers** list on RDBMS tool Properties.

## Manageability

- **Repository search** lets you search the repository by object name or other text, and refine that search by various criteria, including limiting the search to selected object types and finding objects last modified within a defined time window or by a specific user.
- Data Management now supports **bi-directional synchronization** between the repository and disk, available from both the command-line and the client. All repository objects can be synchronized to folders or ZIP files. This makes it easier to promote projects between environments—from development to test and then on to production. It also enables a more robust configuration/release project management process—especially when used with a version control system like Git.

## Platforms and data environments

- Support for Windows Server 2019
- Support for MapR 6.1, CentOS 7, and Amazon Linux v2.
- Support for Microsoft's newly-released OleDb.
- RDBMS JDBC connections to Netezza 7.2.1 and AWS.
- Support for both site-wide and user-level credential/encryption sets for access to AWS S3 buckets.
- MySQL databases now support IAM authentication.
- Support for Hive and HBase on ESP-enabled HDI 3.6.
- The RDBMS Output tool now supports a field delimiter configuration option for the Redshift and Azure bulk loader types.
- Native RDBMS support for Snowflake, including bulk loader provider. *The standalone Snowflake tools have been deprecated.*
- Support for Snowflake authentication with Okta.

## Miscellaneous

- An embedded Java Runtime Environment (JRE) is now included in the Data Management installation, ensuring Java version compatibility.
- The JMS tools introduced in version 8.4.1 can now connect to Apache ActiveMQ or IBM MQ providers. The MQ tools have been deprecated.
- Improved analysis of textual input file formats (JSON, XML, CSV, Flat). Some of the improvements include defaulting minimum text size to 32 instead of 8; defaulting to Unicode; int and float data types default to a size of 8; JSON and XML analyzers default to producing only Unicode, rather than Textvar.
- Throttling and queuing for auto-triggered diagnostics capture has been added, limiting the consumption of Execution Server resources.
- **Inactivity timeout** can now be configured to a maximum setting of 10,080 minutes (one week).

## Improvements

Version 9.04 of Redpoint Data Management includes the following fixes:

- Support for accessing Amazon S3 buckets using EC2 instance profiles with IAM roles.

- Snowflake bulk loader can load data containing embedded newlines or carriage returns.
- Capture Diagnostics command includes execution server trace files in diagnostics archive file.
- Retry logic is now more robust when using Azure SQL JDBC.
- Assorted stability improvements.

Version 9.03 of Redpoint Data Management included the following fixes:

- Out-of-range integer values inserted into SQL Server and related databases are now handled correctly.
- Parallel sections processing data beyond a certain combination of size, complexity, and number of tasks no longer abort with "queue limit reached" error.
- Multiple instances of the Snowflake and Azure Data Warehouse bulk loaders can run concurrently without risk of contention for temporary files.
- The Snowflake JDBC bulk loader now correctly handles embedded quote characters.
- The Snowflake JDBC bulk loader correctly handles column field values containing the single character "N".
- Improved support for DB2 via JDBC.
- The Salesforce Output tool now correctly translates Boolean values to Checkbox fields in bulk mode.
- DLD files can now be read from Azure Blob Storage using wildcard-pattern file names.
- Changes to Automation/Project settings are updated on Automation/Project save.
- Password control on Automation User Interface step supports variable assignment.
- RDBMS Output tools now correctly perform variable replacement for scripting actions.
- Assorted fixes to Tool Containers.

Version 9.02 of Redpoint Data Management included the following fixes:

- Records containing multiple very large Unicode fields are fully supported.
- Save to Archive action now includes child items from unopened repository folders.
- RDBMS connections can access Microsoft Excel and Microsoft Access via the ACE OLEDB 12.0 driver.
- RDBMS output connections that use the MSSQL OLEDB provider now handle memory allocation correctly.
- Viewing and interacting with AWS S3 buckets containing folders with more 1000 or more objects is fully supported.
- Extract Layout action supports Unicode characters in field names.
- Passwords entered into tool, macros, data connections, and automation steps are always masked.
- Automations recognize Site Settings variables referenced using the `${settings.variablename}` construction.
- Name Parse macro handles names containing apostrophes correctly.
- RDBMS Output tools configured to use Snowflake temp tables can perform Update operations.

- Azure Data Warehouse bulk loader now reports accurate inserted record counts when load fails.
- Reloaded macros retain their initial canvas position.
- The message “The option to send errors to Message Viewer will override the unselected Report warnings option” is no longer displayed for DLD Input tools.

Version 9.01 of Redpoint Data Management included the following fixes:

- Data Management services no longer occasionally fail to start on Linux system reboot.
- Stopping the redpointdm8\_\* services on RHEL-based systems does not stop all rpdm\_\*server processes on the system, regardless of Data Management version and path.
- When the RDBMS Output tool is configured with the JDBC bulk provider for AWS Redshift and the **Clear table before load** option enabled, the table is not truncated after load.
- Tools with multiple output nodes do not obscure record count displays.
- Projects do not fail with "no handler wants this request" error when running a mix of in-Hadoop and in-server processing under load.
- Offline license activation works correctly.
- Snowflake data connectors handle embedded commas correctly.
- The command-line argument `-dynconfig` can be used to disable/enable controls inside a macro.
- Stored credentials on Linux are not corrupted due to serial number “drift” when virtual machines are rebooted.
- Web Service Call tool no longer limits the send/receive size to 2GB.
- Project parameters warn when leading/trailing spaces are found in field names.
- Wildcarded server names are supported in published web services.
- RDBMS Output tool configured for Redshift does not create table columns that are larger than those in the source database table.
- Int and Bigint data types read from Microsoft Azure SQL Data Warehouse are correctly resolved, eliminating mismatches between source and target databases.
- Expression editor does not cause Project Server to fail when a deprecated expression is entered.
- Expression editor does not delete Calculate tool when characters are selected and **Delete** button is pressed.
- The Data Management Client, rpdm\_shell, and rpdm\_cmd never silently fall back to plain connections.
- SDK tools do not show passwords in Advanced Edit mode when **Override Site Settings** is configured.
- The `web_service_proxy_address` parameter is set only in the CoreCfg.properties file, rather than ExecutionServerCfg.properties.
- Long-running Data Management sessions do not retain and attempt to re-use Kerberos tickets past their expiry.
- Attaching to a project in read-only mode displays the correct warning message.



- Refresh command also refreshes database connections with cached schemas.

## Known issues

### *Mismatch between Site Server and Execution Server security settings causes failures*

It is possible to create a multi-host Data Management site with incompatible security settings. An Execution Server configured with different core security settings than its Site Server will not run correctly.

#### Workaround

Check that the Data Management configuration file **CoreCfg.properties** (located in the Data Management installation folder) has the same core security settings as the Site Server for the following properties:

- `use_ssl`
- `validate_ssl_cert`
- `enable_advanced_security`
- `administrator_os_user`

If necessary, stop the Execution service, edit the **CoreCfg.properties** file using a text editor, and then restart the Execution service.

### *Cannot change repository view to specific time*

Setting the Site Server's Repository View property to a defined **Label** or **Specific time** has no effect.

#### Workaround

Select the desired view from the **Version** tab on the repository object's Properties pane.

### *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.

### *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.

### *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 Data Management versions prior to 8.1, 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 pre-8.1 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 pre-8.1 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 20421-20426) and kill it, then restart the Execution Server.

### ***RDBMS Output tools configured for block-insert transactions on Netezza can fail with error***

In RDBMS Output tools using transactions, Netezza can be sensitive to prior failures in the sequence of operations within the tool.

#### Workaround

Avoid adding pre-execution steps that are likely to fail, as these can cause subsequent operations running in transactions to fail.

### ***Database-qualified table names containing dot character cause failures under Hive 1.2***

Hive version 1.2.0 changed the way the dot character is interpreted in HQL. If you are running Hive 1.2 or later and have custom HQL scripts generated using previous versions of Redpoint Data Management, you may encounter the error “Table or database name may not contain dot(.) character.”

#### Workaround

Edit custom HQL scripts, changing instances of `'database.table'` to `'database'.'table'`.

### ***System variables referencing project directory or path now contain full URI***

In Data Management versions prior to 9.0, the following system variables contained only the “native path” part of a variable, regardless of its URI scheme:

- ProjectPath
- AutomationPath
- ProjectDirectory
- AutomationDirectory

For example, if an automation was located in the folder `hdfs://servername/projects`, the variable **system.AutomationDirectory** would be set to `/servername/projects`.

These variables now return full URIs for local file system paths (for example, `file:///F:/projects` rather than `F:/projects`). Projects or automations that rely on the previous behavior may show an error of the form: project `'/servername/projects/...' does not exist`.

#### Workaround

Projects or automations that rely on the previous behavior can be edited to use the **NativePathFromURI** function, for example: `NativePathFromURI(system.projectDirectory)`.

#### ***Save YARN logs link unavailable in project messages***

When a project is executed using the **Run on Hadoop** feature (distinct from a project with a Parallel Section), the **Click to save YARN log** message does not appear in the project's Message Viewer.

#### Workaround

The **Double-click here to save YARN log** message is available in the project log when the project completes. To access it:

1. Open the Management Dashboard.
2. Navigate to **Logs > Projects**
3. Select your project's log, and find a message like `DIAGNOSTIC INFO: Double-click here to save YARN logs for Hadoop job`.

#### ***Expression using dollar-brace string literal does not evaluate correctly***

Using the dollar-brace string literal `"${"` in an expression such as `"${" + UpperCase(Trim(Name)) + "}"` does not work as expected, because Data Management interprets anything containing this character pair as a configuration-replacement construct.

#### Workaround

Escape the characters separately as `"$ " + "{ "`.  
For example: `"$" + "{" + UpperCase(Trim(Name)) + "}"`.

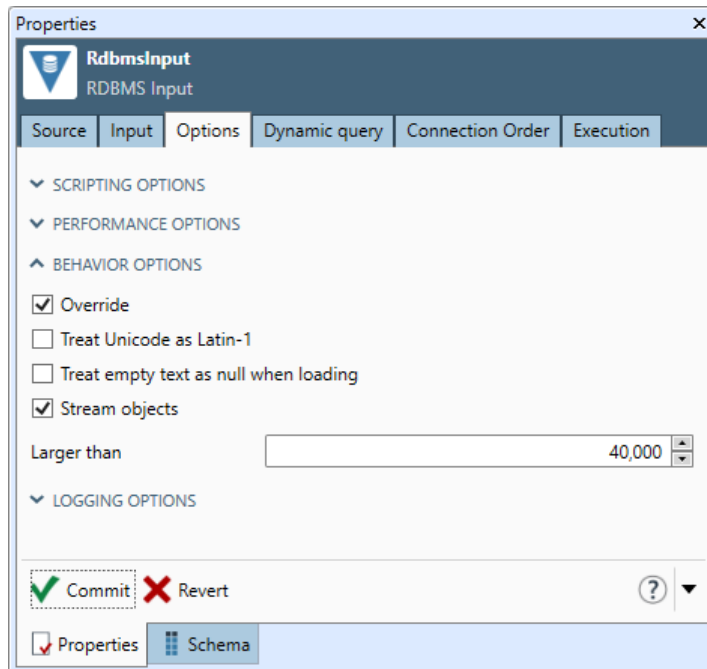
#### ***"Field exceeds maximum allowable buffer size" error in RDBMS Input***

RDBMS Input tools configured with ODBC connections and SQL queries that return large Unicode text fields may fail with "Field exceeds maximum allowable buffer size" error.

#### Workaround

Set the streaming threshold on the RDBMS Input tool or data connection to 40,000:

1. Select the **Options** tab of the RDBMS Input tool or data connection.
2. Under **Behavior Options**, select **Override**.
3. Select **Stream Objects**, and then enter **40000** in the **Larger than** box:



### *Microsoft SQL Server OleDB provider converts Time data types to TextVar*

RDBMS Input tools configured with MSSQL OleDB providers convert fields of type TIME to a TextVar data type.

#### Workaround

Replace the MSSQL OleDB provider with a SQL Server ODBC or JDBC provider, or use a Change Field Types tool downstream of the RDBMS Input to convert the field back to a Time data type.

### *Precision loss with Microsoft SQL Server OleDB provider and floating point number data types*

RDBMS tools configured with MSSQL OleDB providers may experience precision loss when handling fields of type Float.

#### Workaround

Replace the MSSQL OleDB provider with a SQL Server ODBC or JDBC provider.

### *Repository archive created with Version 9.0.2 client and Version 9.0.1 server cannot be loaded*

Repository archives created using a Version 9.0.2 Data Management client and a Version 9.0.1 Site Server are empty and cannot be used with the Load Archive command.

#### Workaround

Upgrade the Data Management Site Server to Version 9.0.2 or later.

