

Amazon RDS Migration Tool User Guide and Reference
Version 1.7

May 2015

Copyright © 2016, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

Your use of the Amazon RDS Migration Tool is subject to the terms of the existing AWS Customer Agreement or other agreement with AWS governing your use of AWS services, and the Additional Terms and Conditions for the Amazon RDS Migration Tool located at <https://s3-us-west-2.amazonaws.com/rdmt/RDSMigrationToolTermsConditions.pdf>. If you do not agree to the foregoing terms, you may not use the Amazon RDS Migration Tool.

Contents

1 Installing Amazon RDS Migration Tool

Installation Prerequisites	1-1
Software Requirements	1-1
Windows Software Requirements	1-1
Linux Software Requirements	1-2
Additional Required Software	1-2
Recommended Hardware Configuration	1-3
Supported Databases	1-4
Installing Amazon RDS Migration Tool	1-4
Installing Amazon RDS Migration Tool on Windows	1-4
Starting and Stopping the Amazon RDS Migration Server on Windows	1-4
Silently Installing Amazon RDS Migration Tool	1-5
Creating a Response File	1-5
Running the Silent Install	1-6
Silently Upgrading Amazon RDS Migration Tool	1-6
Creating a Response File	1-6
Running a Silent Upgrade	1-6
Silently Uninstalling Amazon RDS Migration Tool	1-8
Creating a Response File	1-8
Running a Silent Uninstall	1-8
Changing the Data Directory Location on Windows	1-8
Installing Amazon RDS Migration Tool on Linux	1-10
Pre-Installation Tasks	1-10
Installation Procedures	1-10
Amazon RDS Migration Server Procedures	1-12
Uninstalling Amazon RDS Migration Tool	1-12
Changing the Data Directory Location on Linux	1-13
Accessing the Amazon RDS Migration Console	1-13
Multiple Users Connecting to a Single Console	1-14

2 Getting Started - An Amazon RDS Migration Tool Tutorial

What You Need	2-1
Open the Amazon RDS Migration Console	2-2
Add an Oracle Database as a Source	2-2
Add a Microsoft SQL Server Database as a Target	2-4

Add a Replication Task	2-5
Add a Replication Task to the Amazon RDS Migration Console	2-6
Add the Source and Target Databases to the Task	2-6
Select Tables for the Replication Task.....	2-7
Run and Monitor the Replication Task	2-9
View the Replicated Tables in Microsoft SQL Server	2-10

3 Concepts and Design

Introducing Amazon RDS Migration Tool	3-1
Amazon RDS Migration Tool Features	3-1
What is Replication	3-2
System Architecture	3-3
What is a Replication Task	3-3
Full Load and CDC Processes	3-4
Replication Topologies	3-4
One-to-One	3-5
Logical Independence	3-5
Hub and Spoke.....	3-5

4 Introducing Amazon RDS Migration Tool Endpoints

Databases Supported as Amazon RDS Migration Tool Endpoints	4-1
Amazon RDS Migration Tool Data Types	4-2
Support for Large Object Data Types (LOBs).....	4-3
Amazon RDS Migration Tool Supported DDL Statements	4-3

5 Using an Oracle Database as a Source or Target

Supported Oracle Database Editions	5-1
Limitations	5-2
Using an Oracle Database as a Source	5-4
Supported Compression Methods	5-4
Configuration Properties for the Oracle Source.....	5-4
Security Requirements	5-5
Oracle Source Data Types.....	5-8
Non-Supported Data Types	5-10
Configuring an Oracle Database as an Amazon RDS Migration Tool Source.....	5-10
Provide Oracle Account Access	5-10
Ensure that ARCHIVELOG Mode is On	5-10
Setting up Supplemental Logging	5-11
Working with Oracle on Amazon RDS.....	5-12
Setting up an Oracle Database as a Source in Amazon RDS Migration Tool	5-14
Using Advanced Properties for an Oracle Source.....	5-16
Using an Oracle Database as a Target	5-20
Configuration Properties for the Oracle Target	5-20
Security Requirements	5-21
Oracle Target Data Types	5-23
Configuring an Oracle Database as an Amazon RDS Migration Tool Target	5-24

Provide Oracle Account Access	5-24
Setting up an Oracle Database as a Target in Amazon RDS Migration Tool.....	5-25
Using Advanced Properties for an Oracle Target	5-27

6 Using a Microsoft SQL Server Database as a Source or Target

Prerequisites	6-1
Limitations	6-2
Using a Microsoft SQL Server Database as a Source	6-4
Supported Compression Methods	6-5
Configuration Properties for a Microsoft SQL Server Source	6-5
Security Considerations for a Microsoft SQL Server Source	6-6
Working with Microsoft SQL Server AlwaysOn Availability Groups	6-6
Microsoft SQL Server Source Data Types	6-7
Non-Supported Data Types	6-9
Configuring a Microsoft SQL Server Database as an Amazon RDS Migration Tool Source	6-9
Preparing Microsoft SQL Server Backup and Recovery	6-10
Setting up Microsoft SQL Server for Replication	6-10
Replicating Tables that do not have a Primary Key.....	6-10
Defining Microsoft SQL Server Database Settings.....	6-11
Working with Windows Authentication	6-12
Setting up a Microsoft SQL Server Database as an Amazon RDS Migration Tool Source	6-12
Using Advanced Properties for a Microsoft SQL Server Source Database	6-14
Using a Microsoft SQL Server Database as a Target	6-17
Configuration Properties for a Microsoft SQL Server Target	6-17
Security Considerations for a Microsoft SQL Server Target	6-19
Microsoft SQL Server Target Data Types.....	6-19
Setting up a Microsoft SQL Server Database as an Amazon RDS Migration Tool Target.	6-21
Using Advanced Properties for a Microsoft SQL Server Target Database.....	6-23

7 Using a SAP Sybase ASE Database as a Source or Target

Prerequisites	7-1
Limitations	7-2
Using a SAP Sybase ASE Database as a Source	7-2
Configuration Properties for the SAP Sybase ASE Source	7-2
Security Requirements	7-3
SAP Sybase ASE Database Source Data Types.....	7-3
Non-Supported Data Types	7-4
Setting up a SAP Sybase ASE Database as a Source in Amazon RDS Migration Tool.....	7-4
Using Advanced Properties for a SAP Sybase ASE Source	7-6
Using a SAP Sybase ASE Database as a Target	7-8
Configuration Properties for the SAP Sybase ASE Target.....	7-8
Security Requirements	7-8
SAP Sybase ASE Database Target Data Types	7-8
Non-Supported Data Types	7-9
Setting up a SAP Sybase ASE Database as a Target in Amazon RDS Migration Tool	7-9
Using Advanced Properties for a SAP Sybase ASE Target.....	7-12

8 Using a MySQL Database as a Source or Target

Prerequisites	8-1
General Prerequisites	8-1
Amazon RDS Migration Tool Server for Windows	8-1
Amazon RDS Migration Tool Server for Linux.....	8-2
Enable Binary Logging.....	8-2
Cluster Prerequisites	8-3
Limitations	8-4
Using a MySQL Database as a Source	8-4
Security Requirements	8-4
Setting up Amazon RDS MySQL for CDC (Change Data Capture)	8-5
Configuration Properties for the MySQL Source.....	8-5
MySQL Database Source Data Types	8-6
Data Types Supported in Full Load Only	8-7
Setting up a MySQL Database as a Source in Amazon RDS Migration Tool	8-9
Selecting a Schema	8-10
Using Advanced Properties for a MySQL Source	8-11
Using a MySQL Database as a Target	8-13
Configuration Properties for the MySQL Target	8-13
Security Requirements	8-13
MySQL Database Target Data Types.....	8-13
Setting up a MySQL Database as a Target in Amazon RDS Migration Tool.....	8-15
Using Advanced Properties for a MySQL Target	8-18

9 Using ODBC to Connect to a Source

Prerequisites	9-1
Limitations	9-1
Using ODBC to Connect to a Source	9-1
Configuration Properties for the ODBC Source	9-2
ODBC Source Data Types.....	9-3
Configuring ODBC Databases to work as an Amazon RDS Migration Tool Source.....	9-6
Using Advanced Properties when Using ODBC Databases as a Source.....	9-8

10 Using Amazon Redshift as a Target

Introducing the Amazon Redshift Target Endpoint	10-1
Amazon Redshift Endpoint Prerequisites	10-1
Get Started with Amazon Redshift	10-2
Sign up for an Amazon S3 Bucket.....	10-2
Open the Required Firewall Ports.....	10-2
Amazon Redshift Data Types	10-3
Setting up Amazon Redshift as a Target in Amazon RDS Migration Tool	10-5
Using Advanced Properties for an Amazon Redshift Target	10-6

11 Using PostgreSQL Database as a Source or Target

Using PostgreSQL Database as a Source	11-1
Source Prerequisites	11-2

Client Side	11-2
Server Side.....	11-2
Security Requirements	11-3
Source Limitations	11-3
PostgreSQL Source Data Types	11-4
Setting up PostgreSQL database as a Source in Amazon RDS Migration Tool.....	11-6
Using Advanced Properties for a PostgreSQL Source.....	11-8
Removing RDS Migration Tool Artifacts from the Source Database.....	11-9
Using PostgreSQL as a Target	11-10
Target Prerequisites	11-10
Security Requirements	11-10
PostgreSQL Database Target Data Types	11-10
Data Types when Replicating from a PostgreSQL Source.....	11-11
Setting up PostgreSQL database as a Target in Amazon RDS Migration Tool.....	11-12
Using Advanced Properties for a PostgreSQL database Target.....	11-14
12 Using the Amazon RDS Migration Tool File Channel	
Setting Up Amazon RDS Migration Tool File Channel Tasks	12-1
Local Task	12-1
Remote Task	12-2
Replicating to Multiple Targets (Distribution).....	12-2
Adding Tables to a Running Remote Task.....	12-4
Working with the File Channel Data Files	12-4
File-Channel Directory Structure	12-4
Amazon RDS Migration Tool Installation Requirements for the File Channel	12-5
Security	12-6
Limitations	12-6
Using the File Channel as a Source.....	12-6
File Channel Source Properties.....	12-7
Setting up a File Channel Source using Amazon RDS Migration Tool.....	12-8
Using Advanced Properties for a File-Channel Source.....	12-10
Using the File Channel as a Target.....	12-11
File Channel Target Properties	12-11
Setting up a File Channel Target using Amazon RDS Migration Tool.....	12-13
Using Advanced Properties for a File-Channel Target	12-15
13 Using the Amazon RDS Migration Console	
The Tasks View	13-2
Task Icons.....	13-3
Viewing Specific Tasks.....	13-4
Design Mode.....	13-4
Monitor Mode	13-6
The Server View.....	13-8
Security Roles.....	13-9
Creating Active Directory Groups	13-9
Associating Multiple Groups with the Same Role	13-10

Security Roles and Console Elements.....	13-10
--	-------

14 Designing Tasks

Setting up Tasks	14-1
Working with Databases	14-2
Adding a Database	14-3
Editing Database Configuration Information.....	14-3
Viewing Database Configuration Information.....	14-3
Testing a Database Connection.....	14-4
Adding a Source and Target Database to a Task	14-4
Adding Tables to a Task	14-5
Search for Tables to use in a Replication Task.....	14-6
Selecting Specific Tables for Replication	14-7
Removing Specific Tables from a Replication Task.....	14-7
Create Table Selection Patterns.....	14-7
Editing a Replication Task	14-9
Deleting a Replication Task	14-9
Importing and Exporting Tasks	14-9
How to: Export and Import a Task	14-10
Editing an Exported (json) File	14-11
Making Changes to the Database Connection Information.....	14-11

15 Customizing Tasks

Table Settings	15-1
Carry out General Tasks for a Single Table	15-2
Define Transformations on a Single Table.....	15-3
Limitations	15-3
Using the Transform Tab	15-4
Creating an Expression for Transforms	15-6
Using SQLite Syntax with Transforms.....	15-6
Using Filters.....	15-7
Limitations	15-7
Opening the Filter Table Window	15-8
Creating a Filter Condition for a Specified Column	15-9
Creating a Record Selection Condition for One or More Columns	15-10
Using the Range Builder	15-11
Using SQLite Syntax with Filtering.....	15-12
Defining Global Transformations	15-13
Limitations.....	15-13
Starting the New Transformation Rule Wizard	15-14
Selecting the Transformation Type	15-14
Defining What to Transform.....	15-15
Defining the Transformation Rule	15-18
Limitations	15-18
Rename Schema.....	15-19
Rename Table.....	15-22
Rename Column.....	15-25

Add Column	15-27
Drop Column	15-28
Convert Data Type	15-28
Using the Global Transformation Rules List	15-28
Edit a Global Transformation Rule	15-29
Delete a Global Transformation Rule	15-29
Editing Global Transformation Rules	15-29
Using the Expression Builder (for Filters, Transforms, and Global Transformations)	15-30
Overview of the Expression Builder	15-30
Build an Expression	15-32
Evaluate an Expression	15-33
Test an Expression	15-33
Using Elements in the Expression Builder	15-34
Input Columns (Transforms and Filters only)	15-34
Metadata (Global Transformations only)	15-35
Operators	15-35
Functions	15-37
Header Columns	15-46
Task Settings	15-46
Metadata	15-47
Target Metadata	15-47
Control Tables	15-47
Full Load	15-48
Full Load Settings	15-48
Full Load Tuning	15-50
Change Processing	15-50
Apply Changes Settings	15-50
Changes Processing Tuning	15-51
Error Handling	15-52
Error Handling Settings	15-53
Environmental Errors	15-53
Data Error	15-53
Table Error	15-54
Apply Conflicts	15-55
Logging	15-56
Logging Level	15-56
Log File	15-56

16 Working with Tasks at Runtime

Running a Task	16-1
How to run a Task	16-1
Using the Run Button Options	16-2
Start Replication	16-2
Reload Target	16-3
Use Advanced Run Options	16-3
Viewing the Task Status	16-5
Reading Messages about a Task	16-5

Viewing Notifications	16-5
Using the Notifications List	16-5
View Log Messages for a Task.....	16-6
Using the Log Messages List	16-7
Viewing the Log file in the Log Viewer	16-7

17 Monitoring and Controlling Replication Tasks

Viewing Information in the Monitor	17-1
Monitoring the Full-Load Operation	17-1
General Information for a Full Load	17-2
Detailed Information for the Full Load	17-2
General Information for Completed Task	17-3
Information for Each Table in the Task.....	17-3
Information for Tables that have Completed Loading.....	17-4
Information for Tables that are Currently Loading	17-4
Information for Tables that are in the Loading Queue.....	17-5
Information for Tables with Errors.....	17-6
Monitoring Throughput in a Full Load Operation.....	17-6
Monitoring Change Processing Operations	17-7
General Change Processing Information	17-7
Detailed Change Processing Information	17-8
Information about Incoming Changes	17-9
Information about Applied Changes	17-10
Information about Change Processing Throughput.....	17-11
Information about Apply Latency	17-11
Viewing Messages	17-12
Using the Monitor Tools	17-12
Viewing History Information	17-13
Working with Task Logs at Runtime	17-13
Set the Task Logging Level.....	17-13
Carry Out Task Log File Actions	17-14
View the Task Log File	17-14

18 Amazon RDS Migration Server Settings

Notifications Settings	18-1
Defining Notifications.....	18-2
Creating a New Notification.....	18-2
Using the Notification List.....	18-11
Editing a Notification	18-12
Deleting a Notification	18-12
Setting up Mail Parameters	18-12
Creating a Default Recipient List	18-13
License Settings	18-14
Requesting a License	18-14
Using the Advanced License Request Option	18-15
Registering a License.....	18-15
Viewing a License	18-16

Error Handling	18-16
Logging Settings (Server)	18-17
Set the Server Logging Level.....	18-17
Configure Log Files	18-18
Schedule Maintenance Operations for Log Files.....	18-18
Carry out Log File Actions.....	18-19
View a Server Log	18-19
File Transfer Service	18-20
Defining a File Transfer Service.....	18-20
Editing a File Transfer Service	18-21
Deleting a File Transfer Service	18-21
Scheduling Jobs	18-23
A Error Handling	
Error Types	19-1
Error Handling Properties	19-1
Environmental Errors.....	19-2
Data Errors.....	19-2
Table Errors	19-4
Apply Errors.....	19-4
Fatal Errors.....	19-6
Abnormal Termination	19-7
B Using Encryption	
Setting up HTTPS with IIS	20-1
Setting up HTTPS without IIS	20-3
Creating and Installing a Self-Signed Server Authentication Certificate	20-3
Setting up the Self-Signed Certificate on the Server	20-3
Removing the Self-Signed Certificate from the Server	20-4
Setting Up the Self-Signed Certificate on a Client.....	20-4
Changing the Server Password	20-5
Protecting RDS Migration Tool Passwords	20-5
Changing and Protecting the Master Key	20-6
Master Key Replacement	20-6
Protecting the Mk.dat File from Misuse	20-6
Master Key Protection	20-7
Exporting and Importing Tasks	20-7
C Setting up Amazon RDS Migration Tool in a Cluster Environment	
Setting up Amazon RDS Migration Tool in a Windows Server Cluster (HA)	21-1
Preparing to Set Up the Cluster Environment.....	21-1
Create a Failover Cluster and a Service to Use in Your Environment.....	21-2
Add Storage.....	21-2
Define Client Access.....	21-2
Install Amazon RDS Migration Tool.....	21-3
Add the Amazon RDS Migration Tool Services.....	21-3

Define the Dependencies for Each Service.....	21-4
Enabling Different Console Configurations in a High Availability Environment.....	21-5
Setting up Amazon RDS Migration Tool in a Linux Cluster	21-5

D Control Tables

Apply Exceptions	22-1
Replication Status	22-1
Suspended Tables	22-2
Replication History	22-3

E Impact of DST Change on Amazon RDS Migration Tool

Glossary

1



Installing Amazon RDS Migration Tool

The following chapter describes how to prepare your system for Amazon RDS Migration Tool, how to install Amazon RDS Migration Tool and how to access the Amazon RDS Migration Console.

- [Installation Prerequisites](#)
- [Installing Amazon RDS Migration Tool](#)
- [Accessing the Amazon RDS Migration Console](#)

Installation Prerequisites

This section describes how to prepare your system to use Amazon RDS Migration Tool.

- [Software Requirements](#)
- [Additional Required Software](#)
- [Supported Databases](#)
- [Installing Amazon RDS Migration Tool](#)
- [Accessing the Amazon RDS Migration Console](#)

Software Requirements

This section describes what software is required to work with Amazon RDS Migration Tool.

- [Windows Software Requirements](#)
- [Linux Software Requirements](#)

Windows Software Requirements

To install the Amazon RDS Migration Server and Console on a Windows computer, you must have the following installed on your system:

- Any of the following operating systems:
 - Windows Server 2008 x64
 - Windows Server 2008 R2 x64
 - Windows 7 x64
 - Windows 8 RTM
 - Windows 8 SP1

- Windows Server 2012 RTM and SP1
- .NET framework 4 or 4.5

Linux Software Requirements

To install the Amazon RDS Migration Server on a Linux computer, you must have the following installed on your system:

- Red Hat Enterprise Linux 64-bit version 5.4 or higher.

Additional Required Software

To use Amazon RDS Migration Tool Console, make sure that the following is installed on the computer where you installed Amazon RDS Migration Tool:

- **.NET Framework 4 or 4.5** (you must use this and not the .NET Framework 4 or 4.5 client profile). If the computer with Amazon RDS Migration Tool does not have .NET Framework 4 or 4.5 installed, you can download and install it from the Microsoft site using this link:

<http://www.microsoft.com/en-us/download/details.aspx?id=17718>

- To use the Amazon RDS Migration Console, you must make sure that you use a supported browser. You can use one of the following browsers:
 - **Microsoft Internet Explorer** version 9 and above
 - **Mozilla Firefox** version 38 and above
 - **Google Chrome**
- **Adobe Flash Player** version 10.2.5 or later (for Microsoft Internet Explorer and Mozilla Firefox)

Recommended Hardware Configuration

This section describes the recommended hardware configurations for using Amazon RDS Migration Tool. For information on the software requirements for using Amazon RDS Migration Tool, see [Software Requirements](#) and [Additional Required Software](#).

The following table describes the recommended hardware configuration for installing Amazon RDS Migration Tool on Windows and Linux operating systems.

Table 1–1 Recommended Hardware Configuration

	Basic System	Large System	Extra-Large System	Notes:
Processor	Quad core	Quad core base Dual-core per task	8-core base Quad core per task	<p>Additional cores are useful in any of the following situations:</p> <ul style="list-style-type: none"> Many tasks running in parallel Full-load performance priority Multiple full-load processes running in parallel
Memory	8 GB	32 GB	64 GB	<p>More memory is useful in any of the following situations:</p> <ul style="list-style-type: none"> Many tasks running in parallel Long-running transactions on the source database (for example, monthly batch processing) Many active users on the source system
Disk requirements	320 GB 7200 RPM	500 GB 10,000 RPM RAID	500 GB 15,000 RPM RAID	<p>A faster disk is useful in any of the following situations:</p> <ul style="list-style-type: none"> Using a file-based target, such as Greenplum or Vectorwise Long-running source transactions that may not fit into memory Using tasks that are set up to continue processing during target outage <p>A larger disk is required in any of the following situations:</p> <ul style="list-style-type: none"> Using tasks that are set up to continue processing during target outage Very large source transactions that do not fit into memory <p>RAID is recommended for system recoverability in case of disk failure for all configurations.</p>
Network	1 Gb	10 Gb	Two 10 Gb	

Supported Databases

To replicate data using Amazon RDS Migration Tool, you must be sure to have a supported version of the database you are working with available. For information about the databases you can use with Amazon RDS Migration Tool, see [Introducing Amazon RDS Migration Tool Endpoints](#).

Installing Amazon RDS Migration Tool

You can install Amazon RDS Migration Tool on either Windows or Linux platforms. This section describes the following:

- [Installing Amazon RDS Migration Tool on Windows](#)
- [Installing Amazon RDS Migration Tool on Linux](#)

For information about recommended hardware requirements for working with Amazon RDS Migration Tool, see [Recommended Hardware Configuration](#).

Installing Amazon RDS Migration Tool on Windows

Install Amazon RDS Migration Tool using the AmazonRDSMigrationTool_1.1.x.xx_X64.exe installation kit. This kit runs on Windows 64-bit (x64) environments. For a list of the Windows versions supported by Amazon RDS Migration Tool, see [Windows Software Requirements](#).

Follow the instructions in the Install wizard to install Amazon RDS Migration Tool. Later if you need to start or stop the Amazon RDS Migration Server, see the following section:

[Starting and Stopping the Amazon RDS Migration Server on Windows](#)

Note: If you are running the Amazon RDS Migration Server on a Linux computer, in the Replication Server Location page of the installation window select **Connect to a remote Linux Amazon RDS Migration Server**.

When you complete the Windows installation, then follow the directions in [Starting and Stopping the Amazon RDS Migration Server on Windows](#).

All of the data that is created when you use Amazon RDS Migration Tool is stored in a folder called **data**. By default, this folder is located in the installation folder where you install Amazon RDS Migration Tool. If you want to create the data folder in a different location, select this option in the installation wizard.

If you select this option, all command line actions must be prefixed with `reptcl -d <path to the data folder>`

Starting and Stopping the Amazon RDS Migration Server on Windows

In some cases you may need to stop and start the Amazon RDS Migration Server. You must do this from the Windows computer where Amazon RDS Migration Tool is installed.

To stop and start the Amazon RDS Migration Tool Server on Windows

- From the **Start** menu on the Windows computer where Amazon RDS Migration Tool is installed, find Amazon RDS Migration Tool then select either **Stop Amazon RDS Migration Server** or **Start Amazon RDS Migration Server**.

Silently Installing Amazon RDS Migration Tool

Amazon RDS Migration Tool can be installed silently (i.e. without requiring user interaction). This option is useful, for example, if you need to install Amazon RDS Migration Tool on several machines throughout your organization.

The installation process consists of two stages:

- Creating a response file.
- Running the installation.

Creating a Response File

Before starting the installation, you need to create a response file. To create the response file:

- Copy the response file text below into a text editor.

Response file text:

```
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-DlgOrder]
Dlg0={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdWelcome-0
Count=7
Dlg1={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdLicense2Rtf-0
Dlg2={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdAskDestPath-0
Dlg3={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdAskDestPath-1
Dlg4={9C614355-28A0-4C2A-98DF-DB9FD674826F}-AskOptions-0
Dlg5={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdStartCopy-0
Dlg6={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdFinish-0
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdWelcome-0]
Result=1
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdLicense2Rtf-0]
Result=1
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdAskDestPath-0]
szDir=C:\Program Files\Amazon\RDS Migration Tool
Result=1
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdAskDestPath-1]
szDir=C:\Program Files\Amazon\RDS Migration Tool\data
Result=1
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-AskOptions-0]
Result=1
Sel-0=1
Sel-1=0
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdStartCopy-0]
Result=1
[[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdFinish-0]
Result=1
bOpt1=0
bOpt2=0
```

- To change the default installation directory, edit the *first* **szDir** value as necessary.
- To change the default data directory, edit the *second* **szDir** value as necessary.
- Save the file as **<name>.iss** e.g. **silent_inst_64.iss**

Running the Silent Install

To silently install Amazon RDS Migration Tool, open a command prompt and change the working directory to the directory containing the Amazon RDS Migration Tool setup file. Then issue the following command (where <response file> is the path to the response file you created earlier):

Syntax:

```
<RDS MIGRATION TOOL_KIT> /s /f1<RESPONSE_FILE> /f2<LOG_FILE>
```

Example:

```
C:\>AmazonRDSMigrationTool_1.1.x.xx_X64.exe /s /f1C:\temp\1\rdsmigrationtool_x64_ins.iss /f2C:\temp\1\silent_x64_ins.log
```

If the installation was successful, the log file should contain the following rows:

```
[ResponseResult]
ResultCode=0
```

Silently Upgrading Amazon RDS Migration Tool

Silently upgrading Amazon RDS Migration Tool consists of two stages:

1. Creating a response file.
2. Running the upgrade.

Creating a Response File

Before starting the upgrade, you need to create a response file. To create the response file:

1. Copy the response file text below into a text editor.

Response file text:

```
[[9C614355-28A0-4C2A-98DF-DB9FD674826F]-DlgOrder]
Dlg0={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdWelcome-0
Count=2
Dlg1={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdFinish-0
[[9C614355-28A0-4C2A-98DF-DB9FD674826F]-SdWelcome-0]
Result=1
[[9C614355-28A0-4C2A-98DF-DB9FD674826F]-SdFinish-0]
Result=1
bOpt1=0
bOpt2=0
```

2. Save the file as <name>.iss e.g. **rep_upg_64.iss**

Running a Silent Upgrade

To silently upgrade Amazon RDS Migration Tool, open a command prompt and change the working directory to the directory containing the Amazon RDS Migration Tool setup file. Then issue the following command (where <response file> is the path to the response file you created earlier):

Syntax:

```
<RDS MIGRATION TOOL_KIT> /s /f1<RESPONSE_FILE> /f2<LOG_FILE>
```

Example:

```
C:\>AmazonRDSMigrationTool_1.1.x.xx_X64.exe /s /f1C:\temp\1\rep_x64_up.iss
```

/f2C:\temp\1\silent_x64_up.log

If the upgrade was successful, the log file should contain the following rows:

```
[ResponseResult]  
ResultCode=0
```

Silently Uninstalling Amazon RDS Migration Tool

Silently uninstalling Amazon RDS Migration Tool consists of two stages:

1. Creating a response file.
2. Running the uninstall.

Creating a Response File

Before starting the uninstall, you need to create a response file. To create the response file:

1. Copy the response file text below into a text editor.

Response file text:

```
[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-DlgOrder]
Dlg0={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdWelcomeMaint-0
Count=3
Dlg1={9C614355-28A0-4C2A-98DF-DB9FD674826F}-MessageBox-0
Dlg2={9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdFinish-0
[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdWelcomeMaint-0]
Result=303
[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-MessageBox-0]
Result=6
[{9C614355-28A0-4C2A-98DF-DB9FD674826F}-SdFinish-0]
Result=1
bOpt1=0
bOpt2=0
```

2. Save the file as **<name>.iss** e.g. **silent_uninst_64.iss**

Running a Silent Uninstall

To silently uninstall Amazon RDS Migration Tool, open a command prompt and change the working directory to the directory containing the Amazon RDS Migration Tool setup file. Then issue the following command (where **<response file>** is the path to the response file you created earlier):

Syntax:

```
<RDS MIGRATION TOOL_KIT> /s /f1<RESPONSE_FILE> /f2<LOG_FILE>
```

Example:

```
C:\>AmazonRDSMigrationTool_1.1.x.xx_X64.exe /s /f1C:\temp\1\rep_x64_un.iss
/f2C:\temp\1\silent_x64_un.log
```

If the uninstall was successful, the log file should contain the following rows:

```
[ResponseResult]
ResultCode=0
```

Changing the Data Directory Location on Windows

This section explains how to change the location of the Amazon RDS Migration Tool Data Directory. Such a procedure may need to be performed if the drive on which the current folder resides has insufficient space or if you are moving from a temporary POC setup to production, for example.

To change the Data Directory Location:

1. Stop the **Amazon RDS Migration Console** and **Amazon RDS Migration Server** services.
2. Move the data folder to a new location. For example:
`C:\Program Files\Amazon\RDS Migration Tool\Data2`
3. Open the Registry and perform the following procedure:
 - a. Browse to:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\AmazonRDSMigrationToolConsole
 - b. Modify the **ImagePath** string as follows:
`"C:\Program Files\Amazon\RDS Migration Tool\bin\WebCmd.exe" service -d "C:\Program Files\Amazon\RDS Migration Tool\Data2" run`
 - c. Browse to:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\AmazonRDSMigrationToolServer
 - d. Open the **ImagePath** string and add **-d <path_for_new_data_directory>** after the `repctle.exe` path. For example:
`"C:\Program Files\Amazon\RDS Migration Tool\bin\repctl.exe" -d "C:\Program Files\Amazon\RDS Migration Tool\Data2" service start name=Server address=127.0.0.1 port=3553`
4. Start the Amazon RDS Migration Tool services.

Installing Amazon RDS Migration Tool on Linux

This section describes how to install the Amazon RDS Migration Server on Linux. For information on supported Linux platforms, see [Installation Prerequisites](#). This section contains the following:

- [Pre-Installation Tasks](#)
- [Installation Procedures](#)
- [Amazon RDS Migration Server Procedures](#)
- [Uninstalling Amazon RDS Migration Tool](#)
- [Changing the Data Directory Location on Linux](#)

Notes:

When using the Amazon RDS Migration Server on Linux, you must use the Amazon RDS Migration Console on a Windows computer. See [Windows Software Requirements](#) for a list of the supported Windows versions.

See [Linux Software Requirements](#) for a list of the Linux versions supported by Amazon RDS Migration Tool.

Pre-Installation Tasks

Before you begin the installation, you must carry out the following tasks:

- Install Amazon RDS Migration (Console) on any Windows computer. For information on installing Amazon RDS Migration Tool on a Windows computer, see [Installing Amazon RDS Migration Tool on Windows](#).
- Copy the RPM file to any location on the Linux computer.
- Amazon RDS Migration needs a dedicated account. The default account is created with `amazon` user and `amazon` group.

To use a different user/group, do the following:

- a. Create the user that you want to be the installation folder owner.
- b. Create a new file named `inst.dat` that contains the following lines:

```
AREPCTUSER=<YOUR_DESIRED_USERNAME>
AREPCTGROUP=<YOUR_DESIRED_GROUPNAME>
```

- c. Save the file to:

```
/opt/aws/inst.dat
```

Installation Procedures

Before you carry out the following procedures, make sure you carry out the [Pre-Installation Tasks](#) first.

You can install Amazon RDS Migration Tool either as ROOT or as a SUDOER user.

To install Amazon RDS Migration Tool as ROOT:

Issue the following command:

```
$ rpm -i [path to the RPM file]
```

To install Amazon RDS Migration Tool as SUDOER:

Issue the following command:

```
$ sudo rpm -i [path to the RPM file]
```

Setup will perform the following actions:

- Automatically create a new user and group named “aws” (unless you chose to use a different user and group or a user named “aws” already exists).
- Change the Amazon RDS Migration Tool installation folder owner to the “aws” user and group or to your preferred user and group. For instructions on changing the default user and group, see [Pre-Installation Tasks](#).
- Install the required files.

Linux Installation Examples

```
[repectl@lin1~]$ rpm -i ardsmigrationtool-<n>.rpm
```

-OR- (to specify a custom installation folder)

```
$ rpm -i --prefix=<FULL_PATH_TO_TARGET_FOLDER> <PATH_TO_RPM_KIT>
```

When installing on SUSE Linux, type the following command:

```
rpm -ivh --nodeps [path to the RPM file]
```

Example:

```
rpm -ivh --nodeps ards migration tool-<n>.x86_64.rpm
```

-OR- (to specify a custom installation folder)

```
rpm -ivh --nodeps --prefix=<FULL_PATH_TO_TARGET_FOLDER> <PATH_TO_RPM_KIT>
```

The output will be similar to this:

```
arepctl: Configured to point to the product directory /opt/amazon/rds migration
tool
Amazon RDS Migration server was started as PID 5071
```

Amazon RDS Migration Server Procedures

This section describes how to verify that the Amazon RDS Migration Server is running as well as how to start and stop the Amazon RDS Migration Server.

Verifying that the Amazon RDS Migration Server is Running

To verify that the Amazon RDS Migration Server is running, issue the following command:

```
ps -ef | grep repctl
```

A list similar to the following is displayed:

```
amazon    5071      1  0 17:32 pts/0    00:00:00
/opt/amazon/rdsmigrationtool/bin/...
amazon    5072  5071  0 17:32 ?          00:00:00
/opt/amazon/rdsmigrationtool/bin/...
root      5075  4736  0 17:32 pts/0    00:00:00 grep repctl
```

Starting the Amazon RDS Migration Server Process

To start the Amazon RDS Migration Tool Server service, run the following command (shown using the default installation path):

```
/opt/amazon/rdsmigrationtool/bin/arepctl start
```

You will receive a confirmation similar to:

```
Amazon RDS Migration server was started as PID 5100
```

Note: If you are running as `amazon` (or the user selected during installation) then you can simply run the `arepctl start` command. Otherwise, you need to be a root user or a `sudoer`.

For an explanation of how to verify that the Amazon RDS Migration Server service is running, see [Verifying that the Amazon RDS Migration Server is Running](#).

Stopping the Amazon RDS Migration Server Processes

To stop the Amazon RDS Migration Server service, run the following command (shown using the default installation path):

```
/opt/amazon/rdsmigrationtool/bin/arepctl stop
```

You will receive a confirmation similar to:

```
Amazon RDS Migration server was sent a stop signal
Waiting for Amazon RDS Migration server to stop (2 seconds)
Waiting for Amazon RDS Migration server to stop (17 seconds)
Amazon RDS Migration server is no longer running
[service command] Succeeded
```

Uninstalling Amazon RDS Migration Tool

To uninstall Amazon RDS Migration Tool, type the following at the Linux prompt:

```
rpm -e rdsmigrationtool
```

You will receive a confirmation similar to:

```
Amazon RDS Migration server was sent a stop signal
```



```
[service command] Succeeded
```

To ensure that Amazon RDS Migration Tool was removed from the computer, type the following to list the sub-directories in the `rds migration tool` directory.

```
ls /opt/amazon/rdsmigrationtool
```

Only the `tmp` and `data` directories should be listed.

Changing the Data Directory Location on Linux

This section explains how to change the location of the Amazon RDS Migration Tool Data Directory. Such a procedure may need to be performed if the drive on which the current folder resides has insufficient space or if you are moving from a temporary POC setup to production, for example.

To change the Data Directory Location:

1. Open the Linux terminal and issue the following command:

```
cd /opt/amazon/rdsmigrationtool/bin
```

2. Stop the `repctl` services on the Linux by running:

```
/opt/amazon/rdsmigrationtool/bin/arepctl stop
```

3. Make sure all the `repctl` services have stopped by running:

```
ps -ef | grep repctl
```

4. Move the data folder from current location (`/opt/amazon/rdsmigrationtool/data`) to your desired location.

5. Create a file named `site_arep_login.sh` in the Amazon RDS Migration Tool bin folder.

6. Add the following command to the file:

```
export AREP_DATA=<new data folder path>
```

Example:

```
export AREP_DATA=/opt/sys232/repdata
```

7. Start the Amazon RDS Migration Tool Server services as normal (see [Starting the Amazon RDS Migration Server Process](#)).

Accessing the Amazon RDS Migration Console

You browse to the Amazon RDS Migration Console using a supported Web browser from a computer in the same network as the computer with the Amazon RDS Migration Server. For information on supported browsers, see [Additional Required Software](#).

You can access the Console from the **Start** menu of the computer where you installed Amazon RDS Migration Tool.

To enable and control access to Amazon RDS Migration Tool, you can create user roles as described in [Security Roles](#).

To access Amazon RDS Migration Tool

Click **Start** and from the **All Programs** section point to Amazon RDS Migration Tool and select **Amazon RDS Migration Console**.

If you are using the Amazon RDS Migration Server on Linux, you must access the Console from a Windows computer in the same network. See the note below for information on how to access Amazon RDS Migration Tool from a remote computer.

Notes:

- You can also access Amazon RDS Migration Tool from any computer in your network.

To access the Console from a remote computer, type the following address in the address bar in your Web browser:

```
https://<computer name>/amazonrdsmigrationconsole
```

Where *<computer name>* is the name of the computer, including the Windows domain name, where you installed Amazon RDS Migration Tool.

- The person logged in to the computer where you are accessing the Console must be authorized as an Amazon RDS Migration Tool user. For more information, see [Security Roles](#).
-
-

To work with Amazon RDS Migration Tool over a secure connection (strongly recommended), you need to set it up for `https` and replace `http` with `https` in the URL above. In this case, you must change the default URL address. For more information on setting up HTTPS connections, see [Appendix B, "Using Encryption"](#).

The default URL address is defined in a file called `ServiceConfiguration.xml`. This file is located on the computer where you installed Amazon RDS Migration Tool in the following folder:

```
C:\Program Files\Amazon\RDS Migration Tool\data
```

Multiple Users Connecting to a Single Console

Multiple users can connect to a single Amazon RDS Migration Console using a Web browser.

1. Install Amazon RDS Migration Tool on the computer that will serve as the Amazon RDS Migration Console.
2. If Amazon RDS Migration Server is installed on another computer (Linux for example), on the console machine, edit the Amazon RDS Migration Tool URL (and port if required) in the `ServiceConfiguration.xml` file to point to that machine.

By default, the file is located in the following folder:

```
C:\Program Files\Amazon\RDS Migration Tool\data
```

3. Open the Windows Services console and restart the **Amazon RDS Migration Console** service.
4. Connect as described in [Accessing the Amazon RDS Migration Console](#) above.

2



Getting Started - An Amazon RDS Migration Tool Tutorial

This chapter provides a tutorial of how to set up a basic Replication task. This task will show how to set up a task that copies the data from an Oracle source database and replicates the database in Microsoft SQL Server.

Before you begin to work with this tutorial, see [What You Need](#).

Then do the following:

- [Open the Amazon RDS Migration Console](#)
- [Add an Oracle Database as a Source](#)
- [Add a Microsoft SQL Server Database as a Target](#)
- [Add a Replication Task](#)
- [Run and Monitor the Replication Task](#)
- [View the Replicated Tables in Microsoft SQL Server](#)

What You Need

In this tutorial you will need the following:

- Amazon RDS Migration Tool installed on a computer in your network.
- For the Oracle source, Oracle Database with the default:
 - Access to the HR schema tables that are part of the Oracle Database installation.

Note: If these tables are not available see your Oracle Database Administrator to resolve the issue.

- `system/<password>` for an admin user
- For the target, Microsoft SQL Server database with:
 - The default `tempdb` system database (used to store the target tables). This can be installed on your local computer.

For additional information about installing the components necessary for working with this tutorial and working with Amazon RDS Migration Tool, see the [Installation Prerequisites](#).

Open the Amazon RDS Migration Console

To start working with Amazon RDS Migration Tool, you need to be sure that you have one of the following Browsers:

- Microsoft Internet Explorer version 8 and above
- Firefox version 3.6 and above
- Google Chrome

To open the Amazon RDS Migration Console:

From the Windows **Start** menu, select **All Programs > Amazon RDS Migration Tool > Amazon RDS Migration Console**.

Notes:

- You can access Amazon RDS Migration Tool from any computer in your system.

To access the Console from a remote computer, type the following address in the address bar in your Web browser:

```
http://<computer name>/amazonrdsmigrationconsole
```

Where *<computer name>* is the name or IP address of the computer (including the Windows domain name) on which the Amazon RDS Migration Server is installed.

- The person logged in to the computer where you are accessing the Console must be authorized as an Amazon RDS Migration Tool user. For more information, see [Security Roles](#).
-
-

Add an Oracle Database as a Source

Once you have attached a server, you begin to add the databases you are working with.

To add a database

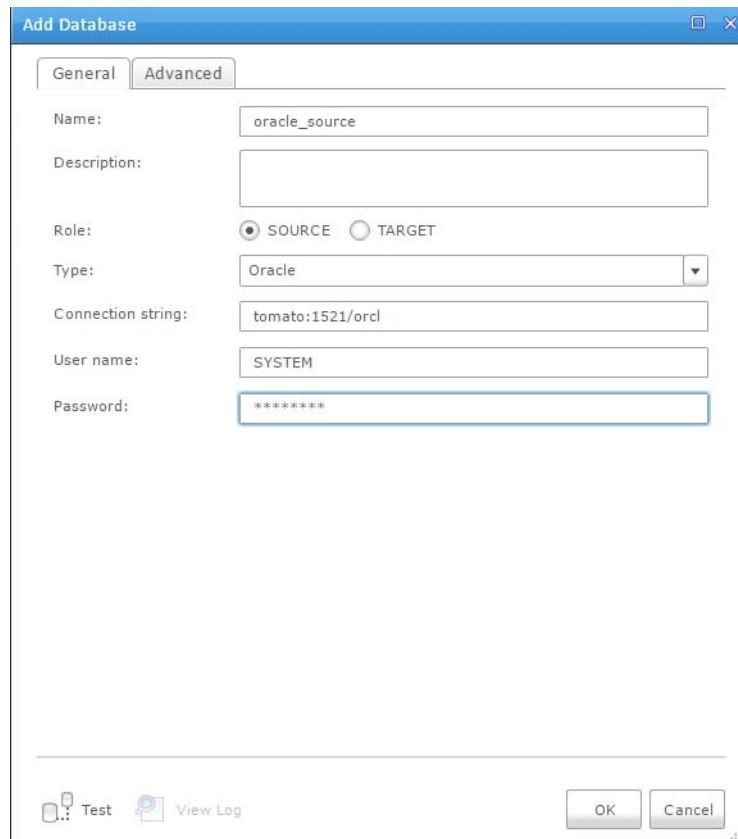
1. In Tasks view, click **Manage Databases**.

The **Manage Databases** dialog box opens.

2. Click **Add Database** toolbar button.

The following figure shows the data that you need to enter in the Add Database dialog box when you add an Oracle source database in this tutorial.

Figure 2–1 Add an Oracle Source Database



To add an Oracle Source database

1. In the **Add Database** dialog box, enter the following information:

- **Name:** oracle_source.
- **Description:** Leave this blank.
- **Role:** Select **Source**.
- **Type:** Select **Oracle** from the list.
- **Connection string:** Type the connect string to the Oracle database you are working with. Type the connect string in any Oracle format.

For example, if you are connecting to an Oracle database on a computer called **tomato**, using the default Oracle port and default service name, the connect string looks like this:

```
tomato:1521/orcl
```

- **User Name:** Type **SYSTEM**. This is the Oracle default administrator user name. If you know that the Oracle database you are working with has a different user name, then type that user name.
- **Password:** Type **manager**. This is the Oracle default administrator password. If you know that the Oracle database you are working with has a different password, then type that password.
- Click **Test** to see if you entered the information correctly and whether the Oracle database is available. Then click **OK** to add the database.

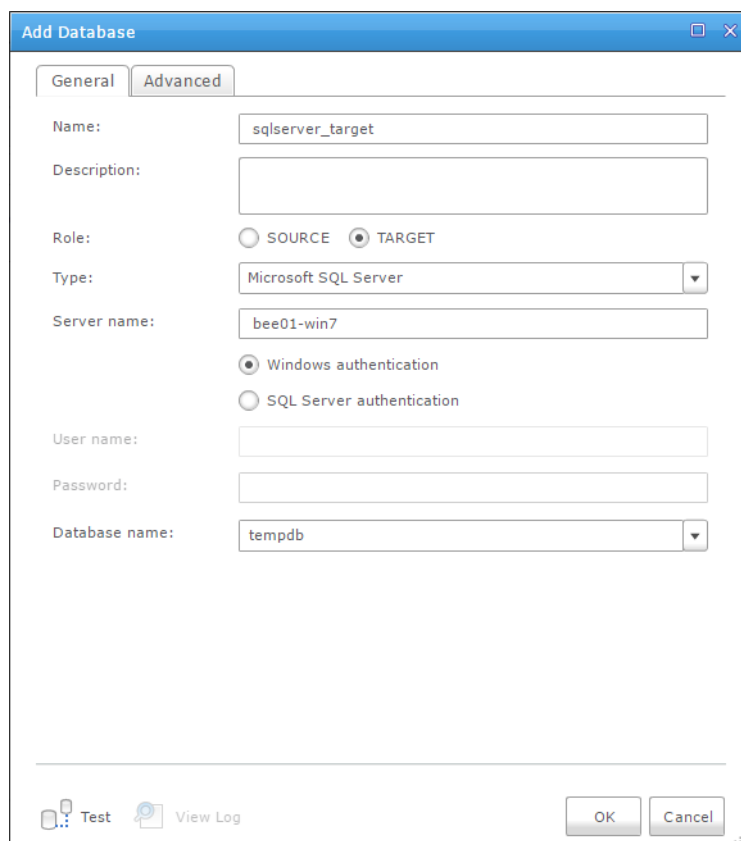
You can also set advanced settings for the Oracle database, although this is not necessary for this tutorial. For information on the advanced settings for an Oracle database, see [Using Advanced Properties for an Oracle Source](#).

For information on adding other types of databases, see the chapter for the database you are using. For a list of supported databases, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

Add a Microsoft SQL Server Database as a Target

To add a target database, return to the Task view and follow the directions as in [To add a database](#). This time you will add a Microsoft SQL Server database to use as a target. The following figure shows the Add Database dialog box with the information necessary for adding a SQL target.

Figure 2–2 Add a Microsoft SQL Server Target Database



To add a Microsoft SQL Server target database

1. From the top left of the Amazon RDS Migration Console, click **Add Database**.
2. In the Add Database dialog box, type the following:
 - **Name:** sqlserver_target.
 - **Description:** Leave this blank.
 - **Role:** Select **Target**.
 - **Type:** Select **Microsoft SQL Server** from the list.

- **Microsoft SQL Server name:** Type the name of the computer where the Microsoft SQL Server database you are working with is installed. For example if the computer where the Microsoft SQL Server database is installed is called `bee`, then type `bee`.
- Select **Windows Authentication**. In this case be sure that the Microsoft SQL Server database you are using is configured to accept Windows authentication. Leave **Username** and **Password** blank.

Note: If the Microsoft SQL Server database you are working with is not configured to accept Windows authentication, then select **Username/password** and type the correct user name and password for the Microsoft SQL Server database you are using.

- In the **Database name** type or select `tempdb`. This is the name of the database to where you are replicating the data.
Note: You can create a new database in Microsoft SQL Server to use as your target. If you do this, type or select the name of the database you created.
- Click **Test** to verify that you entered the information correctly and that the Microsoft SQL Server database is available. Then click **OK** to add the database

You can also set advance settings for the Microsoft SQL Server database, although this is not necessary for this tutorial. For information on the advanced settings for a Microsoft SQL Server database, see [Using Advanced Properties for a Microsoft SQL Server Source Database](#).

For information on adding other types of databases, see the chapter for the database you are using. For a list of supported databases, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

Add a Replication Task

Once you have added a source and a target database, you can begin to define the replication task. Note that if you have not already added a source and target database, you have the option to add them when you define the replication task.

In this tutorial, you will define a task that will copy the data from the `HR.EMPLOYEES` and `HR.JOBS` tables. The `HR` schema is provided by default in your Oracle database. You will make a copy of the same tables in your Microsoft SQL Server `tempdb`. The `EMPLOYEES` and `JOBS` tables created in Microsoft SQL Server will be identical to the Oracle tables. For information on how to use Transformations and Filters when creating a replication task, see [Define Transformations on a Single Table](#) and [Using Filters](#).


Adding a replication task has the following steps:

- [Add a Replication Task to the Amazon RDS Migration Console](#)
- [Add the Source and Target Databases to the Task](#)
- [Select Tables for the Replication Task](#)

Add a Replication Task to the Amazon RDS Migration Console

To add a replication task, make sure **Tasks** is selected in the upper left corner of the Amazon RDS Migration Console as shown in the following figure.

Figure 2–3 *Tasks*



Click **New Task** to open the New Task dialog box. In the **Name** field for this dialog box, type `Task_1`. Leave the **Description** field blank and click **OK**.

The Amazon RDS Migration Console displays a diagram in the middle where you add the source and target databases used in the task. In addition a pane on the right side lets you select that tables you want to work with and carry out transforms and filtering operations. For additional information, see [The Tasks View](#), [Viewing Specific Tasks](#) and [Designing Tasks](#).

You can also make changes to the default task settings if you need to. For information on changing the task settings, see [Task Settings](#).

Add the Source and Target Databases to the Task

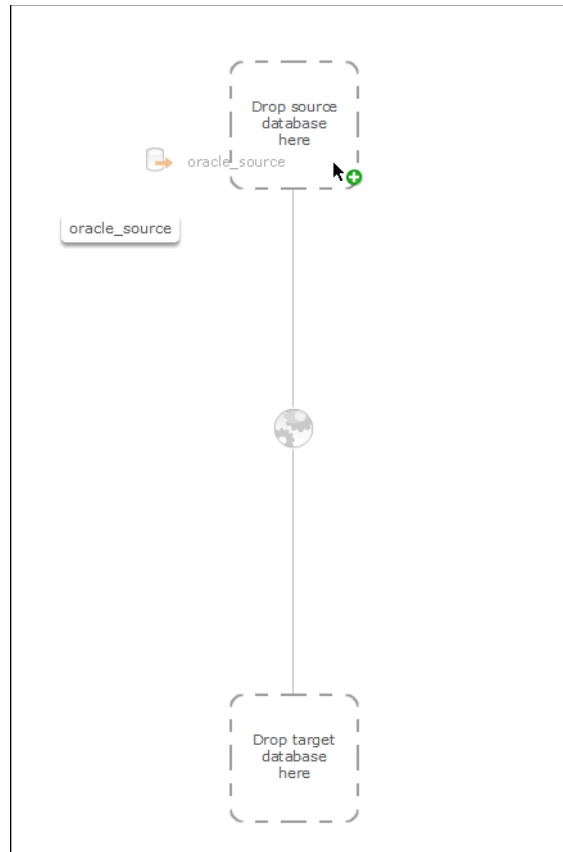
To add a source and target to the replication task, you drag a source database from the left pane to the Source area and a target database to the Target area in the task diagram. The source databases icon consists of a database with an orange arrow pointing away from the database whereas the target database icon consists of a database with a blue arrow pointing towards the database.

Since we named our databases in a way that we can recognize whether it is a source or target database we can display All Databases.

To add the source or target databases to the task

1. Drag the `oracle_source` database to the source area in the task diagram as shown in the figure below:

Figure 2-4 Create a Task



2. Drag the `sqlserver_target` database to the Target area as shown in the figure in the previous step.

When you have added the source and target databases to the task, you can then select the tables from the source database to use in the replication task.

For more information, see [Designing Tasks](#).

Select Tables for the Replication Task

You select specific tables from the source database with the data you want to copy to the target database. Amazon RDS Migration Tool lets you copy data from one type of database to another. For example, in this tutorial we will copy data from an Oracle database into a Microsoft SQL Server database.

In this tutorial we will explicitly select the `HR.EMPLOYEES` and `HR.JOBS` tables from the Oracle source. This will take all of the data from these tables "as is" and copy them to the Microsoft SQL Server target.

Notes:

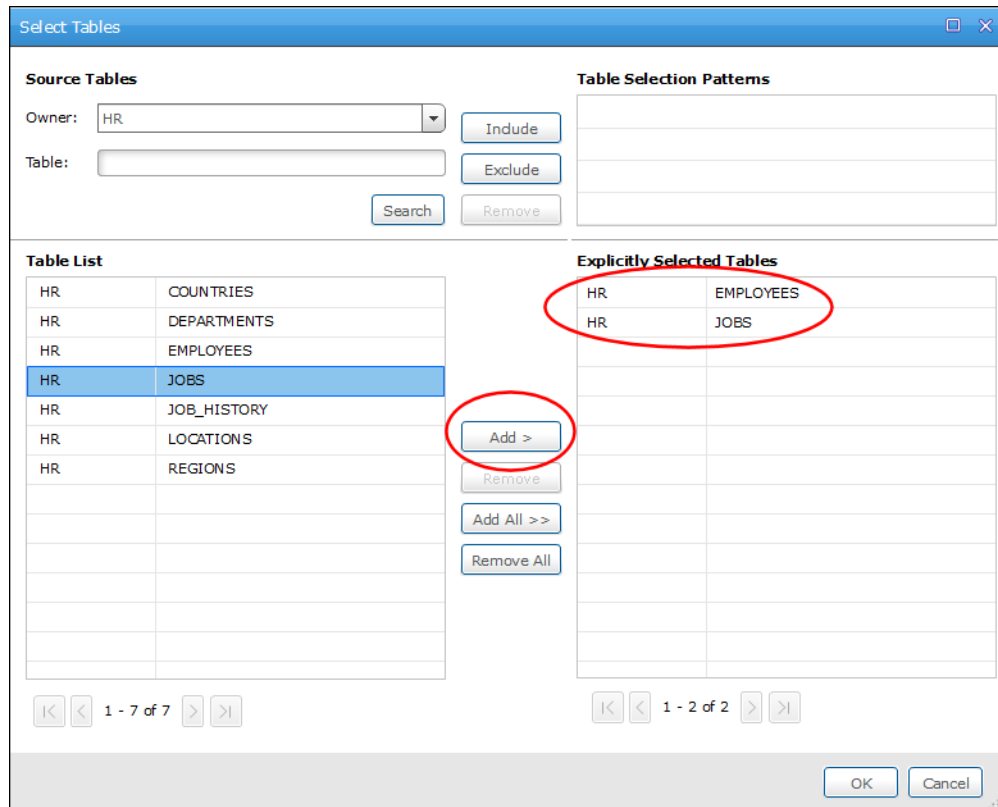
There may be circumstances where you need to copy only some of the data to the target database. In this case you need to use a filter. For information on using filters, see [Using Filters](#).

There may be circumstances where you need to copy the data into the target using different rows or columns than those in the source. In this case you need to use transforms. For information on using transforms, see [Define Transformations on a Single Table](#).

To add the tables to the replication task

1. In the right pane of the Amazon RDS Migration Console, click **Table Selection**. The **Select Tables** dialog box is displayed.

Figure 2–5 Add the Tables to the Task



2. Do the following in the Select Tables dialog box:
 - In the **Owner** field, select **HR** from the list, then click **Search**.
 - From the **Table List**, select **HR.EMPLOYEES** and then click the **ADD** to select that table. Repeat this for the **HR.JOBS** table.
 - Click **OK**.

In this tutorial we will copy the data "as is." There is no need to use transformation, filters, or select table patterns. For information on using these features, see [Adding Tables to a Task](#).

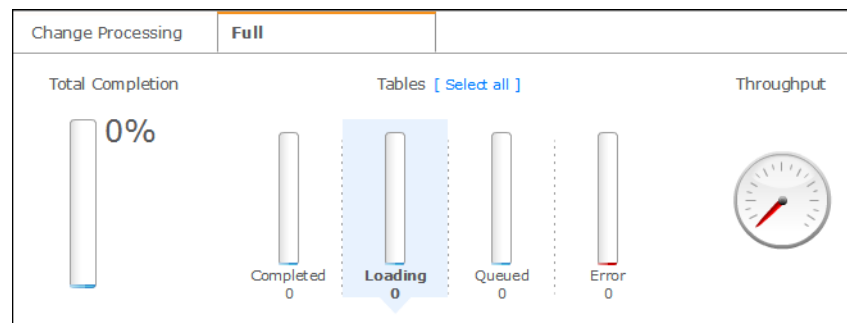
Run and Monitor the Replication Task

Once you create the replication task, you can run the task and see the results of the replication in real time. In this tutorial, you will run the replication task as a full load and view the progress in the Monitor. Additional run options are also available. For information on the available run options, see [Using the Run Button Options](#).

To run and monitor the replication task

1. Select **Tasks** from the list at the top left of the Amazon RDS Migration Console. If Tasks is already selected, you can click the button.
2. Select **Task_1** from the task list. This may be the only available task if you are just beginning to use Amazon RDS Migration Tool.
3. Click **Open** from the button bar at the top of the Task list.
4. From the right side of the Amazon RDS Migration Console, click **Monitor**. The monitoring status is displayed using various gauges and graphs. These are displayed on the left side of the console as shown in the figure below.

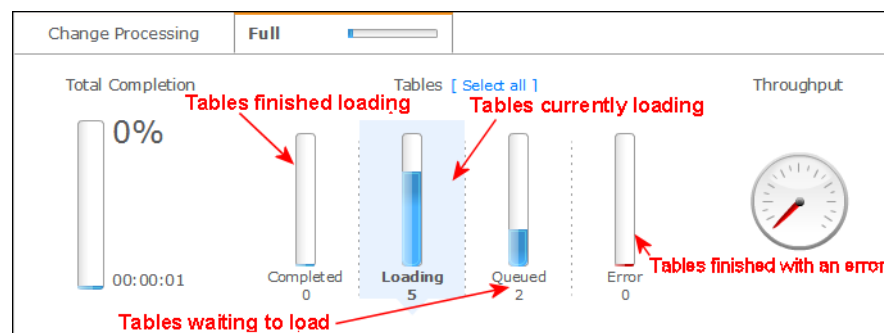
Figure 2–6 Monitoring a Task



5. In the toolbar at the top of the console you will see the run buttons. Click **Run**.
6. Note the gauges and bars as the task runs. You will see the loading and error information in the bar graphs and the throughput gauge.

The left pane has two tabs. In addition to the Full-Load tab, that we just saw, Change Processing tab shows information about real-time changes that occur after the full load is complete. For information on reading the data presented in these sections, see [Viewing Information in the Monitor](#). The following figure shows the Full-Load monitoring while a task is running.

Figure 2–7 Task Running



7. At the top of the Full Load graph area, click the **Select all** link. A table is displayed below the graphs with information about each of the tables being processed in the task.
8. Click **Stop** to stop all processing for your task.
9. Click the **Run** button arrow and select **Resume processing**. Click **OK** in the dialog box that appears and watch the data change for each table.

In the figure below, you can see that the HR.REGIONS, HR.LOCATIONS and HR.JOBS tables completed processing. The other tables are still being processed as shown by the incomplete bar graph in the **Progress** column.

Figure 2–8 Status Table

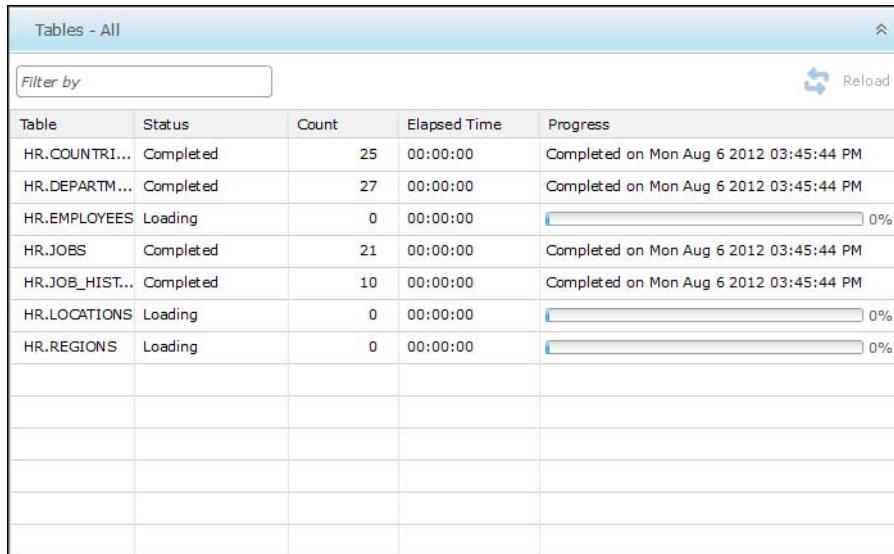


Table	Status	Count	Elapsed Time	Progress
HR.COUNTRI...	Completed	25	00:00:00	Completed on Mon Aug 6 2012 03:45:44 PM
HR.DEPARTM...	Completed	27	00:00:00	Completed on Mon Aug 6 2012 03:45:44 PM
HR.EMPLOYEEES	Loading	0	00:00:00	<div style="width: 0%; border: 1px solid #ccc;">0%</div>
HR.JOBS	Completed	21	00:00:00	Completed on Mon Aug 6 2012 03:45:44 PM
HR.JOB_HIST...	Completed	10	00:00:00	Completed on Mon Aug 6 2012 03:45:44 PM
HR.LOCATIONS	Loading	0	00:00:00	<div style="width: 0%; border: 1px solid #ccc;">0%</div>
HR.REGIONS	Loading	0	00:00:00	<div style="width: 0%; border: 1px solid #ccc;">0%</div>

10. Click on the other bar graphs, such as the Completed graph and the Loading graph to see additional information.
For information about the data supplied in these tables, see [Monitoring the Full-Load Operation](#).

View the Replicated Tables in Microsoft SQL Server

In this part of the tutorial, you will view the tempdb database in Microsoft SQL Server. You will note that two new tables, HR.EMPLOYEEES and HR.JOBS now exist in this database.

To view the replicated tables in Microsoft SQL Server

1. From the Windows **Start** menu, go to **All Programs, Microsoft SQL Server**, then select **Microsoft SQL Server Management Studio**.
2. In the Object Explorer, find the Microsoft SQL Server target computer you are working with. Expand the **Databases** folder for that computer, then expand the **System Databases** folder, then expand the tempdb database. The EMPLOYEEES and JOBS tables should now appear in the list.
3. Right-click the EMPLOYEEES table and select **Select Top 1000 Rows**. Check that there is data in the table.

4. Right-click the JOBS table and select **Select Top 1000 Rows**. Check that there is data in the table.

Concepts and Design

This section describes the main concepts of data replication and shows the main components of Amazon RDS Migration Tool. It has the following topics:

- [Introducing Amazon RDS Migration Tool](#)
- [Amazon RDS Migration Tool Features](#)
- [What is Replication](#)
- [System Architecture](#)
- [What is a Replication Task](#)
- [Full Load and CDC Processes](#)
- [Replication Topologies](#)

Introducing Amazon RDS Migration Tool

Amazon RDS Migration Tool consists of a Web-based console and a replication server to replicate data across heterogeneous data sources. These capabilities provide users with instant visibility to current and historical exceptions, status, performance, and resource usage information.

Amazon RDS Migration Tool can execute replication between enterprise databases including Oracle, Microsoft SQL Server, and IBM DB2. It uses a "Click-2-Replicate" design that simplifies the replication process by automating the steps required to build a replication solution. The Amazon RDS Migration Tool software is easy to learn and to implement.

Replication is log based, which means that only the changes are read. This reduces the impact on the source databases.

Amazon RDS Migration Tool can carry out two types of replication: Full Load and Change Processing (CDC).

Amazon RDS Migration Tool Features

Amazon RDS Migration Tool is a simple but powerful solution that provides replication between various databases. Amazon RDS Migration Tool lets you do the following:

- Load data efficiently and quickly to operational data stores/warehouses.
- Create copies of production databases.
- Distribute data across databases.

Amazon RDS Migration Tool has high throughput, speed, and scale. It is designed to scale and support large scale enterprise data replication scenarios with a multi-server, multi-task, and multi-threaded architecture.

When you run an Amazon RDS Migration Tool task, two types of replication can occur:

- **Full Load:** The full load process creates files or tables at the target database, automatically defines the metadata that is required at the target, and populates the tables with data from the source.
- **Change Processing (CDC):** Change processing captures changes in the source data or metadata as they occur and applies them to the target database as soon as possible in near-real-time.

What is Replication

Replication is a procedure whereby two or more collections of computerized information are kept identically synchronized. Possible reasons for synchronization include:

- **Load reduction:** It is recommended that you have a copy of all or of a subset of a collection on a different server to reduce the load on the main server.
- **Improved service:** Users of the copy of the information may get better access to the copy of the data than to the original.
- **Security considerations:** Some users might be allowed access to a subset of the data and only this subset is made available as a replicated copy to those users.
- **Geographic distribution:** The enterprise (for example, a chain of retail stores or warehouses) may be widely distributed and each node uses primarily its own subset of the data (in addition to all of the data being available at a central location for less common use).
- **Disaster Recovery:** A copy of the main data is required for rapid failover (the capability to switch over to a redundant or standby computer server, in case of failure of the main system).
- **Support the need for implementing "cloud" computing.**

The information used is usually stored as files or in a database. In the case of files, the structure and contents of a file are known only to the specialized programs that use the file. Databases are managed by database management systems (DBMS), that make use of standardized descriptions of the structure of the information (such as tables, columns, rows, and data types). These descriptions are known collectively as metadata, and allow a general-purpose replicator to carry out relevant operations (for example, filtering, data transformations) without need to know anything about the contents or "meaning" of the data. Because file systems do not contain metadata, operations available for replication are more limited.

During replication, a collection of data is copied from system A to system B. A is known as the source (for this collection), B is known as the target. A system can be either a source or a target or even both (within certain restrictions). When a number of sources and targets and data collections are defined, the replication topology can be quite complex.

Since copies cannot be made or maintained instantaneously, and since it is assumed that the source computer cannot be stopped and the information cannot be kept "frozen," replication must take into account that the source data may be changing while it is being copied. This problem includes three main issues:

- Integrity: Make sure that the data in the target actually reflects the completed result of a change in the source and not some intermediate invalid result.
- Consistency: Make sure that if the change affects several different tables or rows, the copy reflects a consistent state all were changed or none).
- Latency: How out-of-date is the copy?

The first two issues are the responsibility of the replicator. While some latency is unavoidable in any system, a good replicator will aim not to exceed several seconds of latency as a general rule.

System Architecture

In the initial load process, Amazon RDS Migration Tool reads a filtered stream of rows (with relevant columns only) and passes them to the transformation process for further filtering writing to the target database (in the expected output format).

The CDC (Change Data Capture) process obtains a stream of filtered events or changes in data or metadata from the transaction log file. One of its most important functions is to buffer all of the changes for a given transaction into a single unit before they can be sent forward to the target when the transaction commits. During the initial load process CDC also buffers all of the changes that occur within a transaction until after all of tables affected have finished being loaded.

The Designer/Console server, part of the Replication server, is a Web-based application that serves all of the user Interface dealing with designing or modifying the replication system as well as displaying and controlling its operation.

What is a Replication Task

An Amazon RDS Migration Tool task is one instance of table synchronization activity. It must be defined by the user or database administrator, using the browser-based Amazon RDS Migration Console. The definition of a task consists of:

- Specifying the source and target databases
- Specifying the source and target tables to be kept in sync
- Specifying the relevant source table columns
- Specifying filtering conditions (if any) for each source table, as Boolean predicates on the values one or more source columns (the predicates are in SQLite syntax)
- Listing the target table columns and (optionally) specifying their data types and values (as expressions or functions over the values of one or more source or target columns, using SQL syntax). If not specified, the same column names and values as the source tables are used, with default mapping of the source DBMS data types onto the target DBMS data types. Amazon RDS Migration Tool automatically takes care of the required filtering, transformations and computations during the Load or CDC execution.

The simplest specification of a task may not mention of the target data, with only the source tables (or ALL, or a mask) specified. In this case, the target tables are identical to the source tables, using the default mappings between the source and target DBMS data types. In this way, the entire definition process could be accomplished by a single click, referred to as "Click to Replicate".

Once a task is defined, it can be activated immediately. The target tables with the necessary metadata definitions are automatically created and loaded, and the CDC is

activated. The replication activity can then be monitored, stopped, or restarted using the Amazon RDS Migration Console.

Using Multiple Tasks

Several different replication tasks can be defined and activated. This best if:

- The source tables for the tasks are different
- Some of the tables in the source may be the same, but the filtering conditions on the source rows are different
- The target tables updated by the tasks are all different.

Updating the same target table and row by two different replication tasks would not be a good practice and may cause unpredictable results.

The different replication tasks work independently and run concurrently. Each one has its own Initial Load, CDC, and Log Reader processes.

Full Load and CDC Processes

The full load process creates files or tables at the target database, automatically defines the metadata that is required at the target, and populates the tables with data from the source. Unlike the CDC process the data is loaded one entire table or file at a time for efficiency purposes. It is important to understand that the source tables may be subject to update activity during the Load process. *There is no need to stop processing in the source.* Because of this, and to guarantee the consistency of the data, the CDC process is automatically activated as soon as the Load starts, but the changes are not applied at the target until after the load of a table completes. The data on the target might not be consistent while the Load is active, but consistency and integrity of the target data at the conclusion of the Load are guaranteed.

The Load process can be interrupted and when restarted it continues from wherever it was stopped. New tables can be added to an existing target without reloading the existing tables. Similarly, columns in previously-populated target tables can be added or dropped without requiring reloading.

The Change Data Capture (CDC) process captures changes in the source data or metadata as they occur and applies them to the target database as soon as possible in near-real-time. The changes are captured and applied as units of single committed transactions, and several different target tables can be updated as the result of a single source commit. This guarantees transactional integrity in the target database. The CDC process for any file or table starts as soon as the data load operation for the file or table begins.

CDC operates by reading the recovery log file of the source database management system and grouping together the entries for each transaction. Various techniques are employed to ensure that this is done in an efficient manner without seriously impacting the latency of the target data. If the changes cannot be applied to the target within a reasonable time (for example, if the target is not accessible), the changes are buffered on the Replication server for as long as is necessary. There is no need to re-read the source DBMS logs, which may take a large amount of time.

Replication Topologies

A replication task can be set up in various ways. The following topologies are currently supported by Amazon RDS Migration Tool.

- [One-to-One](#)

- [Logical Independence](#)
- [Hub and Spoke](#)

One-to-One

In a one-one topology, there is one source and one target database. Using two different replication tasks, the databases may switch roles, allowing two-way synchronization.

Note: If the same row in a table is updated by two different replication tasks, the result of a two-way synchronization may be unpredictable. A problem can occur even if two different rows are referentially-related, that is if some application updates a row based on reading a value in a different row. If the rows are updated concurrently on the source and the target, the result may be unpredictable¹. Such occurrences are rare, but can occur.

¹ CDC has no way of knowing exactly when a row was read by an application on one system relative to its having been changed on another system. Read operations are typically not logged.

When the source and target databases are distinct (one-way replication) concurrent-update problems do not occur, and transactional integrity and consistency can be guaranteed.

Logical Independence

Two-way replication works best when updates of a row on a source and on a target are entirely autonomous and do not affect each other. There is an assumption that any table or a horizontal or a vertical segment of a partitioned table can only be updated in one source. We allow the same row to be updated at several places, but in this case the columns being updated must be distinct. Another assumption is that if a data value in one row depends on or is derived from a value in another row, the values can be changed only on the same server but nowhere else (except by the Replicator). This is called logical independence.

Under logical independence assumptions, no concurrent update conflicts can occur during replication.

Hub and Spoke

Many-to-one and one-to many can be combined into a hub-and-spoke topology, where data can be merged into multiple targets and then distributed to other targets. No cycles or multiple paths for propagating changes are allowed. The replication topology is actually one of an acyclic directed graph.

Introducing Amazon RDS Migration Tool Endpoints

Amazon RDS Migration Tool lets you work with many databases that you already have. There is no need to install any additional software.

This chapter includes the following topics:

- [Databases Supported as Amazon RDS Migration Tool Endpoints](#)
- [Amazon RDS Migration Tool Data Types](#)
- [Amazon RDS Migration Tool Supported DDL Statements](#)

Databases Supported as Amazon RDS Migration Tool Endpoints

Amazon RDS Migration Tool can replicate data from many types of databases and copy it to various targets. To use these databases as an Amazon RDS Migration Tool endpoint, you do not need to install any other Amazon software. Simply install the databases and install Amazon RDS Migration Tool. That's all.

An Amazon RDS Migration Tool endpoint can be either a source or a target. A source database contains the original data, the data you want to copy. A target database is where the replicated data is stored. Amazon RDS Migration Tool makes it possible to take data from one database and replicate the data into a completely different database.

You can also use a various file types as both a source or a target.

The following chapters provide information on using Amazon RDS Migration Tool endpoints.

- [Using an Oracle Database as a Source or Target](#)
- [Using a Microsoft SQL Server Database as a Source or Target](#)
- [Using a SAP Sybase ASE Database as a Source or Target](#)
- [Using a MySQL Database as a Source or Target](#)
- [Using ODBC to Connect to a Source](#)
- [Using Amazon Redshift as a Target](#)
- [Using the Amazon RDS Migration Tool File Channel](#)

Amazon RDS Migration Tool Data Types

Source data in an Amazon RDS Migration Tool task is converted to an Amazon RDS Migration Tool data type. Data that is not convertible to an Amazon RDS Migration Tool data type is not supported and an error is returned.

To see how a data type is mapped from source to target:

1. See the chapter for the source target database you are using. In the section on data types, see the mapping table to see the Amazon RDS Migration Tool data type.
2. See the chapter for the target database you are using. In the section on data types, see the mapping table to see how the Amazon RDS Migration Tool data type maps to the target.

The following example describes this procedure.

When replicating data from an Oracle source database to a Microsoft SQL Server target database, see the Oracle section to view how the data types are mapped from Oracle to Amazon RDS Migration Tool data types. For example, the Oracle data type `BINARY` maps to the Amazon RDS Migration Tool data type `BYTES`. For the Microsoft SQL Server target, the Amazon RDS Migration Tool data type `BYTES` maps to the Microsoft SQL Server data type `VARBINARY (Length)`.

The following table describes the Amazon RDS Migration Tool data types. Some of the data types have precision and scale information that applies to them.

Table 4–1 Amazon RDS Migration Tool Data Types Description

Amazon RDS Migration Tool Data Types	Description
STRING	A character string.
WSTRING	A double-byte character string.
BOOLEAN	A Boolean value.
BYTES	A binary data value.
DATE	A date value: Year, Month, Day.
TIME	A time value: Hour, Minutes, Seconds.
DATETIME	A timestamp value: Year, Month, Day, Hour, Minute, Second, Fractional Seconds. The fractional seconds have a maximum scale of 9 digits.
INT1	A one-byte, signed integer.
INT2	A two-byte, signed integer.
INT4	A four-byte, signed integer.
INT8	An eight-byte, signed integer.
NUMERIC	An exact numeric value with a fixed precision and scale.
REAL4	A single-precision floating-point value.
REAL8	A double-precision floating-point value.
UINT1	A one-byte, unsigned integer.
UINT2	A two-byte, unsigned integer.
UINT4	A four-byte, unsigned integer.
UINT8	An eight-byte, unsigned integer.

Table 4–1 (Cont.) Amazon RDS Migration Tool Data Types Description

Amazon RDS Migration Tool Data Types	Description
BLOB	<p>Binary Large Object.</p> <p>This data type can be used only with Oracle or Microsoft SQL Server endpoints.</p> <p>For a list of endpoints that support using LOBs, see Support for Large Object Data Types (LOBs).</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
CLOB	<p>Character Large Object.</p> <p>For a list of endpoints that support using LOBs, see Support for Large Object Data Types (LOBs).</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
NCLOB	<p>Native Character Large Object.</p> <p>For a list of endpoints that support using LOBs, see Support for Large Object Data Types (LOBs).</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>

Support for Large Object Data Types (LOBs)

Amazon RDS Migration Tool provides full support for using large object data types (BLOBs, CLOBs, and NCLOBs).

The following source endpoints have full LOB support:

- Oracle (See [Oracle Source Data Types](#) for more information)
- Microsoft SQL Server (See [Microsoft SQL Server Source Data Types](#) for more information)
- ODBC (See [ODBC Source Data Types](#) for more information)

The following target endpoints have full LOB support:

- Oracle (See [Oracle Target Data Types](#) for more information)
- Microsoft SQL Server (See [Microsoft SQL Server Target Data Types](#) for more information)

The following target endpoints have limited LOB support. You cannot use an unlimited LOB size for these target endpoints. For information on how to set the LOB size for a task, see [Task Settings/Metadata](#).

- Amazon Redshift (See [Amazon Redshift Data Types](#) for more information)

For endpoints that have full LOB support, you can also set a size limit for LOB data types. For information on how to do this and on how to configure a task for using LOB data types, see [Task Settings/Metadata](#).

Amazon RDS Migration Tool Supported DDL Statements

Amazon RDS Migration Tool automatically changes the metadata of the target table to reflect DDL statements performed on the source endpoint.

Supported DDL statements include:

- Create table
- Drop table
- Rename table
- Add column
- Drop column
- Rename column
- Change column data type

For information about which DDL statements are supported for a specific endpoint, refer to the chapter describing that endpoint.

For more information about DDL settings, see [Apply Changes Settings](#).

Using an Oracle Database as a Source or Target

Amazon RDS Migration Tool supports two types of Oracle Agents, a LogMiner agent and a binary agent. The configurations that you make to your Oracle database to work with Amazon RDS Migration Tool depend on which agent you are using.

This chapter describes how to set up and use an Oracle database as a source or target endpoint in a replication task.

The chapter contains the following topics:

- [Supported Oracle Database Editions](#)
- [Limitations](#)
- [Using an Oracle Database as a Source](#)
- [Using an Oracle Database as a Target](#)

Supported Oracle Database Editions

Before you begin to work with an Oracle database as a source or target in Amazon RDS Migration Tool, make sure that the Oracle database with the tables that are necessary for replication is available in your system. Amazon RDS Migration Tool supports the following Oracle database editions:

- Oracle Enterprise Edition
- Oracle Standard Edition
- Oracle Express Edition
- Oracle Personal Edition

Note:

- The Oracle database is supported on all operating systems and platforms.
 - The Oracle database can be installed on any computer in your network.
 - An Oracle account with the specific access privileges is required. See [Security Requirements](#) for more information.
 - If you are using the Oracle endpoint with log miner, you must set up supplemental logging as described in [Setting up Supplemental Logging](#).
-
-

Install an Oracle client on the computer where Amazon RDS Migration Tool is located. Install the following:

- On Windows systems, install Oracle Instant Client for Microsoft Windows (x64) Version 11.2.0.3.0.

Note: Support for the XMLTYPE data type requires the full Oracle Client.

- On Linux systems, install Oracle Instant Client for Linux (x86-64) Version 11.2.0.3.0.

Note: Support for the XMLTYPE data type requires the full Oracle Client.

In addition, if not already included in your system, you need to create a symbolic link in the `$Oracle_Home\lib` directory. This link should be called `libclntsh.so`, and should point to a specific version of this file. For example:

- Oracle Version 11 client:

```
lrwxrwxrwx 1 oracle oracle 63 Oct 2 14:16 libclntsh.so ->
/u01/app/oracle/home/lib/libclntsh.so.11.1
```

- Oracle Version 10 client:

```
lrwxrwxrwx 1 oracle oracle 63 Oct 2 14:16 libclntsh.so ->
/u01/app/oracle/home/lib/libclntsh.so.10.1
```

Additionally, the `LD_LIBRARY_PATH` environment variable should be appended with the Oracle `lib` directory and added to the `site_arep_login.sh` script.

Limitations

- The Amazon RDS Migration Tool Oracle endpoint cannot create a new schema on the Oracle database. Therefore, if you are replicating data to an Oracle target and you want to change the schema name, the new schema name must already exist on the Oracle database. If it does not exist, you must create the schema on the database, then you can use that schema name in Amazon RDS Migration Tool.
- Amazon RDS Migration Tool only supports Oracle TDE tablespace encryption. All other forms of encryption are not supported.
- If you are managing supplemental logging and you carry out transformations on any of the columns, you must be sure that supplemental logging is activated for all fields and columns.

- The AR_H_USER header column is supported only for Oracle database version 11.2.0.3 and higher. In Oracle database version 10, the value for this column may not be correct. For information on using header columns, see [Header Columns](#).
- The rename table `<table name>` to `<new table name>` syntax is supported by Amazon RDS Migration Tool when using Oracle version 11 and higher.
- Data changes resulting from partition/sub-partition operations (ADD, DROP, EXCHANGE and TRUNCATE) will not be replicated. To replicate such changes, you need to reload the table. Any future data changes to newly added partitions will be replicated to the target *without* needing to reload the table again. However, operations to UPDATE old data records (in these partitions) will generate a "0 rows affected" warning.
- LogMiner Agent does not support LOB update when only LOB columns have been updated in the specific UPDATE statement.
- On Oracle 11, LogMiner Agent does not support the UPDATE statement for XMLTYPE and LOB columns.
- The DDL statement: "ALTER TABLE ADD <column> <data_type> DEFAULT <>" does not replicate the default value to the target and the new column is set to NULL. If the new column is nullable, Oracle updates all the table rows before logging the DDL itself. As a result, Amazon RDS Migration Tool captures the changes (the counters have changed) but does not update the target. As the new column is set to NULL, if the target table has no Primary Key/Unique Index, subsequent updates will generate a "zero rows affected" message.
- Data changes resulting from the "CREATE TABLE AS..." statement is not supported. However, the new table *will* be created on the target.
- On Oracle 12c, Log Miner does not support LOB columns.
- When Limited-size LOB mode is enabled, empty LOBs on the Oracle source are replicated as NULL values. For more information on Limited-size LOB mode, see [Task Settings | Metadata](#).
- Changes made by the Oracle DBMS_REDEFINITION package - e.g. table metadata and the OBJECT_ID) - will not be captured by Amazon RDS Migration Tool.
- Index-organized tables with an overflow segment are not supported in CDC mode when using Binary Reader (i.e. when not using LogMiner to access the redo logs).
- Table clusters are not supported when using Oracle Binary Reader.
- **Oracle Target:** The **Use direct path full load** option is not supported for tables with INDEXTYPE CONTEXT.
Workaround:
Use Array Load.
- **Oracle Target:** In Batch Optimized Apply mode, loading into the net changes table uses Direct Path which does not support XMLType.
Workaround:
Use Transactional Apply mode.
- **Oracle Target:** Amazon RDS Migration Tool cannot create a new schema on the Oracle database. To replicate to a new schema, the new schema name must already exist on the Oracle target. You can then specify the new schema name in the Task Settings' [Target Metadata](#) and [Control Tables](#) tabs as required.

- **Oracle Source:** Changes cannot be captured from redo logs larger than 4 GB.

Using an Oracle Database as a Source

The following topics provide information pertinent to using an Oracle database as the source endpoint in an Amazon RDS Migration Tool task.

- [Supported Compression Methods](#)
- [Configuration Properties for the Oracle Source](#)
- [Security Requirements](#)
- [Oracle Source Data Types](#)
- [Non-Supported Data Types](#)
- [Configuring an Oracle Database as an Amazon RDS Migration Tool Source](#)

Supported Compression Methods

The table below lists which compression methods Amazon RDS Migration Tool supports when working with an Oracle source endpoint. As the table shows, compression support depends both on your Oracle database version and whether or not Amazon RDS Migration Tool is configured to use LogMiner to access the redo logs.

Table 5–1 Supported Oracle Compression Methods

Version	OLTP	Basic	Others
Oracle 10	n/a	No	No
Oracle 11 and above - Binary Reader	Yes	Yes	No
Oracle 11 and above - LogMiner	Yes	Yes	*Yes

*Any compression method supported by LogMiner

Configuration Properties for the Oracle Source

The following table describes the configuration properties available for an Oracle source. For information on how to set these properties, see [Using Advanced Properties for an Oracle Source](#).

Table 5–2 Oracle Connect String Standard Properties

Property	Description/Format	Default Value
server	The name of the computer with the Oracle instance you are working with. This is written in any recognized Oracle format. server=tomato:1521/orcl	
username	The name of the user authorized to access the Oracle account you are using. username=john	
password	The password for the authorized user. password=<sample>	

Table 5–2 (Cont.) Oracle Connect String Standard Properties

Property	Description/Format	Default Value
addSupplementalLogging	If true (Y) supplemental logging is automatically configured on the Oracle database. In this case, you do <i>not</i> have to manually set up supplemental logging. For more information on setting up supplemental logging, see Setting up Supplemental Logging .	Y
useLogminerReader	Indicates whether changes written to the redo log are read using LogMiner. If true (Y), LogMiner is used. useLogminerReader=Y	Y
asm_server	The connect string to the Automated Storage Management (ASM) server. This is written in any recognized Oracle format. asm_server=tomatoASM This property is used when the Oracle redo logs you are using are stored using Automated Storage Management (ASM). This is not available if you are using the LogMiner to read the logs.	
asm_user	The user name for the ASM user. asm_user=johnj	
asm_password	The password for the ASM user. asm_password=<sampleASM>	
retryInterval	The number of seconds that the system waits before resending a query.	1
numberDatatypeScale	The scale used when reading or writing an Oracle NUMBER data type. For example: numberDatatypeScale=8 uses a scale of 8 when Amazon RDS Migration Tool reads or writes NUMBER data types. Acceptable values are 0 through 38 or Enter a value of -1 to handle a NUMBER data type in Oracle as a FLOAT.	10

Security Requirements

In order to use an Oracle source in an Amazon RDS Migration Tool task, the user specified in the Amazon RDS Migration Tool Oracle database definitions must be granted the following privileges in the Oracle database:

- SELECT ANY TRANSACTION
- SELECT on V_\$ARCHIVED_LOG
- SELECT on V_\$LOG
- SELECT on V_\$LOGFILE

- SELECT on V_\$DATABASE
- SELECT on V_\$THREAD
- SELECT on V_\$PARAMETER
- SELECT on V_\$NLS_PARAMETERS
- SELECT on V_\$TIMEZONE_NAMES
- SELECT on V_\$TRANSACTION
- SELECT on ALL_INDEXES
- SELECT on ALL_OBJECTS
- SELECT on ALL_TABLES
- SELECT on ALL_USERS
- SELECT on ALL_CATALOG
- SELECT on ALL_CONSTRAINTS
- SELECT on ALL_CONS_COLUMNS
- SELECT on ALL_TAB_COLS
- SELECT on ALL_IND_COLUMNS
- SELECT on ALL_LOG_GROUPS
- SELECT on SYS.DBA_REGISTRY
- SELECT on SYS.OBJ\$
- SELECT on DBA_TABLESPACES
- SELECT on ALL_TAB_PARTITIONS
- SELECT on ALL_ENCRYPTED_COLUMNS
- If views are exposed: SELECT on ALL_VIEWS

Grant the following additional privilege (for each replicated table) when you are using a specific table list.

- SELECT on <any-replicated-table>;

Grant the following additional privilege when using a pattern for the table list.

- SELECT ANY TABLE;

Grant the following additional privilege (for each replicated table) when Amazon RDS Migration Tool adds supplemental logging automatically (the default behavior) and you are using a specific table list. For information on how to turn off supplemental logging, see [Using Advanced Properties for an Oracle Source](#).

- ALTER on <any-replicated-table>;

Grant the following additional privilege when Amazon RDS Migration Tool adds supplemental logging automatically (the default behavior). For information on how to turn off supplemental logging, see [Using Advanced Properties for an Oracle Source](#).

- ALTER ANY TABLE;

Access Privileges when using LogMiner to Access the Redo Logs

If you are using LogMiner to access the Redo logs, grant the following privileges.

For Oracle versions before 12c:

- CREATE SESSION
- EXECUTE on DBMS_LOGMNR
- SELECT on V_\$LOGMNR_LOGS
- SELECT on V_\$LOGMNR_CONTENTS

For Oracle 12c or above:

- LOGMINING

For information on how to turn on LogMiner, see [Using Advanced Properties for an Oracle Source](#).

Access Privileges when using Binary Reader to Access the Redo Logs

The following privileges should be granted when using Binary Reader to access the Redo logs:

- SELECT on v_\$transportable_platform

Grant the SELECT on v_\$transportable_platform privilege if the Redo logs are stored in ASM and accessed by RDS Migration Tool from ASM.

- CREATE ANY DIRECTORY

Amazon RDS Migration Tool uses following Oracle file access features:

- BFILE read - Used when RDS Migration Tool does not have file-level access to the Redo logs, and the Redo logs are not accessed from ASM.
- DBMS_FILE_TRANSFER package - Used to copy the Redo log files to a temporary folder (in which case the EXECUTE ON DBMS_FILE_TRANSFER privilege needs to be granted as well)
- DBMS_FILE_GROUP package - Used to delete the Redo log files from a temporary/alternate folder (in which case the EXECUTE ON DBMS_FILE_GROUP privilege needs to be granted as well).

Oracle file features work together with Oracle directories. Each Oracle directory object includes the name of the folder containing the files which need to be processed.

If you want RDS Migration Tool to create and manage the Oracle directories, you need to grant the CREATE ANY DIRECTORY privilege specified above. Note that the directory names will be prefixed with `amazon_`. If you do not grant this privilege, you need to create the corresponding directories manually. In you create the directories manually and the Oracle user specified in the Oracle Source endpoint is not the user that created the Oracle Directories, grant the READ on DIRECTORY privilege as well.

If the Oracle source endpoint is configured to copy the Redo log files to a temporary folder, and the Oracle user specified in the Oracle source endpoint is not the user that created the Oracle directories, the following additional privileges are required:

- READ on the Oracle directory object specified as the source directory
- WRITE on the directory object specified as the destination directory in the copy process

See also: [Using Advanced Properties for an Oracle Source](#).

Oracle Source Data Types

The Oracle endpoint for Amazon RDS Migration Tool supports most Oracle data types. The following table shows the Oracle source data types that are supported when using Amazon RDS Migration Tool and the default mapping to Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 5–3 Supported Oracle Data Types with Mapping to Amazon RDS Migration Tool Data Types

Oracle Data Types	Amazon RDS Migration Tool Data Types
BINARY_FLOAT	REAL4
BINARY_DOUBLE	REAL8
BINARY	BYTES
FLOAT (P)	Precision < or = 24: REAL4 Precision > 24: REAL8
NUMBER (P,S)	When scale is < 0: REAL8
NUMBER according to the "Expose number as" property in the Amazon RDS Migration Tool Oracle source database settings.	When scale is 0 and: Precision = 0: REAL8 Precision < or = 2: INT1 Precision >2 and <or = 4: INT2 Precision >4 and <or = 9: INT4 Precision > 9: NUMERIC If precision > or = scale: NUMERIC In all other cases: REAL8
DATE	DATETIME
INTERVAL_YEAR TO MONTH	STRING (with interval_year_to_month indication)
INTERVAL_DAY TO SECOND	STRING (with interval_day_to_second indication)
TIME	DATETIME
TIMESTAMP	DATETIME
TIMESTAMP WITH TIME ZONE	STRING (with timestamp_with_timezone indication)
TIMESTAMP WITH LOCAL TIME ZONE	STRING (with timestamp_with_local_timezone indication)
CHAR	STRING
VARCHAR2	STRING
NCHAR	WSTRING
NVARCHAR2	WSTRING
RAW	BYTES
REAL	REAL8

Table 5–3 (Cont.) Supported Oracle Data Types with Mapping to Amazon RDS Migration Tool Data Types

Oracle Data Types	Amazon RDS Migration Tool Data Types
<p>BLOB</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of BLOBs for a specific task.</p> <p>BLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	BLOB
<p>CLOB</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of CLOBs for a specific task.</p> <p>During CDC, CLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	CLOB
<p>NCLOB</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of NCLOBs for a specific task.</p> <p>During CDC, NCLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	NCLOB
<p>LONG</p> <p>The LONG data type is not supported in Batch optimized apply (TurboStream CDC) mode.</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of LOBs for a specific task.</p> <p>During CDC, LOB data types are supported only in tables that have a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	CLOB
<p>LONG RAW</p> <p>The LONG RAW data type is not supported in Batch optimized apply (TurboStream CDC) mode.</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of LOBs for a specific task.</p> <p>During CDC, LOB data types are supported only in tables that have a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	BLOB

Table 5–3 (Cont.) Supported Oracle Data Types with Mapping to Amazon RDS Migration Tool Data Types

Oracle Data Types	Amazon RDS Migration Tool Data Types
XMLTYPE	CLOB
<p>Note: Support for the XMLTYPE data type requires the full Oracle Client (as opposed to the Oracle Instance Client).</p> <p>When the target column is a CLOB, both full lob mode and limited lob mode are supported (depending on the target).</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	

Non-Supported Data Types

Source Oracle tables with columns of the following Oracle data types cannot be replicated. Replicated columns with these data types will show as null.

- BFILE
- ROWID
- REF
- UROWID
- Nested Table
- User-defined data types
- ANYDATA

Note: Virtual columns are not supported.

Configuring an Oracle Database as an Amazon RDS Migration Tool Source

The following topics describe the configuration requirements for using an Oracle database with Amazon RDS Migration Tool as a source. An Oracle DBA should know how to carry out these tasks.

- [Provide Oracle Account Access](#)
- [Ensure that ARCHIVELOG Mode is On](#)
- [Setting up Supplemental Logging](#)

Provide Oracle Account Access

You must provide Oracle account access to the Amazon RDS Migration Tool user. This user must have read/write privileges on the Oracle database. For information on setting up access to the Oracle account, see [Security Requirements](#).

Ensure that ARCHIVELOG Mode is On

Oracle can be run in two different modes: the ARCHIVELOG mode and the NOARCHIVELOG mode. To use the Oracle logs with Amazon RDS Migration Tool, run the database in ARCHIVELOG mode. If the log is not set to ARCHIVELOG mode, then execute the following query:

```
ALTER DATABASE ARCHIVELOG
```

Note that if your Oracle database instance is on Amazon RDS, a different command needs to be executed. For more information, see [Enabling Automatic Backups](#) and [Setting Up Archiving](#) in *Working with Oracle on Amazon RDS*.

Setting up Supplemental Logging

You must be sure that supplemental logging is enabled for the Oracle database.

Note: You can automatically set up supplemental logging in the Advanced tab of the Add Oracle Database dialog box. If you select this option, you do not have to carry out the following procedure. For more information, see [Using Advanced Properties for an Oracle Source](#).

Set up supplemental logging as described in the steps below.

Step 1: Check that supplemental logging is enabled for the database

- a. Run the following query:

```
SELECT name, value, description FROM v$parameter WHERE  
name = 'compatible';
```

The returned result should be from GE to 9.0.0.

- b. Run the following query:

```
SELECT supplemental_log_data_min FROM v$database;
```

The returned result should be YES or IMPLICIT.

Enable supplemental logging by executing the following query:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA
```

Note that if your Oracle database instance is on Amazon RDS, a different command needs to be executed. For more information, see [Working with Oracle on Amazon RDS](#).

Step 2: Make sure that the required supplemental logging is added for each table

- a. If a Primary Key exists, supplemental logging must be added for the Primary Key either by using the format to add supplemental logging on the Primary Key, or by adding supplemental logging on the Primary Key columns.
- b. If no Primary Key exists and the table has a single Unique Index, then all of the Unique Index's columns must be added to the supplemental log. Using `SUPPLEMENTAL LOG DATA (UNIQUE INDEX) COLUMNS` does not add the Unique Index columns to the log.
- c. If no Primary Key exists and the table has multiple Unique Indexes, Amazon RDS Migration Tool will select the first Unique Index. Amazon RDS Migration Tool will use the first index in an alphabetically ordered ascending list. Supplemental logging must be added on the selected index's columns. Using `SUPPLEMENTAL LOG DATA (UNIQUE INDEX) COLUMNS` does not add the Unique Index columns to the log.
- d. If there is no Primary Key and no Unique Index, supplemental logging must be added on all columns.

Note: When the target table Primary Key/Unique Index is different than the source table Primary Key/Unique Index, the user needs to add supplemental logging manually on the source table columns that comprise the target table Primary Key/Unique Index.

- e. If you change the target table primary key, the supplemental logging must be added on the selected index's columns instead of the columns of the original primary key/unique index.

Step 3: If a filter or transformation is defined for the table, additional logging might be necessary:

Note: If ALL COLUMNS supplemental logging has been added to the table, there is no need to add any additional logging.

If the table has a Unique Index or a Primary Key, you also need to add supplemental logging on each column that is involved in a filter or transformation (if those columns are different than the Primary Key or Unique Index columns).

Note: If a transformation uses only one column, this column may not be added to a supplemental logging group. For example, "A+B" needs both columns to be added, whereas `substring(A, 10)` does not need "A" to be added.

One method of setting up both Primary Key/Unique Index supplemental logging and supplemental logging on specific columns is to add `USER_LOG_GROUP` supplemental logging only on the Primary Key/Unique Index columns and on the columns that are filtered or transformed.

For example, to replicate a table named `EXAMPLE.TABLE` with Primary Key `ID` and filter by column `NAME`, you can run a command similar to the one below to create the log group supplemental logging:

```
ALTER TABLE EXAMPLE.TABLE ADD SUPPLEMENTAL LOG GROUP example_log_group (ID,NAME) ALWAYS;
```

Working with Oracle on Amazon RDS

Amazon RDS Migration Tool supports Oracle 11.2.0.2.v7 and above on Amazon RDS as a source endpoint and in LogMiner mode only. This section details the requirements for working with Oracle on Amazon RDS.

Setting Up Supplemental Logging

Amazon RDS Migration Tool requires database-level supplemental logging to be enabled. To enable database-level supplemental logging, execute the following command:

```
exec rdsadmin.rdsadmin_util.alter_supplemental_logging('ADD');
```

Additional commands that you can execute to change the supplemental logging attributes include:

```
exec rdsadmin.rdsadmin_util.alter_supplemental_logging('ADD', 'ALL');
```

```
exec rdsadmin.rdsadmin_util.alter_supplemental_logging('DROP','PRIMARY KEY');
```

Enabling Automatic Backups

In **Step 5: Management Options** of setting up your Oracle database instance, set the **Enabled Automatic Backups** option to **Yes**.

Setting Up Archiving

To retain archived redo logs of your Oracle database instance (which will allow Amazon RDS Migration Tool to retrieve the log information using LogMiner), execute the following command (example 24 hours):

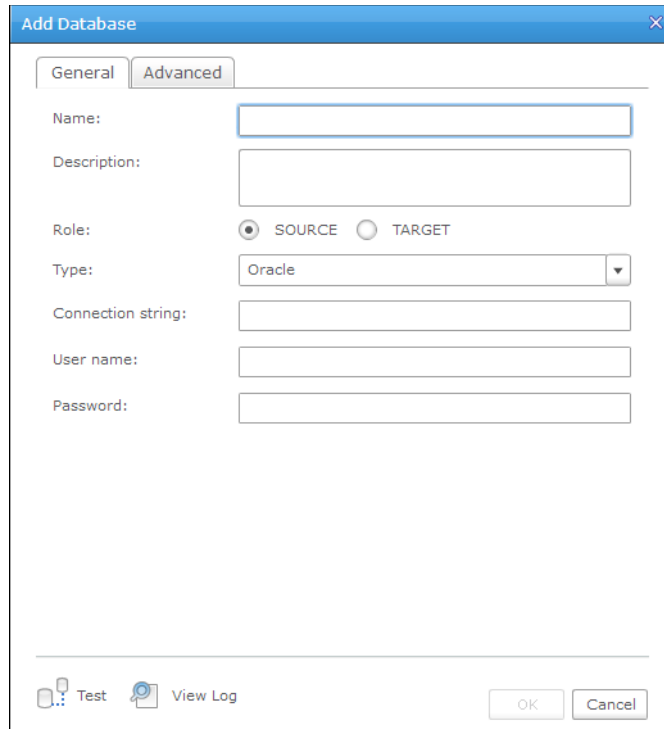
```
exec rdsadmin.rdsadmin_util.set_configuration('archive_log_retention_hours',24);
```

Make sure that your storage has sufficient space for the archived redo logs during the specified period.

Setting up an Oracle Database as a Source in Amazon RDS Migration Tool

You can add an Oracle database to Amazon RDS Migration Tool to use as a source. For information on how to add databases, see [Working with Databases](#). When you select **Oracle Source** as the database type the following dialog box is displayed:

Figure 5–1 Oracle Source Database



Notes:

- Oracle can also be used as a target database. For information on using Oracle as a target, see [Setting up an Oracle Database as a Target in Amazon RDS Migration Tool](#).
 - You can also use Oracle files as a source or target. For more information, see [Using Files](#).
-
-

To add an Oracle source database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Tool console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the Oracle database. This is optional.
4. Select **SOURCE** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **Oracle** as the database **Type**.

6. Type the Oracle **Connection String** for the Oracle Database you want to work with. You can type the connect string in any Oracle format, for example:

```
//host:port/service name
```

Where:

- **host:** This is the name or IP address for the computer with the Oracle database that you are using. For example, `johnboy_w7` or `255.255.255.0`.
- **port:** (optional) This is the TNS Listener Port number for the computer with the Oracle database that you are using. If you do not enter a port number the default Oracle TNS Listener port is used.
- **service name:** (optional) This is the service name for the computer with the Oracle database you are using. If you do not enter a service name the default service name is used.

You can also enter an Oracle Net keyword-value pair. For example:

```
"(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=dlsun242) (PORT=5521))
(CONNECT_DATA=(SERVICE_NAME=bjava21)))"
```

Note: When working with a Multitenant environment, the connection string should specify a specific PDB.

Limitations:

- Connecting to the CDB is not supported.
 - Oracle does not support using PDB with LogMiner. Therefore, if you want to connect to a PDB, make sure that the **Use LogMiner to access redo logs** option is disabled in the **Advanced** tab.
-
-

Specifying Separate Connection Strings for Different RAC Threads

In Oracle Binary Reader, you can use separate connection strings for different RAC threads.

When the redo logs are stored in ASM, the connection string syntax is as follows:

```
[<common ASM connection string>,<thread id> <thread ASM connection string>,<thread id> <thread ASM connection string>...
```

Note: If no `<common ASM connection string>` is specified, all the RAC threads should be defined in the ASM connection.

When using Oracle Binary Reader to access the redo logs, the connection string syntax is as follows:

```
<Oracle connection string>[,<thread id> <thread BFILE connection string>,<thread id> <thread BFILE connection string> ...]
```

`<Oracle connection string>` is mandatory. If specified, the `<thread BFILE connection string>` will be used instead of the `<Oracle connection string>`.

7. Type the Oracle authentication information (User Name, Password) for the authorized user for this Oracle database. If you do not know this information, see your Oracle Database Administrator (DBA).

Note:

- This information is required. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. See [Using Advanced Properties for an Oracle Source](#) for more information.
 - This information is case sensitive.
 - **Important:** Make sure that the Oracle user entered in the Oracle Authentication section has the correct access privileges. For information on how to provide the required privileges, see [Security Requirements](#).
 - If you want to set custom properties for this database, see [Using Advanced Properties for an Oracle Source](#).
-
-

Using Advanced Properties for an Oracle Source

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the **Add Database** dialog box. Use the scroll bar on the right of the tab to scroll through the available options.

You can set the following parameters:

- **Automatically add supplemental logging:** Select this to automatically set up supplemental logging for the Oracle database. For more information on supplemental logging, see [Setting up Supplemental Logging](#).
- **Use LogMiner to access redo logs:** Select this to capture change data using the LogMiner utility (the default). Clear this option if you want Amazon RDS Migration Tool to access the redo logs as a binary file.

IMPORTANT: You must enable the **Use LogMiner to access redo logs** option if your Oracle database version is prior to 10.2.0.4 (i.e. version 10.1.x to 10.2.0.3).

- **Secret Store encryption entries:** When "Use LogMiner to access redo logs" is *not* selected and the Oracle instance uses TDE tablespace encryption, you need to specify the encryption key (or keys if more than one was used) that was/were used to encrypt the Redo logs. In order to specify the correct encryption key(s), you first need to find the relevant entry (or entries in the case of multiple keys) in the Oracle Wallet containing the encryption key(s). For instructions on how to do this, see [Finding the Relevant Entries in the Oracle Wallet](#) below. After you find the relevant entry or entries, copy the entry and its value (or entries and values if more than one) into the **Names** and **Values** fields respectively.

Note: To enter multiple values, first copy each entry into a text editor such as Notepad making sure to separate the values with a comma. Then, copy the string containing the values and commas from the text editor and paste it into the **Values** field. There is no need to do this for entries. You can paste the entries directly into the **Entries** field, remembering to separate each entry with a comma.

Finding the Relevant Entries in the Oracle Wallet

Follow the procedure below to locate the correct entries in the Oracle wallet.

1. If the ENCRYPTION_WALLET_LOCATION parameter is defined in the sqlnet.ora file, use the wallet from the directory defined by this parameter.
2. If the WALLET_LOCATION parameter is defined in the sqlnet.ora file, use the wallet from the directory defined by this parameter.
3. In other cases, use the wallet in the default database location.

Note: The name of the wallet should be **ewallet.p12**.

4. Use the “list” option in the Oracle mkstore utility to determine the ORACLE.SECURITY.DB/TS.ENCRYPTION.<SUFFIX> entry name(s), as follows:

```
mkstore -wrl <full wallet name> -list
```

5. If you know which entry/entries is/are used to encrypt the Redo logs, select the entry name(s) and use the “viewEntry” option in the Oracle mkstore utility to determine the entry value, as follows:

```
mkstore -wrl <full wallet name> -viewEntry <entry name>
```

Note: If you do not know which entry is used to encrypt the Redo logs, you can select multiple DB or TS entries and determine their values as described above (and then copy and paste the entry names and values into the **Names** and **Values** fields as described in the [Note](#) above). If the specified entries are not correct, the task will fail and the error message will contain the correct entry name.

Note: If the DBA changes the entry while the task is running, the task will fail and the error message will contain the new entry name. Add the new entry (name and value) to the already specified entries and then resume the task.

- **ASM Parameters (if redo logs are stored in ASM)** - If the Oracle redo logs you are using are stored using Automated Storage Management (ASM), enter the required access information in the designated fields (not available when the **Use LogMiner to access redo logs** option is selected).

Note: To access the redo logs in ASM, you also need to grant the following read privilege: `SELECT ON v_$transportable_platform`

- **ASM Connection String:** The connection string to the ASM instance if your Oracle database is using ASM.
- **ASM user name:** The user name for the ASM user.
- **ASM password:** The password for the ASM user.

See also [Required ASM Privileges](#).

- **To access a redo log as a binary file (i.e. *not* using Use LogMiner):**

Note: To access the redo log as a binary file, you also need to grant the following read privilege: `SELECT ON v_$transportable_platform`

- **Use path as it appears in the database:** Select this to access the redo logs using the path as it appears in the database.
- **Amazon RDS Migration Tool has file-level access to the redo logs:** Select this to access and read the redo logs directly from the file system of the local computer where Amazon RDS Migration Tool is installed.
- **Copy archived redo logs to a temporary folder:** Select this to copy the archived redo logs to a temporary folder and then specify the path of the archived redo logs on the Oracle machine.
 - **Amazon RDS Migration Tool has file-level access to temporary folder:** Select this to access and read the archived redo logs directly from the file system of the local computer where Amazon RDS Migration Tool is installed.
 - **Access archived redo logs in folder:** To enable Amazon RDS Migration Tool to access the temporary folder (when it has file level access), you need to specify the path to the shared temporary folder on the Oracle machine e.g. `\\my.oracle.box\tempshare`.
- **Look for missing archived redo logs in folder:** Type the full path to a location from where you want Amazon RDS Migration Tool to read the archived redo logs if they are not found in the default location. The folder can be located anywhere in the network where Amazon RDS Migration Tool is located, but be sure that the location is accessible to the Amazon RDS Migration Tool user.
 - **Amazon RDS Migration Tool has file-level access to the specified folder:** Select this to access and read the archived redo logs directly from the file system of the local computer where Amazon RDS Migration Tool is installed.
- **Replace path prefix:** You can determine whether to read the redo logs from a different root location while leaving the relative path unchanged.

Type the first part of the path to the current location of the redo logs. For example, `C:\OldFolder`.

You can include one folder or directory level or multiple folder or directories in this field.

By: Type the name of the folder or prefix to replace the existing prefix that you added in the field above. For example, `C:\NewFolder`.

Note: The following are examples of how to change the prefix:

If your redo logs are located in `C:\OldFolder\archive\logs` and you type `C:\NewFolder` in the **By** field, the redo logs will be read from:

`C:\NewFolder\archive\logs`.

If your redo logs are located in

`C:\replicate\oracle\logs\archive\RedoLogs` and:

In the **Replace path prefix** field, you type `C:\replicate\oracle\logs` then,

In the **By** field, you type `C:\companyName` then,

The redo logs will be read from:

`C:\companyName\archive\RedoLogs`

In this case, the new folder or directory called `companyName` replaces all of the first three level folders that you included in the **Replace path prefix** field.

- **Apply prefix replacement to online and archived redo logs:** Select this to apply the prefix replacement to the online and archived redo logs.
 - **Amazon RDS Migration Tool has file-level access to the new location:** Select this to access and read the online and archived redo log files directly from the file system of the local computer where Amazon RDS Migration Tool is installed.
 - **Apply prefix replacement to archived redo logs only:** Select this to apply the prefix replacement to the archived redo logs only (and not to the online redo logs).
 - **Amazon RDS Migration Tool has file-level access to the original online location:** Select this to access and read the original online redo log files directly from the file system of the local computer where Amazon RDS Migration Tool is installed.
 - **Amazon RDS Migration Tool has file-level access to the new archive location:** Select this to access and read the archived redo log files directly from the file system of the local computer where Amazon RDS Migration Tool is installed.
 - **Delete processed archived redo log files:** Select this to delete the copied archived redo log files after they have been read.
- **Retry interval:** Use the counter or type the number of seconds that the system waits before resending a query.
- **Archived redo logs location identifier:** The destination of the archived redo logs. The value should be the same as the `DEST_ID` number in the `$archived_log` table.

Note : When working with multiple log destinations (`DEST_ID`), it is recommended to specify an **Archived redo logs location identifier**. This will improve performance by ensuring that the correct logs are accessed from the outset.

- **Expose NUMBER as:** Select a precision, scale combination from the drop-down list. Amazon RDS Migration Tool always supports a scale up to 38. You can select a scale up to 38 or you can select `FLOAT`.

By default the `NUMBER` data type is converted to precision 38, scale 10.

Note: The "Expose `NUMBER`" definition in the Oracle database is used for the `NUMBER` data type only (without the explicit precision and scale definition).

- **Use archived redo logs only:** When this option is selected, Amazon RDS Migration Tool will only access the archived redo logs. If the archived redo logs are stored on ASM only, the Amazon RDS Migration Tool user needs to be granted the ASM privileges described in [Required ASM Privileges](#).
- **Override connection string parameters:** Select this to create a custom connect string.

The following is an example of an Oracle connection string:

```
server=SRV1:1521/orcl;username=system;password=manager;
```

Where:

`server` is the name of the computer with the Oracle instance you are working with.

`username` is the name of the user authorized to access the Oracle account you are using.

`password` is the password for the Oracle user.

For information on the properties you can include in a custom connect string for AIS, see [Configuration Properties for the Oracle Source](#).

Required ASM Privileges

From Oracle 11g Release 2 (11.2.0.2), Amazon RDS Migration Tool must be granted the SYSASM privilege in order to access the ASM account. For older supported versions, granting Amazon RDS Migration Tool the SYSDBA privilege should be sufficient.

Note: When connecting to ASM, Amazon RDS Migration Tool will first try to log in as SYSDBA and, if unsuccessful, will try to log in as SYSASM.

You can validate ASM account access by opening a command prompt and issuing the following statements:

```
sqlplus asmuser/asmpassword@+asmserver as sysdba
```

-OR-

```
sqlplus asmuser/asmpassword@+asmserver as sysasm
```

Using an Oracle Database as a Target

The following topics describe what you need to use an Oracle database as the target endpoint in an Amazon RDS Migration Tool task.

- [Configuration Properties for the Oracle Target](#)
- [Security Requirements](#)
- [Oracle Target Data Types](#)
- [Configuring an Oracle Database as an Amazon RDS Migration Tool Target](#)
- [Setting up an Oracle Database as a Target in Amazon RDS Migration Tool](#)

Configuration Properties for the Oracle Target

The following table describes the configuration properties available for an Oracle target. For information on how to set these properties, see [Using Advanced Properties for an Oracle Target](#).

Table 5–4 Oracle Connect String Standard Properties

Property	Description/Format	Default Value
server	The name of the computer with the Oracle instance you are working with. This is written in any recognized Oracle format. server=tomato:1521/orcl	
username	The name of the user authorized to access the Oracle account you are using. username=john	
password	The password for the authorized user. password=kuku	
useDirectPathFullLoad	Y indicates that the OCI direct path protocol for bulk loading Oracle tables is used. N indicates that tables are loaded without the bulk-load API.	Y

Security Requirements

A user must have the following privileges granted in the Oracle database to use an Oracle target in an Amazon RDS Migration Tool task:

- SELECT ANY TRANSACTION
- SELECT on V\$NLS_PARAMETERS
- SELECT on V\$TIMEZONE_NAMES
- SELECT on ALL_INDEXES
- SELECT on ALL_OBJECTS
- SELECT on DBA_OBJECTS
- SELECT on ALL_TABLES
- SELECT on ALL_USERS
- SELECT on ALL_CATALOG
- SELECT on ALL_CONSTRAINTS
- SELECT on ALL_CONS_COLUMNS
- SELECT on ALL_TAB_COLS
- SELECT on ALL_IND_COLUMNS
- DROP ANY TABLE
- SELECT ANY TABLE
- INSERT ANY TABLE
- INSERT ANY TABLE
- UPDATE ANY TABLE
- CREATE ANY VIEW
- DROP ANY VIEW
- CREATE ANY PROCEDURE

- ALTER ANY PROCEDURE
- DROP ANY PROCEDURE
- CREATE ANY SEQUENCE
- ALTER ANY SEQUENCE
- DROP ANY SEQUENCE

You can add the following permissions to use a specific table list:

- SELECT on <any-replicated-table>
- ALTER on <any-replicated-table>

The following permission must be granted for logon:

- CREATE SESSION

The following permission must be granted if you are using a direct path:

- LOCK ANY TABLE

If the "DROP and CREATE table" or "TRUNCATE before loading" option is selected in the [Full Load Settings](#) tab *and* the target table schema is different from the Amazon RDS Migration Tool user, the following permission must be granted:

- DROP ANY TABLE

The Amazon RDS Migration Tool user must also be granted read permissions for the following DBA tables:

- SELECT on DBA_USERS
- SELECT on DBA_TAB_PRIVS
- SELECT on DBA_OBJECTS
- SELECT on DBA_SYNONYMS
- SELECT on DBA_SEQUENCES
- SELECT on DBA_TYPES
- SELECT on DBA_INDEXES
- SELECT on DBA_TABLES
- SELECT on DBA_TRIGGERS

Note: If any of the required privileges cannot be granted to a V\$xxx, then grant them to the V_\$.xxx.

Oracle Target Data Types

The Oracle endpoint for Amazon RDS Migration Tool supports most Oracle data types. The following table shows the Oracle target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped from the source, see the section for the source endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 5–5 Supported Oracle Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	Oracle Data Types
BOOLEAN	NUMBER (1)
BYTES	RAW (length)
DATE	DATETIME
TIME	TIMESTAMP (0)
DATETIME	TIMESTAMP (scale)
INT1	NUMBER (3)
INT2	NUMBER (5)
INT4	NUMBER (10)
INT8	NUMBER (19)
NUMERIC	NUMBER (p,s)
REAL4	FLOAT
REAL8	FLOAT
STRING	With date indication: DATE With time indication: TIMESTAMP With timestamp indication: TIMESTAMP With timestamp_with_timezone indication: TIMESTAMP WITH TIMEZONE With timestamp_with_local_timezone indication: TIMESTAMP WITH LOCAL TIMEZONE With interval_year_to_month indication: INTERVAL YEAR TO MONTH with interval_day_to_second indication: INTERVAL DAY TO SECOND In all other cases: VARCHAR2 (Length)
UINT1	NUMBER (3)
UINT2	NUMBER (5)
UINT4	NUMBER (10)
UINT8	NUMBER (19)
WSTRING	NAVARCHAR2 (length)

Table 5–5 (Cont.) Supported Oracle Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	Oracle Data Types
BLOB	<p>BLOB</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of BLOBs for a specific task.</p> <p>BLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
CLOB	<p>CLOB</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of CLOBs for a specific task.</p> <p>During CDC, CLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
NCLOB	<p>NCLOB</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of NCLOBs for a specific task.</p> <p>During CDC, NCLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
The XMLTYPE target data type is only relevant in Oracle-to-Oracle replication tasks. See the note below.	XMLTYPE

Note: When the source database is Oracle, the source data types will be replicated "as is" to the Oracle target. For example, an XMLTYPE data type on the source will be created as an XMLTYPE data type on the target.

Configuring an Oracle Database as an Amazon RDS Migration Tool Target

The following topics describe the configuration requirements for using an Oracle database with Amazon RDS Migration Tool as a target. An Oracle DBA should know how to carry out these tasks.

- [Provide Oracle Account Access](#)

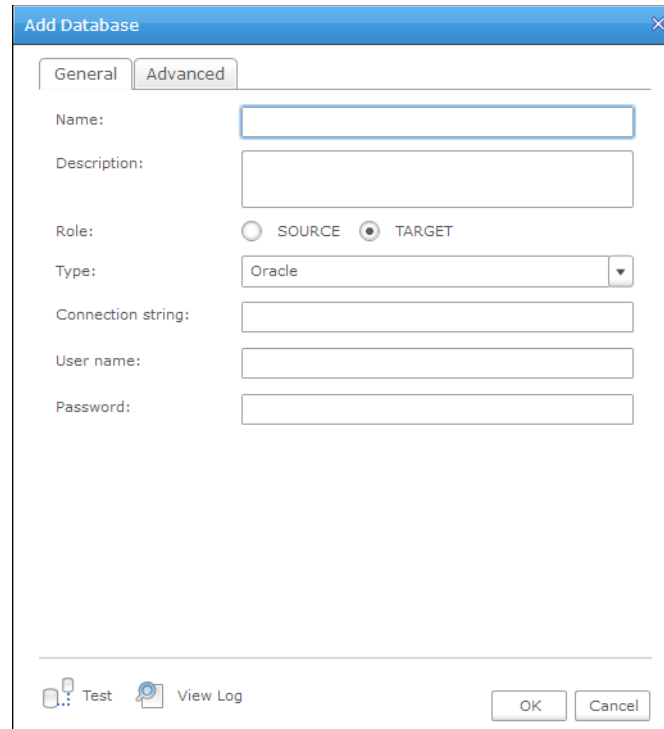
Provide Oracle Account Access

You must provide Oracle account access to the Amazon RDS Migration Tool user. This user must have read/write privileges on the Oracle database. For information on setting up access to the Oracle account, see [Security Requirements](#).

Setting up an Oracle Database as a Target in Amazon RDS Migration Tool

You can add an Oracle database to Amazon RDS Migration Tool to use as a target. For information on how to add databases, see [Working with Databases](#). When you select **Oracle Target** as the database type the following dialog box is displayed:

Figure 5–2 Oracle Target Database



Notes:

- Oracle can also be used as a source database. For information on using Oracle as a source, see [Setting up an Oracle Database as a Source in Amazon RDS Migration Tool](#).
 - You can also use Oracle files as a source or target. For more information, see [Using Files](#).
-
-

To add an Oracle Target database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Tool console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the Oracle database. This is optional.
4. Select **TARGET** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **Oracle** as the database **Type**.

6. Type the Oracle **Connection String** for the Oracle Database you want to work with. You can type the connect string in any Oracle format, for example:

```
//host:port/service name
```

Where:

- **host:** This is the name or IP address for the computer with the Oracle database that you are using. For example, `johnboy_w7` or `255.255.255.0`.
- **port:** (optional) This is the TNS Listener Port number for the computer with the Oracle database that you are using. If you do not enter a port number the default Oracle TNS Listener port is used.
- **service name:** (optional) This is the service name for the computer with the Oracle database you are using. If you do not enter a service name the default service name is used.

You can also enter an Oracle Net keyword-value pair. For example:

```
"(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=dlsun242) (PORT=5521))
(CONNECT_DATA=(SERVICE_NAME=bjava21)))"
```

Note:

- This information is case sensitive.
- You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for an Oracle Source](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

7. Type the Oracle authentication information (**User Name, Password**) for the authorized user for this Oracle database. If you do not know this information, see your Oracle Database Administrator (DBA).

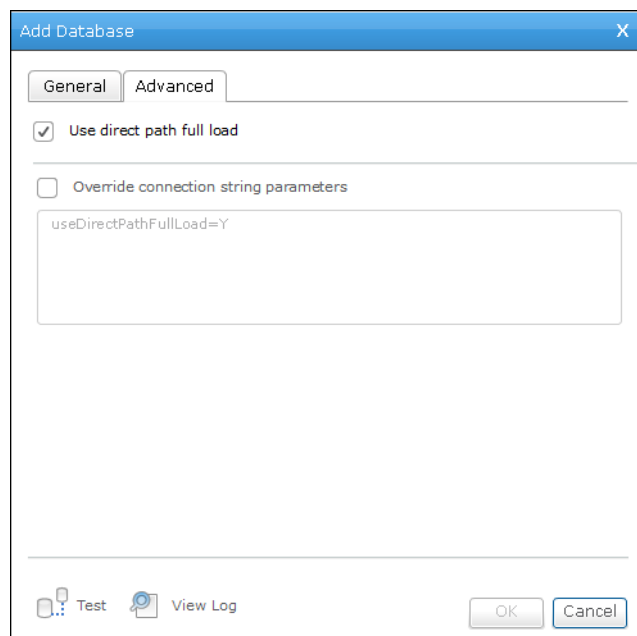
Note:

- This information is required. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. See [Using Advanced Properties for an Oracle Source](#) for more information.
 - This information is case sensitive.
 - **Important:** Make sure that the Oracle user entered in the Oracle Authentication section has the correct access privileges. For information on how to provide the required privileges, see [Security Requirements](#).
 - If you want to set custom properties for this database, see [Using Advanced Properties for an Oracle Source](#).
-
-

Using Advanced Properties for an Oracle Target

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

Figure 5–3 Oracle Target Advanced Tab



You can set the following parameters:

- **Use direct path full load:** Select this to use the OCI direct path protocol for bulk loading Oracle tables. This is the default selection.
- **Override connection string parameters:** Select this to create a custom connect string.

The following is an example of an Oracle connection string:

```
server=SRV1:1521/orcl;username=system;password=manager;
```

Where:

`server` is the name of the computer with the Oracle instance you are working with.

`username` is the name of the user authorized to access the Oracle account you are using.

`password` is the password for the Oracle user.

For information on the properties you can include in a custom connect string for AIS, see [Configuration Properties for the Oracle Source](#).

6



Using a Microsoft SQL Server Database as a Source or Target

This chapter describes how to set up and use a Microsoft SQL Server database as the source or target endpoint in a replication task. It has the following topics:

- [Prerequisites](#)
- [Limitations](#)
- [Using a Microsoft SQL Server Database as a Source](#)
- [Using a Microsoft SQL Server Database as a Target](#)

Prerequisites

Amazon RDS Migration Tool supports the following Microsoft SQL Server editions.

- Enterprise Edition
- Standard Edition
- Workgroup Edition
- Developer Edition

Note: The Microsoft SQL Server database is supported on the following Windows platforms:

- Windows Server 2003, 2008 and 2012
 - Windows XP
 - Windows Vista
 - Windows 7
 - The Microsoft SQL Server database can be installed on any computer in your network.
 - A Microsoft SQL Server account with the specific access privileges is required. See [Security Considerations for a Microsoft SQL Server Source](#) or [Security Considerations for a Microsoft SQL Server Target](#) for more information.
 - Microsoft SQL Server as a source must be configured for a full backup to work with Amazon RDS Migration Tool. For more information, see [Preparing Microsoft SQL Server Backup and Recovery](#).
-

Amazon RDS Migration Tool for Windows: Client Prerequisites (for source and target endpoints):

- For all versions of Microsoft SQL Server, Microsoft SQL Server Native Client 11.0 must be installed on the Amazon RDS Migration Tool machine.

Amazon RDS Migration Tool for Linux: Client Prerequisites (for source and target endpoints):

- For all versions of Microsoft SQL Server, the Simba Microsoft SQL Server ODBC Driver for Linux must be installed on the Amazon RDS Migration Tool machine.

After installing the driver, edit the `simba.sqlserverodbc.ini` file as follows:

- `DriverManagerEncoding=UTF-16`
- `ODBCInstLib=libodbcinst.so`

In addition, you also need to edit the `odbcinst.ini` file as described below.

Note: The Shared Object (.so) file name that you need to specify will either be `libsimbasqlserverodbc_64.so` or `libsimbasqlserverodbc64.so`. Therefore, before editing the file, you should access the Simba installation folder and verify the correct file name.

Edit the `odbcinst.ini` file as follows:

```
[Simba SQL Server ODBC Driver]
Driver =<Simba installation path>/libsimbasqlserverodbc_64.so
```

Example:

```
[Simba SQL Server ODBC Driver]
Driver =/opt/simba/sqlserverodbc/lib/64/libsimbasqlserverodbc_64.so
```

Limitations

- If you are using a Microsoft SQL Server source database in a replication task, the Microsoft SQL Server Replication Publisher definitions for the database that was used in the task are not removed when you remove a task. A Microsoft SQL Server system administrator must delete these definitions from Microsoft SQL Server.
- Tables with more than 8000 bytes of information in a Microsoft SQL Server source including header and mapping information, are not processed correctly due to limitations in buffer size.
- The Microsoft SQL Server source endpoint does not support the use of sparse tables.
- Changes to computed fields in a Microsoft SQL Server source will not be replicated.
- Microsoft SQL Server partition switching is not supported.
- When using the `WRITETEXT` and `UPDATETEXT` utilities, Amazon RDS Migration Tool does not capture events applied on the source database.
- The following DML pattern is not supported:

```
select <*> into <new_table> from <existing_table>
```

- When using Microsoft SQL Server as a source, column-level encryption is not supported.
- Due to a known issue with Microsoft SQL Server 2008, Amazon RDS Migration Tool does not support server level audits on Microsoft SQL Server 2008 as a source endpoint.

For example, running the following command:

```
USE [master]
GO
ALTER SERVER AUDIT [my_audit_test-20140710] WITH (STATE=on)
GO
```

Will cause Amazon RDS Migration Tool to fail.

- The following limitations apply when accessing the backup transaction logs at file level:
 - The backup transaction logs must reside in a shared folder with the appropriate permissions and access rights.
 - Active transaction logs are accessed through the Microsoft SQL Server API (and not at file-level).
 - The Amazon RDS Migration Tool and Microsoft SQL Server machine must reside in the same domain.
 - Compressed backup transaction logs are not supported.
 - Transparent Data Encryption (TDE) is not supported. Note that when accessing the backup transaction logs using SQL Server's native functionality (i.e. not using file-level access), TDE encryption *is* supported.
 - Unix platforms are not supported.
 - Reading the backup logs from multiple stripes is not supported.

For more information on configuring Amazon RDS Migration Tool to access the backup transaction logs at file-level access, see [Using Advanced Properties for a Microsoft SQL Server Source Database](#).

- Microsoft SQL Server backup to multiple disks is not supported.
- Microsoft SQL Server 2012 Target: When the target table is created manually with a computed column, Full Load replication is not supported in BCP mode. Disabling the "Use BCP for loading tables" option in the **Advanced** tab will resolve this issue. For more information on BCP mode, see [Using Advanced Properties for a Microsoft SQL Server Target Database](#).
- When inserting a value into SQL Server spatial data types (GEOGRAPHY and GEOMETRY), one can either ignore the SRID (Spatial Reference System Identifier) property - in which case the default SRID will be used (0 for GEOMETRY and 4326 for GEOGRAPHY) - or specify a different number. When replicating tables with spatial data types, Amazon RDS Migration Tool replaces the SRID that was inserted by user with the default SRID.
- Change Data Capture operations are not supported on memory optimized tables. This limitation applies to Microsoft SQL Server 2014 (when the feature was first introduced) and above.

Using a Microsoft SQL Server Database as a Source

The following topics describe what you need to use a Microsoft SQL Server source in an Amazon RDS Migration Tool task.

- [Supported Compression Methods](#)
- [Configuration Properties for a Microsoft SQL Server Source](#)
- [Security Considerations for a Microsoft SQL Server Source](#)
- [Working with Microsoft SQL Server AlwaysOn Availability Groups](#)
- [Microsoft SQL Server Source Data Types](#)
- [Non-Supported Data Types](#)
- [Configuring a Microsoft SQL Server Database as an Amazon RDS Migration Tool Source](#)
- [Setting up a Microsoft SQL Server Database as an Amazon RDS Migration Tool Source](#)

Supported Compression Methods

The table below lists which compression methods Amazon RDS Migration Tool supports for each Microsoft SQL Server version.

Table 6–1 Supported Microsoft SQL Server Compression Methods

Microsoft SQL Server Version	Row/Page Compression (at Partition Level)	Vardecimal Storage Format	Vardecimal Storage Format Sparse Columns	Vardecimal Storage Format Sparse Columns Columnar Structure Compression
2005	No	No	No	No
2008	Yes	No	No	No
2012	Yes	No	No	No

Configuration Properties for a Microsoft SQL Server Source

The following table describes the configuration properties available for Microsoft SQL Server source endpoints. For information on how to set these properties, see [Using Advanced Properties for a Microsoft SQL Server Source Database](#).

Table 6–2 Microsoft SQL Server Connect String Standard Properties

Property	Description/Format	Value
database	The name of the Microsoft SQL Server database you are using. database=DB1	
WindowsAuthentication	Indicates whether you are using user credentials from the Windows domain. The possible values are: Y indicates that this property is true and Windows authentication is used to access the Microsoft SQL Server database. N indicates that this property is false and Microsoft SQL Server authentication is used to access the Microsoft SQL Server database. WindowsAuthentication=Y	Y (indicates true)
username	The Microsoft SQL Server account user name for the Microsoft SQL Server user. This is required if WindowsAuthentication=N. username=DDD	
password	The password for the Microsoft SQL Server user. This is required if WindowsAuthentiation=N. password=PWd	

Table 6–2 (Cont.) Microsoft SQL Server Connect String Standard Properties

Property	Description/Format	Value
safeguardPolicy	<p>The policy Amazon RDS Migration Tool uses to prevent truncation of the active TLOG. The values can either be:</p> <p>RELY_ON_SQL_SERVER_REPLICATION_AGENT (The default)</p> <p>-OR-</p> <p>EXCLUSIVE_AUTOMATIC_TRUNCATION</p> <p>For more information on these values, see Using Advanced Properties for a Microsoft SQL Server Source Database.</p>	See the description.
safeguardFrequency	<p>The frequency that the safeguardPolicy will be applied.</p> <p>For more information on this property, see Using Advanced Properties for a Microsoft SQL Server Source Database.</p>	300
additionalConnectionProperties	Any additional ODBC connection parameters that you want to specify.	

Security Considerations for a Microsoft SQL Server Source

The following describes the security requirements for using Amazon RDS Migration Tool with a Microsoft SQL Server source or target.

The Amazon RDS Migration Tool user must have the `sysAdmin` fixed server role on the Microsoft SQL Server database you are connecting to.

A Microsoft SQL Server system administrator must provide this permission for all Amazon RDS Migration Tool users.

Working with Microsoft SQL Server AlwaysOn Availability Groups

The Microsoft SQL Server AlwaysOn Availability Groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring.

To use AlwaysOn Availability Groups as a source in Amazon RDS Migration Tool:

1. Enable Distribution on all Microsoft SQL Server instances in your Availability Replicas.
2. In RDS Migration Tool:
 - a. Open the Microsoft SQL Server source database settings.
 - b. In the **Server Name** field, specify the DSN name or IP address that was configured for the Availability Group Listener.

Notes:

- When you start an Amazon RDS Migration Tool task for the first time, it may take longer than usual to start as the creation of the table articles is being duplicated by the Availability Groups Server.

- After you failover, you will see recoverable errors in the log and the task will restart. This is normal as the connection is being forcibly closed by the failover. As soon as the new primary server is available, processing will resume.

Microsoft SQL Server Source Data Types

The Microsoft SQL Server source for Amazon RDS Migration Tool supports most Microsoft SQL Server data types. The following table shows the Microsoft SQL Server source data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 6–3 Microsoft SQL Server Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

Microsoft SQL Server Data Types	Amazon RDS Migration Tool Data Types
BIGINT	INT8
BIT	BOOLEAN
DECIMAL	NUMERIC
INT	INT4
MONEY	NUMERIC
NUMERIC (p,s)	NUMERIC
SMALLINT	INT2
SMALLMONEY	NUMERIC
TINYINT	UINT1
REAL	REAL4
FLOAT	REAL8
DATETIME	DATETIME
DATETIME2 (Microsoft SQL Server 2008 and later)	DATETIME
SMALLDATETIME	DATETIME
DATE	DATE
TIME	TIME
DATETIMEOFFSET	WSTRING
CHAR	STRING
VARCHAR	STRING

Table 6–3 (Cont.) Microsoft SQL Server Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

Microsoft SQL Server Data Types	Amazon RDS Migration Tool Data Types
VARCHAR (max) TEXT To use this data type with Amazon RDS Migration Tool, you must enable the use of CLOBs for a specific task. LOB columns for Microsoft SQL Server tables are updated in the target even for UPDATE statements that did not change the value of the LOB column in Microsoft SQL Server. During CDC, CLOB data types are supported only in tables that include a primary key. For more information, see LOB support in Task Settings/Metadata .	CLOB
NCHAR NVARCHAR (length) NVARCHAR (max) NTEXT To use this data type with Amazon RDS Migration Tool, you must enable the use of NCLOBs for a specific task. LOB columns for Microsoft SQL Server tables are updated in the target even for UPDATE statements that did not change the value of the LOB column in Microsoft SQL Server. During CDC, NCLOB data types are supported only in tables that include a primary key. For more information, see LOB support in Task Settings/Metadata .	WSTRING WSTRING NCLOB
BINARY VARBINARY VARBINARY (max) IMAGE LOB columns for Microsoft SQL Server tables are updated in the target even for UPDATE statements that did not change the value of the LOB column in Microsoft SQL Server. To use this data type with Amazon RDS Migration Tool, you must enable the use of BLOBs for a specific task. BLOB data types are supported only in tables that include a primary key. For more information, see LOB support in Task Settings/Metadata .	BYTES BYTES BLOB
TIMESTAMP UNIQUEIDENTIFIER	BYTES STRING

Table 6–3 (Cont.) Microsoft SQL Server Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

Microsoft SQL Server Data Types	Amazon RDS Migration Tool Data Types
HIERARCHYID	HIERARCHYID - When replicating to Microsoft SQL Server. WSTRING (250) - When replicating to all other endpoints.
XML	NCLOB
LOB columns for Microsoft SQL Server tables are updated in the target even for UPDATE statements that did not change the value of the LOB column in Microsoft SQL Server. To use this data type with Amazon RDS Migration Tool, you must enable the use of NCLOBs for a specific task. During CDC, NCLOB data types are supported only in tables that include a primary key. For more information, see LOB support in Task Settings/Metadata .	
GEOMETRY	GEOMETRY - When replicating to endpoints that support this data type. CLOB - When replicating to endpoints that do not support this data type.
GEOGRAPHY	GEOGRAPHY - When replicating to endpoints that support this data type. CLOB - When replicating to endpoints that do not support this data type.

Non-Supported Data Types

Tables that include fields with the following data types are not supported by Amazon RDS Migration Tool.

- CURSOR
- SQL_VARIANT
- TABLE

Note: User-defined data types are supported according to their base-type. For example a user-defined data type based on DATETIME is handled as a DATETIME data type.

Configuring a Microsoft SQL Server Database as an Amazon RDS Migration Tool Source

This topic describes the configuration requirements for using a Microsoft SQL Server database. A Microsoft SQL Server system administrator should carry out these tasks.

- [Preparing Microsoft SQL Server Backup and Recovery](#)
- [Setting up Microsoft SQL Server for Replication](#)
- [Replicating Tables that do not have a Primary Key](#)

- [Defining Microsoft SQL Server Database Settings](#)

Preparing Microsoft SQL Server Backup and Recovery

Amazon RDS Migration Tool consumes changes captured from the database transaction LOG (TLOG). The TLOG is maintained by Microsoft SQL Server for recovery purposes. All changes made to a database are written to the TLOG. The following happens when recovery is required:

- A backup copy of the database is made.
- Logged events are taken and used in a rolling-forward process where the recorded changes are replayed against that copy.

To prepare for backup and recovery you must make sure that the Microsoft SQL Server Recovery Model is set up. You select the Recovery Model in the Microsoft SQL Server Management Studio. This should be carried out by a Microsoft SQL Server system administrator.

The TLOG data is truncated as soon as it is no longer needed therefore the TLOG is not persistent. However, Amazon RDS Migration Tool guaranteed delivery requires persistency in the changed data. To ensure persistency:

- A full database backup must be carried out before beginning to replicate data.
- The Recovery Model must be set to **Bulk logged** or **Full**.

To set the recovery model

In the database properties **Options** tab, set the Recovery Model to **Bulk logged** or **Full**. In these modes, the transaction Log is more durable.

Setting up Microsoft SQL Server for Replication

If you are using Microsoft SQL Server as the source in an Amazon RDS Migration Tool task, you need to enable your Microsoft SQL Server database for MS-REPLICATION.

In the Microsoft SQL Server's Management Studio, follow the instructions provided by the Configure Distribution wizard to set up replication or see the Microsoft SQL Server documentation.

To open the wizard from Microsoft SQL Server

- In the Microsoft SQL Server Management Studio, right-click the Replication folder and select **Configure Distribution**.

The Configure Distribution wizard opens.

You should make the following selections in the wizard:

- In the Distributor step, select **<Microsoft SQL Server Name> will act as its own distributor; Microsoft SQL Server will create a distribution database and log**

Replicating Tables that do not have a Primary Key

Note: This functionality is supported only for Microsoft SQL Server Enterprise edition and is *not* supported on Microsoft SQL Server 2005.

By default, Amazon RDS Migration Tool automatically sets up MS-REPLICATION for each of the source tables in a replication task. However, MS-REPLICATION requires each of the source tables to have a primary key, which may not always be the case. Therefore, if you need to replicate tables that do not have a primary key, the following options are available:

- [Use MS-CDC](#)
- [Do not Use MS-Replication or MS-CDC](#)

Use MS-CDC

To set up MS-CDC, you first need to enable MS-CDC for the database by running the following command:

```
use [DBname]
EXEC sys.sp_cdc_enable_db
```

Then you need to enable MS-CDC for each of the source tables by running the following command:

```
EXECUTE sys.sp_cdc_enable_table @source_schema = N'MySchema', @source_name = N'MyTable', @role_name = NULL;
```

Note: Replicating tables that do not have a Primary Key or a Unique Index may adversely affect performance (since additional database resources are required to capture the changes). However, you can prevent performance issues related to the absence of Primary Keys or a Unique Index by manually adding indexes to the target tables.

For more information on setting up MS-CDC for specific tables, please refer to the Microsoft website.

Do not Use MS-Replication or MS-CDC

If your database is not set up for MS-REPLICATION or MS-CDC, you can still capture tables that do not have a Primary Key, but bear in mind that in such a setup only INSERT/DELETE DML events will be captured. UPDATE and TRUNCATE TABLE events will be ignored.

It is also important to note that a DELETE statement executed on an UPDATED source record, will *not* be applied on the target.

Defining Microsoft SQL Server Database Settings

Set the following for the Microsoft SQL Server database(s) that you are using as a source:

- From the Object Explorer in the Microsoft SQL Server Management Studio, right click the database and select **Properties**. In the **Options** tab, set the **Recovery model** to **Bulk logged** or **Full**. In this mode, the transaction Log is more durable and truncation occurs less frequently.
- Ensure that there is a full database backup for each Microsoft SQL Server database that you are using as a source.
- When creating a connection string, it is possible to use any parameter supported by Microsoft SQL Server. The Microsoft SQL Server system administrator must ensure that the Microsoft SQL Server instance is configured correctly so that the proper authentication credentials are accepted.

- To be able to work with MS-REPLICATION, each of the source tables must have a primary key.

Working with Windows Authentication

You can configure the Amazon RDS Migration Tool Microsoft SQL Server endpoint to log in to Microsoft SQL Server using Windows authentication.

If you choose this option, you also need to make sure that:

- The Microsoft SQL Server instance is set up to allow Windows log on.
- The Amazon RDS Migration Tool user is specified as the "Log on as" user for the "Amazon RDS Migration Server" service account.

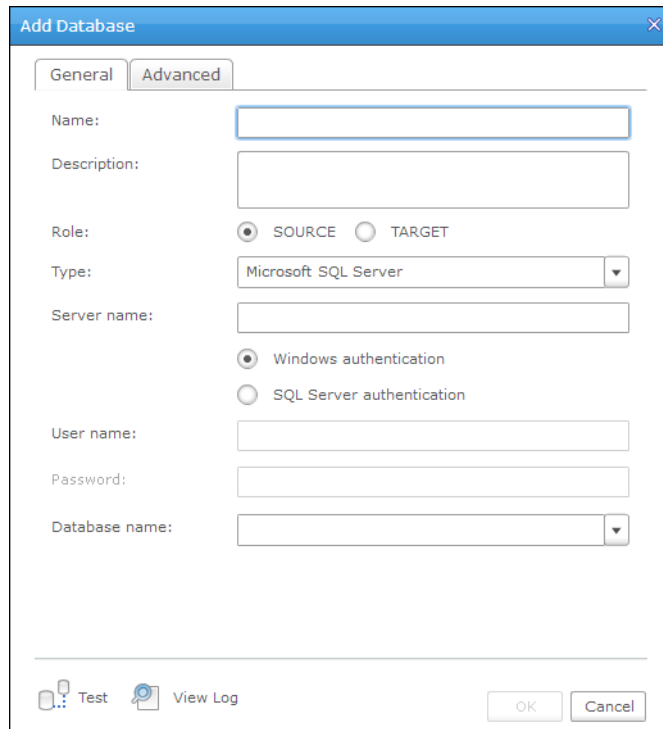
-OR-

Microsoft SQL Server is configured to allow login for the Amazon RDS Migration Server service account.

Setting up a Microsoft SQL Server Database as an Amazon RDS Migration Tool Source

You can add a Microsoft SQL Server database to Amazon RDS Migration Tool to use as a source. For information on how to add databases, see [Working with Databases](#). When you select **Microsoft SQL Server Source** as the database type the following dialog box is displayed:

Figure 6–1 Microsoft SQL Server Source Database



To add a Microsoft SQL Server source database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).

2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the Microsoft SQL Server database. This is optional.
4. Select **SOURCE** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **Microsoft SQL Server** as the database **Type**.
6. Type the **Microsoft SQL Server name**. This is the name of the computer with the Microsoft SQL Server instance you want to work with.

Note:

You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a Microsoft SQL Server Source Database](#).

7. Select **Windows authentication** or **User name/password**.

If you select **Windows authentication**, you will work with the user credentials for the Windows domain. This privilege must be configured in the Microsoft SQL Server database by the system administrator.

Note: When using Windows authentication be sure that the user account that is associated with the **Amazon RDS Migration Server** service has Network read and write permissions. This must be configured by a Windows system administrator.

See also [Working with Windows Authentication](#).

If you select **User name/password**, type the Microsoft SQL Server authentication information (User Name, Password) for the authorized user for this Microsoft SQL Server database. If you do not know this information, see the Microsoft SQL Server System Administrator.

Note:

- The User Name/Password information is mandatory if you select **User name/password**. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. For more information, see [Using Advanced Properties for a Microsoft SQL Server Source Database](#).
- This information is case sensitive.
- **Important:** Make sure that the Microsoft SQL Server user has the correct access privileges. For information on how to provide the required privileges, see [Security Considerations for a Microsoft SQL Server Source](#).
- You can set custom properties in the **Advanced** tab. For more information, see [Using Advanced Properties for a Microsoft SQL Server Source Database](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

8. Type the **Database name** or select one from the list of available databases. This is the name of the database where you are replicating the data.

Using Advanced Properties for a Microsoft SQL Server Source Database

You can use advanced properties and change the default settings for various parameters in a Microsoft SQL Server source or target by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

You can set the following properties in the **Advanced** tab:

- **Prevent truncation of unread changes from TLOG:** For optimal performance, Amazon RDS Migration Tool will try to capture all unread changes from the active transaction log (TLOG). However, sometimes due to truncation, the active TLOG may not contain all of the unread changes. When this occurs, Amazon RDS Migration Tool accesses the backup log to capture the missing changes. To minimize the need to access the backup log, Amazon RDS Migration Tool prevents truncation using one of the following methods:
 - **Start transactions in the database:** This is the default method. When this method is used, Amazon RDS Migration Tool prevents TLOG truncation by mimicking a transaction in the database. As long as such a transaction is open, changes that appear after the transaction started will not be truncated. If you need Microsoft Replication to be enabled in your database, then you must choose this method.
 - **Exclusively use `sp_repldone` within a single task:** When this method is used, Amazon RDS Migration Tool reads the changes and then uses `sp_repldone` to mark the TLOG transactions as ready for truncation. Although this method does not involve any transactional activities, it can only be used when

Microsoft Replication is not running. Also, using this method, only one Amazon RDS Migration Tool task can access the database at any given time. Therefore, if you need to run parallel Amazon RDS Migration Tool tasks against the same database, use the default method.

Note: This method requires the Log Reader Agent to be stopped in the database. If the Log Reader Agent agent is running when the task starts, Amazon RDS Migration Tool will forcibly stop it. Alternatively, you can stop the Log Reader Agent manually, before starting the Amazon RDS Migration Tool task. For instructions on how to do this, refer to the Microsoft SQL Server Management Studio help.

- **Apply TLOG truncation prevention policy every (seconds):** Specify how often to prevent TLOG truncation using one of the methods describes above. Factors that you should consider when determining the policy frequency include storage availability, backup and log routines, and the rate at which Amazon RDS Migration Tool processes events.
- **Alternate backup folder:** The location of the backup logs when using a third-party utility to back up the transaction logs (i.e. instead of Microsoft SQL Server’s own backup mechanism). You can run the backup utility yourself or you can configure Amazon RDS Migration Tool to run it as described in [Backup file preprocessing command](#) below.

Note that the backup files must be exported to the specified location in standard Microsoft SQL Server format.

- **Change processing mode:** Choose one of the following change processing modes:
 - **Read changes from online log** - This is the default. Amazon RDS Migration Tool first reads the online logs for changes and then reads the backup logs.
 - **Read changes from backup if the same record exists in online and backup logs** - When this option is enabled and the same record appears in both the active log and the backup logs, Amazon RDS Migration Tool will read the changes from the backup transaction logs instead of the online transaction log. This can improve performance when reading from the online transaction log is slow (e.g due to lock contention) or when using file-level access to access the backup transaction logs.
 - **Read changes from backup only** - When this option is selected, Amazon RDS Migration Tool will read the changes from the backup transaction logs only. Selecting this method results in increased latency due to the interval between backups. The actual latency time will remain constant, but will vary according to the backup schedule.
- **Amazon RDS Migration Tool has file-level access to the backup log files:** Select this option if Amazon RDS Migration Tool has been granted file-level access to the backup log files in the **Alternate backup folder**.

Note: When Amazon RDS Migration Tool has file-level access to the backup transaction logs, the following rules apply:

- The **Alternate backup folder** must be a common shared network folder, for example: `\\temp\backup`.
- The **Amazon RDS Migration Server** service must be configured to log on using the user name and password specified in the **Backup folder user name** and **Backup folder password** fields.

To do this:

1. In the Windows **Services** console, double-click the **Amazon RDS Migration Server** service.
 2. In the **Log On** tab, select **This account** and then enter the user name and password.
- The specified user must be granted Read permission to the alternate backup folder (i.e. the shared network folder).

For a complete list of the limitations affecting file-level access, see [Limitations](#).

- **Backup folder user name:** The user name required to access the backup folder when Amazon RDS Migration Tool has file-level access.
- **Backup folder password:** The password required to access the backup folder when Amazon RDS Migration Tool has file-level access.
- **Backup file preprocessing command:** You can use a third-party utility to convert the transaction logs to standard Microsoft SQL Server format (if they are in a different format) and back them up to an alternate backup folder. This option should be used in conjunction with the [Alternate backup folder](#) option described above.

Prerequisites and Notes:

The command is invoked via the `XP_CMDSHELL` extended procedure.

- The backup utility is responsible for setting the system return code (0 for success, 1 for failure), assuming that this code is delegated as the `XP_CMDSHELL` return value.
- The backup utility invoked by `XP_CMDSHELL` must have the same security rights as the Microsoft SQL Server service account.
- `XP_CMDSHELL` is normally disabled. It can be enabled and disabled by using the Policy-Based Management or by executing `SP_CONFIGURE`.
- Using this extended procedure requires `CONTROL SERVER` permission (at least).

Command Usage:

The backup utility should provide Amazon RDS Migration Tool with the following parameters:

- `{BACKUP_INFILE}` - The full path to the original backed up transaction log.
- `{ALTDIR_OUTFILE}` - The specifications of the target file to transfer to the alternate backup folder.
- `{BACKUP_SET}` - The backup set to be processed within the backup log.

Example command:

```
C:\Temp\YourBackupUtility.exe -B{BACKUP_INFILE} -A{ALTDIR_OUTFILE}"
```

IMPORTANT: Directory names in the command path or file names in the actual command that contain spaces must be enclosed in double-quotes:

Example:

```
C:\temp\test\"my program\" \"new version\" \converter.exe -A{input file} -B{outfile}
```

- **Delete processed backup logs:** Select this option to delete the backup logs after they have been read.
- **Select virtual backup device types:** When this option is selected, Amazon RDS Migration Tool will read changes from the specified virtual device(s). Usually, this option only needs to be enabled when using a third-party backup utility (which will be recorded as a virtual device).
- **Override connection string parameters:** Select this to use a custom connect string.

The following is an example of a Microsoft SQL Server connection string:

```
server=SRV1;database=DB1;WindowsAuthentication=Y
```

For information on the properties you can include in an custom connect string for Microsoft SQL Server, see [Configuration Properties for a Microsoft SQL Server Source](#).

Using a Microsoft SQL Server Database as a Target

The following topics describe what you need to use a Microsoft SQL Server target in an Amazon RDS Migration Tool task.

- [Configuration Properties for a Microsoft SQL Server Target](#)
- [Security Considerations for a Microsoft SQL Server Target](#)
- [Microsoft SQL Server Target Data Types](#)
- [Setting up a Microsoft SQL Server Database as an Amazon RDS Migration Tool Target](#)

Configuration Properties for a Microsoft SQL Server Target

The following table describes the configuration properties available for Microsoft SQL Server target endpoints. For information on how to set these properties, see [Using Advanced Properties for a Microsoft SQL Server Target Database](#).

Table 6–4 Microsoft SQL Server Connect String Standard Properties

Property	Description/Format	Default Value
server	The name or IP address of the computer on which the Microsoft SQL Server database is installed.	

Table 6–4 (Cont.) Microsoft SQL Server Connect String Standard Properties

Property	Description/Format	Default Value
WindowsAuthentication	<p>Indicates whether you are using user credentials from the Windows domain.</p> <p>The possible values are:</p> <p>Y indicates that this property is true and Windows authentication is used to access the Microsoft SQL Server database.</p> <p>N indicates that this property is false and Microsoft SQL Server authentication is used to access the Microsoft SQL Server database.</p> <p>WindowsAuthentication=Y</p>	Y (indicates true)
username	<p>The Microsoft SQL Server account user name for the Microsoft SQL Server user. This is required if WindowsAuthentication=N.</p> <p>username=DDD</p>	
password	<p>The password for the Microsoft SQL Server user. This is required if WindowsAuthentication=N.</p> <p>password=pwd</p>	
database	<p>The name of the Microsoft SQL Server database you are using.</p> <p>database=DB1</p>	
useBCPFullLoad	<p>Indicates whether full load operations are done using Microsoft SQL Server bulk copy operation.</p> <p>The possible values are:</p> <p>Y indicates that this property is true and full load operations are done using the bulk copy operation.</p> <p>N indicates that this property is false and full load operations are loaded as arrays.</p>	Y (indicates true)
BCPPacketSize	<p>The maximum packet size allowed for Microsoft SQL Server bulk copy operations.</p>	16384
controlTablesFileGroup	<p>You can specify a filegroup for the Amazon RDS Migration Tool internal tables. When the replication task starts, all the internal Amazon RDS Migration Tool control tables (amazon_apply_exception, amazon_apply, amazon_changes) will be created on the specified filegroup.</p> <p>For an example command, see the Filegroup for Amazon RDS Migration Tool internal tables option described below.</p>	
additionalConnectionProperties	<p>Any additional ODBC connection parameters that you want to specify.</p>	

Security Considerations for a Microsoft SQL Server Target

The following describes the security requirements for using Amazon RDS Migration Tool with a Microsoft SQL Server source or target.

The Amazon RDS Migration Tool user must have at least the `db_owner` user role on the Microsoft SQL Server database you are connecting to.

A Microsoft SQL Server system administrator must provide this permission for all Amazon RDS Migration Tool users.

Microsoft SQL Server Target Data Types

The Microsoft SQL Server target for Amazon RDS Migration Tool supports most Microsoft SQL Server data types. The following table shows the Microsoft SQL Server target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped from the source, see the section for the source endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 6–5 Microsoft SQL Server Target Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	Microsoft SQL Server Data Types
BOOLEAN	TINYINT
BYTES	VARBINARY(length)
DATE	For Microsoft SQL Server 2008 and later: DATE For earlier versions: If scale is < or = 3: DATETIME In all other cases: VARCHAR (37)
TIME	For Microsoft SQL Server 2008 and later: DATETIME2 (%d) For earlier versions: If scale is < or = 3: DATETIME In all other cases: VARCHAR (37)
DATETIME	For Microsoft SQL Server 2008 and later: DATETIME2 (scale) For earlier versions: If scale is < or = 3: DATETIME In all other cases: VARCHAR (37)
INT1	SMALLINT
INT2	SMALLINT
INT4	INT
INT8	BIGINT
NUMERIC	NUMERIC (p,s)
REAL4	REAL

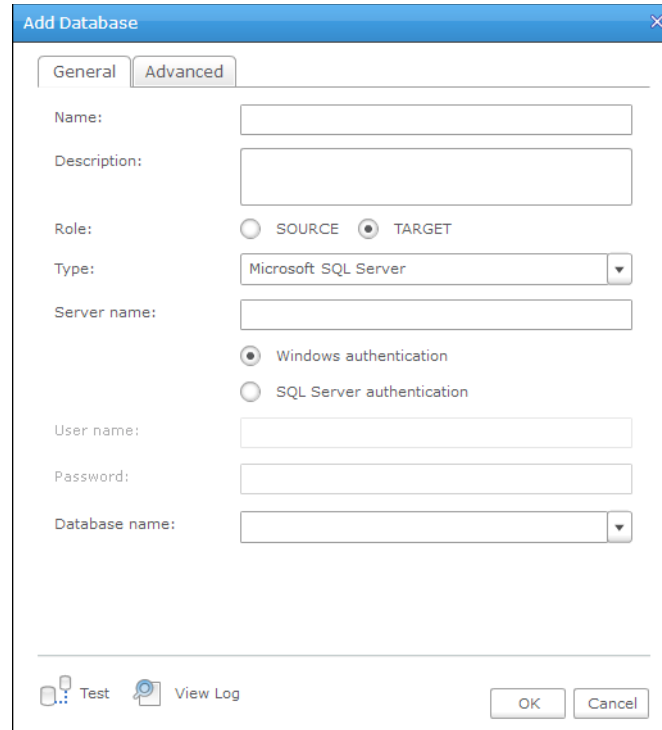
Table 6–5 (Cont.) Microsoft SQL Server Target Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	Microsoft SQL Server Data Types
REAL8	FLOAT
STRING	<p>If column is date or time then:</p> <p>For Microsoft SQL Server 2008 and later: DATETIME2</p> <p>For earlier versions:</p> <p>If scale is < or = 3: DATETIME</p> <p>In all other cases: VARCHAR (37)</p> <p>If the column is <i>not</i> a date or time: VARCHAR (length)</p>
UINT1	TINYINT
UINT2	SMALLINT
UINT4	INT
UINT8	BIGINT
WSTRING	NVARCHAR (length)
BLOB	<p>VARBINARY (max)</p> <p>IMAGE</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of BLOBs for a specific task.</p> <p>BLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
CLOB	<p>VARCHAR (max)</p> <p>TEXT</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of CLOBs for a specific task.</p> <p>During CDC, CLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>
NCLOB	<p>NVARCHAR (max)</p> <p>NTEXT</p> <p>To use this data type with Amazon RDS Migration Tool, you must enable the use of NCLOBs for a specific task.</p> <p>During CDC, NCLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>

Setting up a Microsoft SQL Server Database as an Amazon RDS Migration Tool Target

You can add a Microsoft SQL Server database to Amazon RDS Migration Tool to use as a target. For information on how to add databases, see [Working with Databases](#). When you select **Microsoft SQL Server Target** as the database type the following dialog box is displayed:

Figure 6–2 Microsoft SQL Server Target Database



To add a Microsoft SQL Server target database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the Microsoft SQL Server database. This is optional.
4. Select **TARGET** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **Microsoft SQL Server** as the database **Type**.
6. Type the **Microsoft SQL Server name**. This is the name of the computer with the Microsoft SQL Server instance you want to work with.

Note:

You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a Microsoft SQL Server Target Database](#).

7. Select **Windows authentication** or **User name/password**.

If you select **Windows authentication**, you will work with the user credentials for the Windows domain. This privilege must be configured in the Microsoft SQL Server database by the system administrator.

Note: When using Windows authentication be sure that the user account that is associated with the **Amazon RDS Migration Server** service has Network read and write permissions. This must be configured by a Windows system administrator.

See also [Working with Windows Authentication](#).

If you select **User name/password**, type the Microsoft SQL Server authentication information (User Name, Password) for the authorized user for this Microsoft SQL Server database. If you do not know this information, see the Microsoft SQL Server System Administrator.

Note:

- The User Name/Password information is mandatory if you select **User name/password**. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. For more information, see [Using Advanced Properties for a Microsoft SQL Server Target Database](#).
- This information is case sensitive.
- **Important:** Make sure that the Microsoft SQL Server user has the correct access privileges. For information on how to provide the required privileges, see [Security Considerations for a Microsoft SQL Server Target](#).
- You can set custom properties in the **Advanced** tab. For more information, see [Using Advanced Properties for a Microsoft SQL Server Target Database](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

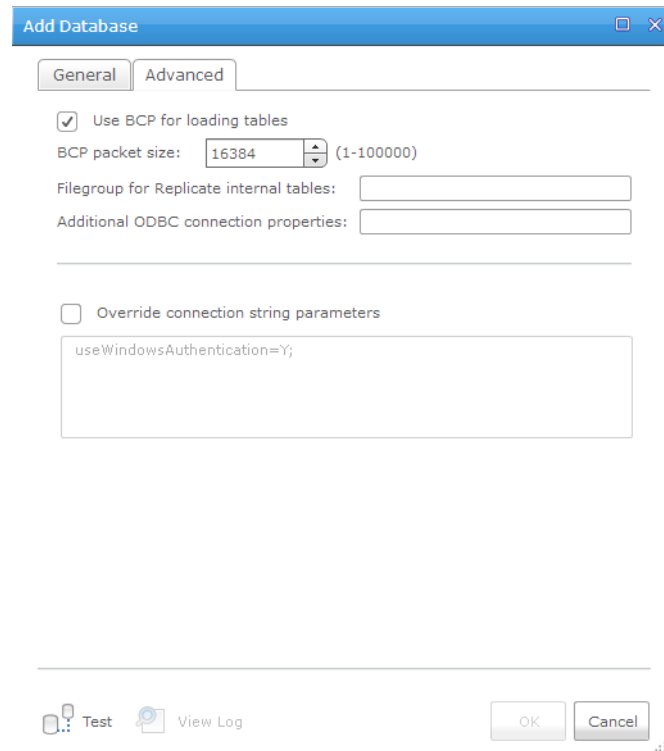
To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

8. Type the **Database name** or select one from the list of available databases. This is the name of the database where you are replicating the data.

Using Advanced Properties for a Microsoft SQL Server Target Database

You can use advanced properties and change the default settings for various parameters in a Microsoft SQL Server source or target by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

Figure 6–3 Microsoft SQL Server Target Advanced Tab



You can set the following properties in the **Advanced** tab:

- **Use BCP for loading tables:** Select this to transfer data for full-load operations using BCP. This feature is not supported when using the Simba Microsoft SQL Server ODBC Driver for Linux.

Note: When the target table contains an identity column that does not exist in the source table, you must disable the **Use BCP for loading tables** option.

- **BCP packet size:** The maximum size of the packets (in bytes) used to transfer data using BCP.
- **Filegroup for Amazon RDS Migration Tool internal tables:** Optionally, specify a filegroup for the Amazon RDS Migration Tool internal tables. When the replication task starts, all the internal Amazon RDS Migration Tool control tables (`amazon_apply_exception`, `amazon_apply`, `amazon_changes`) will be created on the specified filegroup.

The following is an example of a command for creating a filegroup:

```
ALTER DATABASE replicate
ADD FILEGROUP Test1FG1;
GO
```

```
ALTER DATABASE replicate
ADD FILE
(
    NAME = test1dat5,
    FILENAME = 'C:\temp\DATA\t1dat5.ndf',
    SIZE = 5MB,
    MAXSIZE = 100MB,
    FILEGROWTH = 5MB
)
TO FILEGROUP Test1FG1;
GO
```

- **Additional ODBC connection properties:** Specify any additional ODBC connection parameters that you want to use.
- **Override connection string parameters:** Select this to use a custom connect string.

The following is an example of a Microsoft SQL Server connection string:

```
server=123.123.123.12;controlTablesFileGroup=replicate;
useWindowsAuthentication=N;BCPPacketSize=16387;username=erte;password=dewe7d3;
useBCPFullLoad=Y;database=hr;
```

For information on the properties you can include in a custom connect string for Microsoft SQL Server, see [Configuration Properties for a Microsoft SQL Server Source](#).

Using a SAP Sybase ASE Database as a Source or Target

This chapter describes how to set up and use a SAP Sybase ASE database as the source or target endpoint in a replication task. It contains the following topics:

- [Prerequisites](#)
- [Limitations](#)
- [Using a SAP Sybase ASE Database as a Source](#)
- [Using a SAP Sybase ASE Database as a Target](#)

Prerequisites

Before you begin to work with a SAP Sybase ASE database as a source or target in Amazon RDS Migration Tool, make sure that the SAP Sybase ASE database with the tables that are necessary for replication is available in your network.

Note:

- Amazon RDS Migration Tool must be installed on any Windows computer in your network.
 - A SAP Sybase ASE account with the required access privileges is required.
-
-

Make sure the following prerequisites have been met:

- SAP Sybase ASE ODBC 64-bit client installed on the computer where Amazon RDS Migration Tool is located.
- SAP Sybase ASE replication enabled for tables using the `sp_setreptable` command or privileges to enable it automatically.
- RepAgent must be disabled on the SAP Sybase ASE database.
- When replicating to SAP Sybase ASE 15.7 installed on a Windows machine configured with a non-Latin language (e.g. Chinese), Amazon RDS Migration Tool requires Sybase 15.7 SP121 to be installed on the SAP Sybase ASE machine.

Table 7–1 (Cont.) SAP Sybase ASE Connect String Standard Properties

Property	Description/Format	Default Value
username	The name of the user authorized to access the SAP Sybase ASE database.	
password	The password for the authorized user.	
Port	The port through which to connect to the SAP Sybase ASE database.	5000
enableReplication	Set to Y to automatically enable SAP Sybase ASE replication. This is only required if SAP Sybase ASE replication has not been enabled already. For more information, see Prerequisites .	Y
additionalConnectionProperties	Any additional ODBC connection parameters that you want to specify.	

Security Requirements

To use SAP Sybase ASE database as a source in a RDS Migration Tool task, the following permissions are required:

- sa_role
- replication_role
- sybase_ts_role

If the **Automatically enable Sybase replication** option is enabled (in the **Advanced** tab), RDS Migration Tool also needs permission to run the stored procedure `sp_setreptable`.

For information on the **Automatically enable SAP Sybase ASE replication** option, see [Using Advanced Properties for a SAP Sybase ASE Source](#).

SAP Sybase ASE Database Source Data Types

The following table shows the SAP Sybase ASE database source data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 7–2 SAP Sybase ASE Database Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

SAP Sybase ASE Source Data Types	Amazon RDS Migration Tool Data Types
BIGINT	INT8
INT	INT4
SMALLINT	INT2

Table 7–2 (Cont.) SAP Sybase ASE Database Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

SAP Sybase ASE Source Data Types	Amazon RDS Migration Tool Data Types
TINYINT	UINT1
DECIMAL	NUMERIC
NUMERIC	NUMERIC
FLOAT	REAL8
DOUBLE	REAL8
REAL	REAL4
MONEY	NUMERIC
SMALLMONEY	NUMERIC
DATETIME	DATETIME
SMALLDATETIME	DATETIME
DATE	DATE
TIME	TIME
CHAR	STRING
UNICHAR	UNICODE CHARACTER
NCHAR	WSTRING
VARCHAR	STRING
UNIVARCHAR	UNICODE
NVARCHAR	WSTRING
BINARY	BYTES
VARBINARY	BYTES
BIT	BOOLEAN
TEXT	CLOB
UNITEXT	NCLOB
IMAGE	BLOB

Non-Supported Data Types

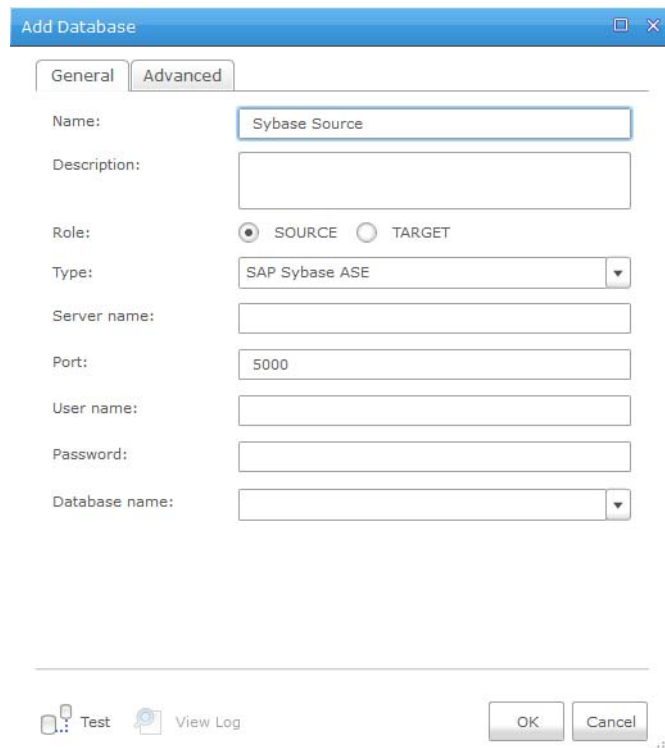
Source SAP Sybase ASE tables with columns of the following SAP Sybase ASE data types cannot be replicated. Replicated columns with these data types will show as null.

- UDT

Setting up a SAP Sybase ASE Database as a Source in Amazon RDS Migration Tool

You can add a SAP Sybase ASE database to Amazon RDS Migration Tool to use as a source. For information on how to add databases, see [Working with Databases](#). When you select **SAP Sybase ASE Source** as the database type the following dialog box is displayed:

Figure 7-1 SAP Sybase ASE Source Database



Notes: You can also use SAP Sybase ASE files as a source. For more information, see [Using the Amazon RDS Migration Tool File Channel](#).

To add a SAP Sybase ASE source database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the SAP Sybase ASE database. This is optional.
4. Select **SOURCE** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **SAP Sybase ASE** as the database **Type**.
6. In the **Server Name** field, enter the host name or IP address of the computer on which the SAP Sybase ASE database is installed.

Note:

- This information is case sensitive.
- You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a SAP Sybase ASE Source](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

7. Optionally, change the default port (5000).
8. Type the SAP Sybase ASE authentication information (**User Name**, **Password**) for the authorized user for this SAP Sybase ASE database. If you do not know this information, see your SAP Sybase ASE Database Administrator (DBA).

Note:

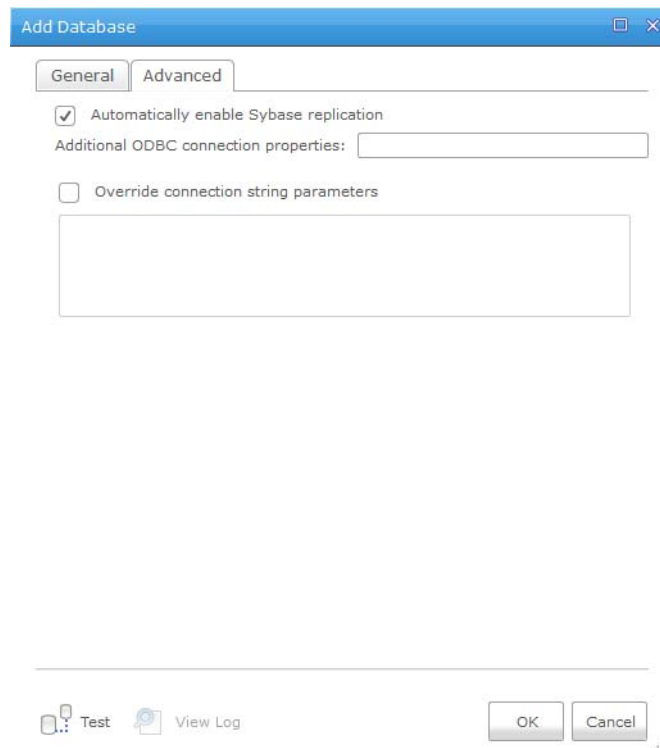
- This information is required. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. See [Using Advanced Properties for a SAP Sybase ASE Source](#) for more information.
 - This information is case sensitive.
 - **Important:** Make sure that the SAP Sybase ASE user entered in the SAP Sybase ASE Authentication section has the correct access privileges. For information on how to provide the required privileges, see [Security Requirements](#).
 - If you want to set custom properties for this database, see [Using Advanced Properties for a SAP Sybase ASE Source](#).
-
-

9. In the **Database name** field, enter the SAP Sybase ASE database name.

Using Advanced Properties for a SAP Sybase ASE Source

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the **Add Database** dialog box.

Figure 7-2 SAP Sybase ASE Source Advanced Tab



You can set the following parameters:

- **Automatically enable SAP Sybase ASE replication:** Select this to automatically enable SAP Sybase ASE replication. This is only required if SAP Sybase ASE replication has not been enabled already. For more information, see [Prerequisites](#).
- **Additional ODBC connection properties:** Specify any additional ODBC connection parameters that you want to use.

Note: If the user name or password specified in the **General** tab contains non-Latin characters (e.g. Chinese), the following property is required:

```
charset=gb18030
```

- **Override connection string parameters:** Select this to create a custom connect string.

The following is an example of a SAP Sybase ASE connection string:

```
server=SybaseASE1;port=5421;database=MyASE;username=Admin;password=7yyfhn85;enableReplication=Y
```

For a description of the supported connection parameters, see [Configuration Properties for the SAP Sybase ASE Source](#).

Using a SAP Sybase ASE Database as a Target

The following topics describe how to use a SAP Sybase ASE database as the target endpoint in an Amazon RDS Migration Tool task.

- [Configuration Properties for the SAP Sybase ASE Target](#)
- [Security Requirements](#)
- [SAP Sybase ASE Database Target Data Types](#)
- [Non-Supported Data Types](#)
- [Setting up a SAP Sybase ASE Database as a Target in Amazon RDS Migration Tool](#)

Configuration Properties for the SAP Sybase ASE Target

The following section describes the configuration properties available for a SAP Sybase ASE target.

Table 7–3 SAP Sybase ASE Connect String Standard Properties

Property	Description/Format	Default Value
server	The name or IP address of the computer on which the SAP Sybase ASE database is installed.	
database	The name of the SAP Sybase ASE database to which you want to connect.	
username	The name of the user authorized to access the SAP Sybase ASE database.	
password	The password for the authorized user.	
Port	The port through which to connect to the SAP Sybase ASE database.	5000
additionalConnectionProperties	Any additional ODBC connection parameters that you want to specify.	

Security Requirements

You must provide SAP Sybase ASE account access to the Amazon RDS Migration Tool user. This user must have read/write privileges in the SAP Sybase ASE database.

SAP Sybase ASE Database Target Data Types

The following table shows the SAP Sybase ASE database target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 7–4 Supported SAP Sybase ASE Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	SAP Sybase ASE Data Types
BOOLEAN	BIT
BYTES	VARBINARY (Length)
DATE	DATE
TIME	TIME
TIMESTAMP	If scale is => 0 and =< 6, then: BIGDATETIME If scale is => 7 and =< 9, then: VARCHAR (37)
INT1	TINYINT
INT2	SMALLINT
INT4	INTEGER
INT8	BIGINT
NUMERIC	NUMERIC (p,s)
REAL4	REAL
REAL8	DOUBLE PRECISION
STRING	VARCHAR (Length)
UINT1	TINYINT
UINT2	UNSIGNED SMALLINT
UINT4	UNSIGNED INTEGER
UINT8	UNSIGNED BIGINT
WSTRING	VARCHAR (Length)
BLOB	IMAGE
CLOB	UNITEXT
NCLOB	TEXT

Non-Supported Data Types

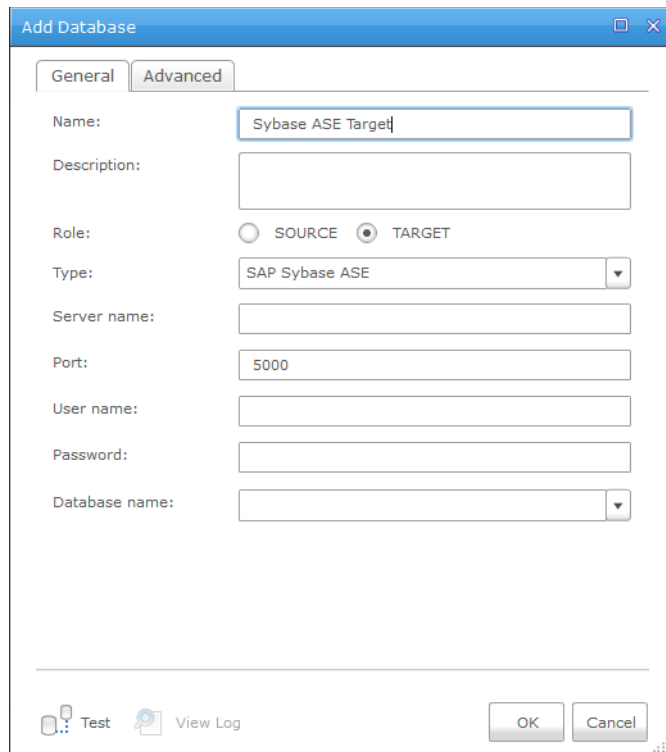
Target SAP Sybase ASE tables with columns of the following SAP Sybase ASE data types cannot be replicated. Replicated columns with these data types will show as null.

- UDT

Setting up a SAP Sybase ASE Database as a Target in Amazon RDS Migration Tool

You can add a SAP Sybase ASE database to Amazon RDS Migration Tool to use as a target. For information on how to add databases, see [Working with Databases](#). When you select **SAP Sybase ASE** as the target database type the following dialog box is displayed:

Figure 7-3 SAP Sybase ASE Target Database: General Tab



To add a SAP Sybase ASE target database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the SAP Sybase ASE database. This is optional.
4. Select **TARGET** as the database **Role**.
 You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **SAP Sybase ASE** as the database **Type**.
6. In the **Server Name** field, enter the host name or IP address of the computer on which the SAP Sybase ASE database is installed.

Note:

- This information is case sensitive.
- You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a SAP Sybase ASE Target](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

7. Optionally, change the default port (5000).
 8. Type the SAP Sybase ASE authentication information (**User Name, Password**) for the authorized user for this SAP Sybase ASE database. If you do not know this information, see your SAP Sybase ASE Database Administrator (DBA).
-
-

Note:

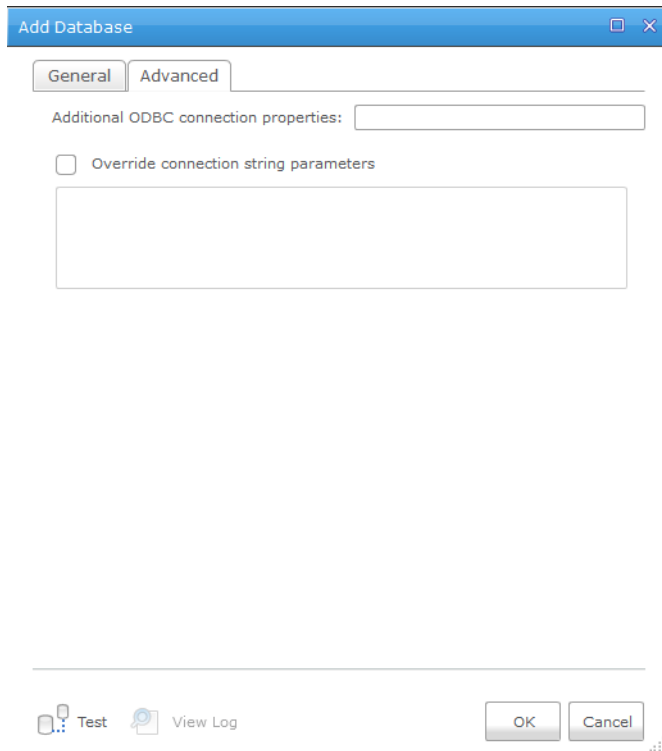
- This information is required. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. See [Using Advanced Properties for a SAP Sybase ASE Target](#) for more information.
 - This information is case sensitive.
 - **Important:** Make sure that the SAP Sybase ASE user entered in the SAP Sybase ASE Authentication section has the correct access privileges. For information on how to provide the required privileges, see [Security Requirements](#).
 - If you want to set custom properties for this database, see [Using Advanced Properties for a SAP Sybase ASE Target](#).
-
-

9. In the **Database name** field, enter the SAP Sybase ASE database name.

Using Advanced Properties for a SAP Sybase ASE Target

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the **Add Database** dialog box.

Figure 7–4 SAP Sybase ASE Target: Advanced Tab



You can set the following parameters:

- Additional ODBC connection properties:** Specify any additional ODBC connection parameters that you want to use.

Note: If the user name or password specified in the **General** tab contains non-Latin characters (e.g. Chinese), the following property is required:

```
charset=gb18030
```

- Override connection string parameters:** Select this to create a custom connect string.

The following is an example of a SAP Sybase ASE connection string:

```
server=SybaseASE1;port=5421;database=MyASE;username=Admin;password=7yyfhn85;
```

For a description of the supported connection parameters, see [Configuration Properties for the SAP Sybase ASE Source](#).

Using a MySQL Database as a Source or Target

This chapter describes how to set up and use a MySQL database as the source or target endpoint in a replication task. It contains the following topics:

- [Prerequisites](#)
- [Limitations](#)
- [Using a MySQL Database as a Source](#)
- [Using a MySQL Database as a Target](#)

Prerequisites

Before you begin to work with a MySQL database as a source or target in Amazon RDS Migration Tool, make sure that the following prerequisites have been met.

General Prerequisites

- Amazon RDS Migration Tool installed on Windows or Linux in your network.
- A MySQL account with the required [Security Requirements](#).
- A MySQL database with the tables that you want to replicate should be accessible in your network.

The following MySQL editions are supported:

- MySQL Community Edition
- MySQL Standard Edition
- MySQL Enterprise Edition
- MySQL Cluster Carrier Grade Edition

Amazon RDS Migration Tool Server for Windows

The following section describes the steps you need to perform to work with Amazon RDS Migration Tool for Windows and MySQL as a source or target endpoint in a Amazon RDS Migration Tool task:

- MySQL ODBC 5.2.6 64-bit client or above must be installed on the same computer as Amazon RDS Migration Tool.

Amazon RDS Migration Tool Server for Linux

The following section describes the steps you need to perform to work with Amazon RDS Migration Tool for Linux and MySQL as a source or target endpoint in an Amazon RDS Migration Tool task:

1. On the Amazon RDS Migration Tool machine, install MySQL Client 5.2 or above for Linux.
2. Makes sure that the `/etc/odbcinst.ini` file contains the following entry for MySQL, as in the following example:

```
[MySQL ODBC 5.2 Unicode Driver]
Driver      = /usr/lib64/libmyodbc5w.so
UsageCount  = 1
```

Enable Binary Logging

To enable binary logging, the following parameters must be configured in MySQL's `my.ini` (Windows) or `my.cnf` (UNIX) files.

Table 8–1 Required my.ini/my.cnf Parameters for Binary Logging

Parameter	Value
server_id	Any value from 1. Example: server_id=1
log-bin=<path>	Path to the binary log file (without an extension). Example: log-bin=E:\MySql_Logs\BinLog
binlog_format	Must be: binlog_format=row
expire_logs_days	To prevent disk space issues, it is strongly recommended not to use the default value (0). Example: expire_logs_days=5
binlog_checksum	Must be: binlog-checksum=none
binlog_row_image	Must be: binlog_row_image=full

Cluster Prerequisites

To be able to replicate clustered (NDB) tables (i.e. by connecting Amazon RDS Migration Tool to any of the cluster nodes), the following parameters must be configured in MySQL's `my.ini` (Windows) or `my.cnf` (UNIX) files.

Table 8–2 Required `my.ini/my.cnf` Parameters for Cluster Replication

Parameter	Value
<code>ndb_log_bin</code>	Must be: <code>ndb_log_bin=on</code> This ensures that changes in clustered tables will be logged to the binary log.
<code>ndb_log_update_as_write</code>	Must be: <code>ndb_log_update_as_write=OFF</code> This prevents writing UPDATES as INSERTs in the binary log.
<code>ndb_log_updated_only</code>	Must be: <code>ndb_log_updated_only=OFF</code> Ensures that the binary log will contain the entire row and not just the changed columns.

Limitations

- The following DDLs are not supported:
 - Truncate Partition
 - Drop Table
 - Rename Table
- Using the `alter table <table_name> add column <column_name>` statement to add columns to the beginning or to the middle of a table is not supported.
- **MySQL Source:** Changes are not captured from tables whose names contain both upper and lower case characters.
- **MySQL Source:** The `AR_H_USER` header column is currently not supported. For information on using header columns, see [Header Columns](#).
- **MySQL Source:** If the MySQL database is stopped during Full Load, the Full Load will end successfully, but the tables on the target may have less rows than the source tables. If this should happen, either restart the task or reload the tables with the missing rows.
- **MySQL Target:** When updating a column's value to its existing value, a zero rows affected is returned from MySQL (unlike Oracle and SQL Server that perform an update of one row). This generates an entry in the `amazon_apply_exceptions` control table and the following warning:

Some changes from the source database had no impact when applied to the target database. See `amazon_apply_exceptions` table for details.

Using a MySQL Database as a Source

The following topics describe how to use a MySQL database as the source endpoint in an Amazon RDS Migration Tool task.

- [Security Requirements](#)
- [Configuration Properties for the MySQL Source](#)
- [MySQL Database Source Data Types](#)
- [Setting up a MySQL Database as a Source in Amazon RDS Migration Tool](#)

Security Requirements

The Amazon RDS Migration Tool user must have the `ReplicationAdmin` role with the following privileges (according to task type):

- `REPLICATION CLIENT` - Required for Change Processing tasks only. In other words, Full Load only tasks do not require this privilege.
- `REPLICATION SLAVE` - Required for Change Processing tasks only. In other words, Full Load only tasks do not require this privilege.
- `SUPER` - Only required in versions *prior* to MySQL 5.6.6.

The Amazon RDS Migration Tool user must also have `SELECT` privileges for the source tables designated for replication.

Setting up Amazon RDS MySQL for CDC (Change Data Capture)

To enable CDC with Amazon RDS MySQL, you need to use Amazon RDS MySQL version 5.6 or higher.

1. Follow the instructions provided by AWS to create a new Parameter Group (see the Binary Logging Format section):
http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_LogAccess.Concepts.MySQL.html
2. When creating the new Parameter Group, set the following values:
 - `binlog_format=row`
 - `binlog_checksum=NONE`
3. Save the new Parameter Group.
4. If you have an existing instance of Amazon RDS MySQL, edit the instance to use the parameters specified in **Step 2** above. Or, if you are provisioning a new instance of Amazon RDS MySQL, reference the new Parameter Group created in **Step 1** above.
5. Use the `binlog retention hours` parameter to specify the number of hours for RDS to retain binary logs. It is recommended to specify at least 12 hours.

To check the current parameter value:

```
CALL mysql.rds_show_configuration;
```

To set the parameter:

```
call mysql.rds_set_configuration('binlog retention hours', 12);
```

Configuration Properties for the MySQL Source

The following section describes the configuration properties available for a MySQL source.

Table 8–3 MySQL Connect String Standard Properties

Property	Description/Format	Default Value
server	The name or IP address of the computer on which the MySQL database is installed.	
database	The name of the MySQL database to which you want to connect.	
username	The name of the user authorized to access the MySQL database.	
password	The password for the authorized user.	
Port	The port through which to connect to the MySQL database.	3306

Table 8–3 (Cont.) MySQL Connect String Standard Properties

Property	Description/Format	Default Value
additionalConnectionProperties	Any additional ODBC connection parameters that may be required. For more information, see Additional ODBC connection properties .	

MySQL Database Source Data Types

The following table shows the MySQL database source data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 8–4 MySQL Database Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

MySQL Source Data Types	Amazon RDS Migration Tool Data Types
INT	INT4
BIGINT	INT8
MEDIUMINT	INT4
TINYINT	INT1
DECIMAL (10)	NUMERIC (10,0)
BINARY	BYTES (1)
BIT	BOOLEAN
BIT (64)	BYTES (8)
BLOB	BYTES (66535)
LONGBLOB	BLOB
MEDIUMBLOB	BLOB
TINYBLOB	BYTES (255)
DATE	DATE
DATETIME	DATETIME
TIME	STRING
TIMESTAMP	DATETIME
YEAR	INT2
DOUBLE	REAL8

Table 8–4 (Cont.) MySQL Database Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

MySQL Source Data Types	Amazon RDS Migration Tool Data Types
FLOAT	REAL (DOUBLE)
<p>If the FLOAT values are not in the range specified below, use a transformation to map FLOAT to STRING. For an explanation of how to do this, see Using the Transform Tab.</p> <p>Supported FLOAT range: - 1.79E+308 to -2.23E-308, 0 and 2.23E-308 to 1.79E+308</p>	
VARCHAR (45)	WSTRING (45)
VARCHAR (2000)	WSTRING (2000)
VARCHAR (4000)	WSTRING (4000)
VARBINARY (4000)	BYTES (4000)
VARBINARY (2000)	BYTES (2000)
CHAR	WSTRING
TEXT	WSTRING (65535)
LONGTEXT	NCLOB
MEDIUMTEXT	NCLOB
TINYTEXT	WSTRING (255)
GEOMETRY	BLOB
POINT	BLOB
LINestring	BLOB
POLYGON	BLOB
MULTIPOINT	BLOB
MULTILINESTRING	BLOB
MULTIPOLYGON	BLOB
GEOMETRYCOLLECTION	BLOB

Note: If the DATETIME and TIMESTAMP data types are specified with a “zero” value (i.e. 0000-00-00), you need to make sure that the target database in the replication task supports “zero” values for the DATETIME and TIMESTAMP data types. If they are not supported, you can use a transformation to specify a supported value (e.g. 1970.) Otherwise, they will be recorded as null on the target.

Data Types Supported in Full Load Only

The following MySQL data types are supported in Full Load only:

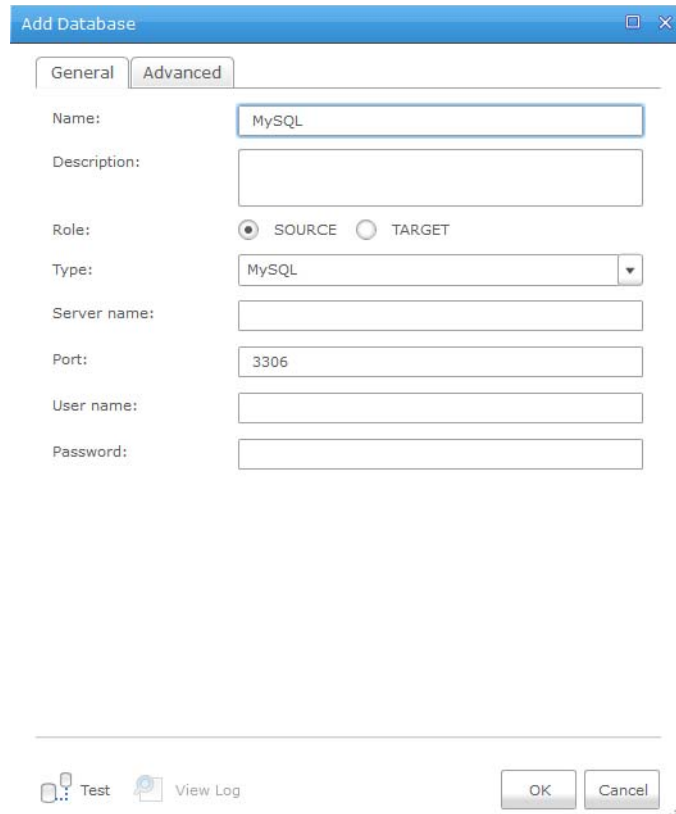
MySQL Source Data Types	Amazon RDS Migration Tool Data Types
ENUM	STRING

MySQL Source Data Types	Amazon RDS Migration Tool Data Types
SET	STRING

Setting up a MySQL Database as a Source in Amazon RDS Migration Tool

You can add a MySQL database to Amazon RDS Migration Tool to use as a source. For information on how to add databases, see [Working with Databases](#). When you select **MySQL Source** as the database type the following dialog box is displayed:

Figure 8–1 MySQL Source Database



Notes: You can also use MySQL files as a source. For more information, see [Using the Amazon RDS Migration Tool File Channel](#).

To add a MySQL source database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the MySQL database. This is optional.
4. Select **SOURCE** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **MySQL** as the database **Type**.

6. In the **Server Name** field, enter the host name or IP address of the computer on which the MySQL database is installed.

Note:

- This information is case sensitive.
- You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a MySQL Source](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

7. Optionally, change the default port (5000).
8. Type the MySQL authentication information (**User Name**, **Password**) for the authorized user for this MySQL database. If you do not know this information, see your MySQL Database Administrator (DBA).

Note:

- This information is required. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. See [Using Advanced Properties for a MySQL Source](#) for more information.
 - This information is case sensitive.
 - **Important:** Make sure that the MySQL user entered in the MySQL Authentication section has the correct access privileges. For information on how to provide the required privileges, see [Security Requirements](#).
 - If you want to set custom properties for this database, see [Using Advanced Properties for a MySQL Source](#).
-
-

Selecting a Schema

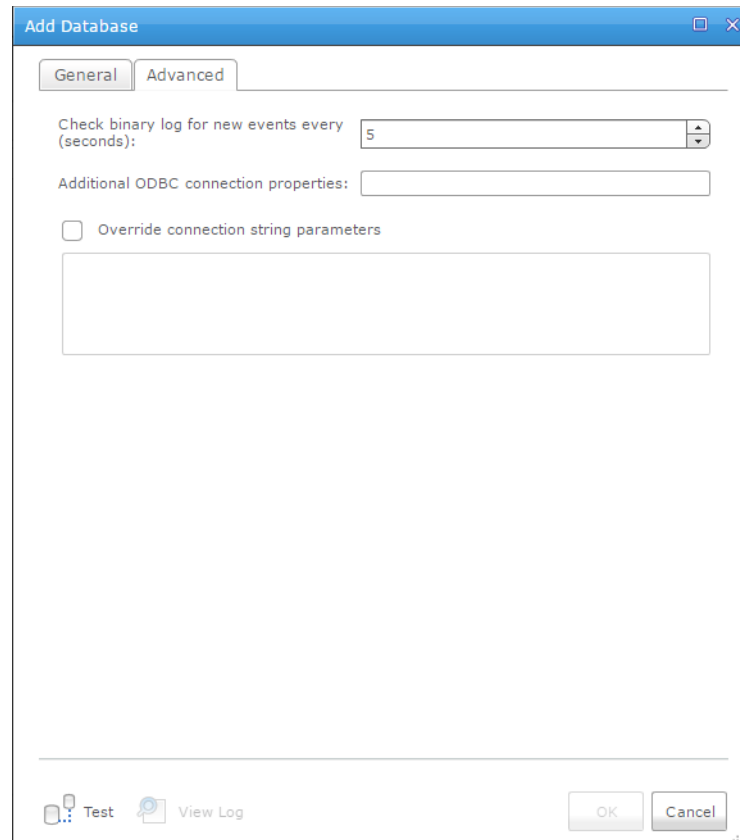
You can choose which MySQL database to access. After configuring the MySQL source database connection settings, open the **Select Tables** dialog box (by clicking the **Table Selection** button on the right of the console) and select which schema to use from the **Schema** drop down list.

See also: [Chapter 14, "Designing Tasks"](#).

Using Advanced Properties for a MySQL Source

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the **Add Database** dialog box.

Figure 8–2 MySQL Source Advanced Tab



You can set the following parameters:

- **Check binary log for new events every:** Specify how often to check the binary log for changes when the databases is idle.
- **Additional ODBC connection properties:** Specify any additional ODBC connection parameters that may be required.

Note: Amazon RDS Migration Tool assumes that MySQL Client 5.2 for Linux or MySQL ODBC Client 5.2.6 64-bit for Windows has been installed on the Amazon RDS Migration Server machine. If a later version has been installed, you need to specify the version number in the **Additional ODBC connection properties** field, as in the following example:

```
provider=MySQL ODBC 5.3 Unicode Driver
```

- **Override connection string parameters:** Select this to create a custom connect string.

The following is an example of a MySQL connection string:

```
server=MySQL1;port=5421;database=MyDB;username=Admin;  
password=7yyfhn85;
```

For a description of the supported connection parameters, see [Configuration Properties for the MySQL Source](#).

Using a MySQL Database as a Target

The following topics describe how to use a MySQL database as the target endpoint in an Amazon RDS Migration Tool task.

- [Configuration Properties for the MySQL Target](#)
- [Security Requirements](#)
- [MySQL Database Target Data Types](#)
- [Setting up a MySQL Database as a Target in Amazon RDS Migration Tool](#)

Configuration Properties for the MySQL Target

The following section describes the configuration properties available for a MySQL target.

Table 8–5 MySQL Connect String Standard Properties

Property	Description/Format	Default Value
server	The name or IP address of the computer on which the MySQL database is installed.	
database	The name of the MySQL database to which you want to connect.	
username	The name of the user authorized to access the MySQL database.	
password	The password for the authorized user.	
Port	The port through which to connect to the MySQL database.	3306
maxFileSize	The maximum size (in KB) of a CSV file before the file is loaded into the MySQL target.	32000 KB
additionalConnectionProperties	Any additional ODBC connection parameters that may be required. For more information, see Additional ODBC connection properties .	

Security Requirements

You must provide MySQL account access to the Amazon RDS Migration Tool user. This user must have read/write privileges in the MySQL database.

MySQL Database Target Data Types

The following table shows the MySQL database target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 8–6 Supported MySQL Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	MySQL Data Types
BOOLEAN	BOOLEAN
BYTES	If length is => 1 and =< 65535, then: VARBINARY (Length) If length is => 65535 and =< 2147483647, then: LONGLOB
DATE	DATE
TIME	TIME
TIMESTAMP	If scale is => 0 and =< 6, then: BIGDATETIME If scale is => 7 and =< 9, then: VARCHAR (37)
INT1	TINYINT
INT2	SMALLINT
INT4	INTEGER
INT8	BIGINT
NUMERIC	DECIMAL (p,s)
REAL4	FLOAT
REAL8	DOUBLE PRECISION
STRING	If length is => 1 and =< 21845, then: VARCHAR (Length) If length is => 21846 and =< 2147483647, then: LONGTEXT
UINT1	UNSIGNED TINYINT
UINT2	UNSIGNED SMALLINT
UINT4	UNSIGNED INTEGER
UINT8	UNSIGNED BIGINT
WSTRING	If length is => 1 and =< 32767, then: VARCHAR (Length) If length is => 32768 and =< 2147483647, then: LONGTEXT

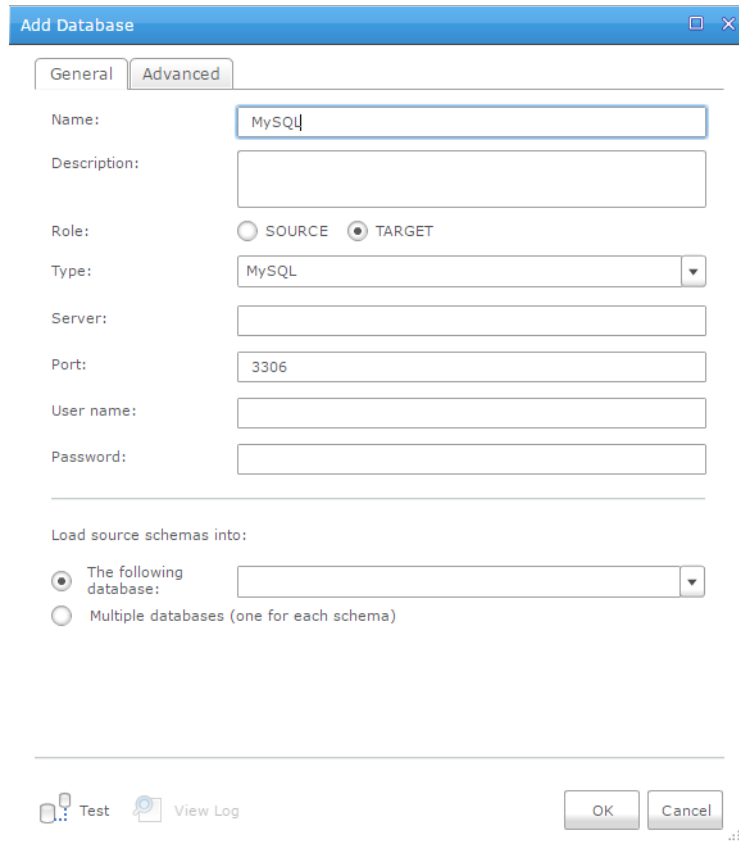
Table 8–6 (Cont.) Supported MySQL Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	MySQL Data Types
BLOB	If length is => 1 and =< 65535, then: BLOB If length is => 65536 and =< 2147483647, then: LONGBLOB If length is => 0 and =< 0, then: LONGBLOB (Full Lob Support)
NCLOB	If length is => 1 and =< 65535, then: TEXT If length is => 65536 and =< 2147483647, then: LONGTEXT - CHARACTER SET: ucs2 If length is => 0 and =< 0, then: LONGTEXT - CHARACTER SET: ucs2 (Full Lob Support)
CLOB	If length is => 1 and =< 65535, then: TEXT If length is => 65536 and =< 2147483647, then: LONGTEXT If length is => 0 and =< 0, then: LONGTEXT (Full Lob Support)

Setting up a MySQL Database as a Target in Amazon RDS Migration Tool

You can add a MySQL database to Amazon RDS Migration Tool to use as a target. For information on how to add databases, see [Working with Databases](#). When you select **MySQL** as the target database type the following dialog box is displayed:

Figure 8–3 MySQL Target Database: General Tab



The screenshot shows the 'Add Database' dialog box with the following fields and options:

- Name:** MySQL
- Description:** (empty text box)
- Role:** SOURCE (radio button), TARGET (radio button, selected)
- Type:** MySQL (dropdown menu)
- Server:** (empty text box)
- Port:** 3306
- User name:** (empty text box)
- Password:** (empty text box)
- Load source schemas into:**
 - The following database: (dropdown menu)
 - Multiple databases (one for each schema) (radio button)
- Buttons:** Test, View Log, OK, Cancel

To add a MySQL target database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the MySQL database. This is optional.
4. Select **TARGET** as the database **Role**.
You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.
5. Select **MySQL** as the database **Type**.
6. In the **Server** field, enter the host name or IP address of the computer on which the MySQL database is installed.

Note:

- This information is case sensitive.
- You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a MySQL Target](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

7. Optionally, change the default port (5000).
8. Type the MySQL authentication information (**User Name**, **Password**) for the authorized user for this MySQL database. If you do not know this information, see your MySQL Database Administrator (DBA).

Note:

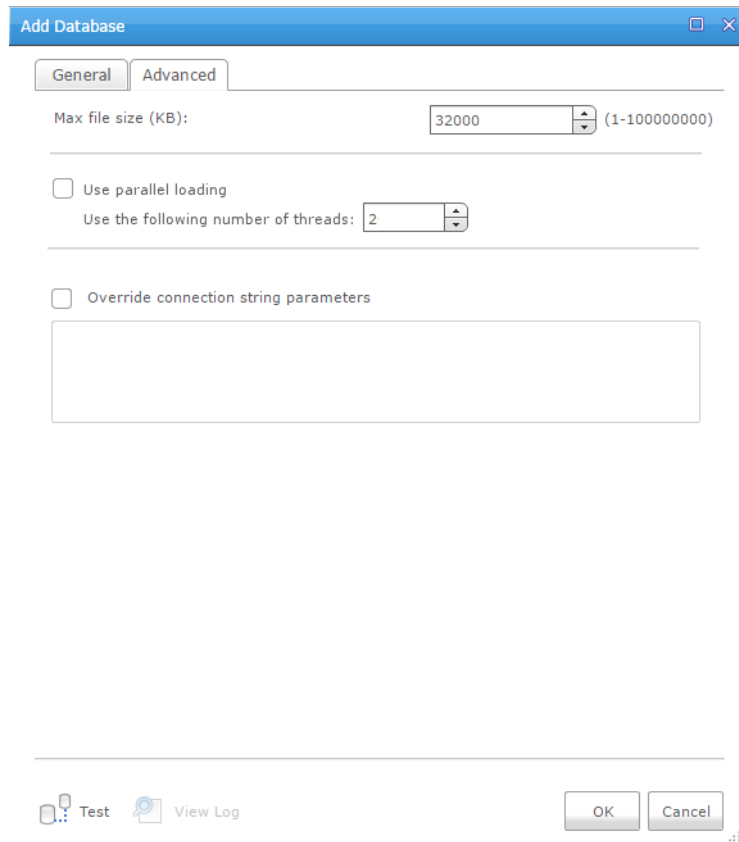
- This information is required. If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. See [Using Advanced Properties for a MySQL Target](#) for more information.
 - This information is case sensitive.
 - **Important:** Make sure that the MySQL user entered in the MySQL Authentication section has the correct access privileges. For information on how to provide the required privileges, see [Security Requirements](#).
 - If you want to set custom properties for this database, see [Using Advanced Properties for a MySQL Target](#).
-
-

9. Select one of the following **Load source schemas into** options:
 - **The following database** - When this option is selected, all source schemas will be loaded into the selected database.
 - **Multiple databases** - When this option is selected, each of the source schemas will be loaded into its corresponding database.

Using Advanced Properties for a MySQL Target

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the **Add Database** dialog box.

Figure 8–4 MySQL Target: Advanced Tab



You can set the following parameters:

- **Max file size (KB):** Select or type the maximum size (in KB) of a CSV file before it is loaded into the MySQL target database. The default value is 32000 KB.
- **Use parallel loading:** Select this option to improve performance when loading data into the MySQL target database.
 - **Use the following number of threads:** Specify how many threads to use to load the data into the MySQL target database. Note that setting a large number of threads may have an adverse effect on database performance since a separate connection is required for each thread.
- **Override connection string parameters:** Select this to create a custom connect string.

The following is an example of a MySQL connection string:

```
server=MySQL1;port=5421;database=MyDB;username=Admin;
password=7yyfhn85;
```

For a description of the supported connection parameters, see [Configuration Properties for the MySQL Source](#).

Using ODBC to Connect to a Source

This chapter describes how to use ODBC connectivity to connect to a source database. It contains the following sections:

- [Prerequisites](#)
- [Limitations](#)
- [Using ODBC to Connect to a Source](#)

Prerequisites

The following section describes the prerequisites for working with Amazon RDS Migration Tool and an ODBC endpoint.

You can connect a database to Amazon RDS Migration Tool using ODBC by indicating the DSN (Data Source Name). In this case you must be sure that a DSN is defined for the ODBC database on the computer where Amazon RDS Migration Tool is installed.

1. Install a database client on the computer where Amazon RDS Migration Tool is installed. The client you install depends on the ODBC provider you are using. For example, if you are using an IBM DB2 database, install an IBM DB2 client.

Note: You must use a 64-bit ODBC provider client to work with Amazon RDS Migration Tool.

2. Use the ODBC Data Source Administrator to create a System DSN. The Data Source is located in the Windows control panel.

Limitations

When using ODBC as a source, the following limitations apply:

- UPDATES to primary key fields are not supported. To update the field, define it as a unique index instead.
- The ODBC Source endpoint supports full-load operations only.
- For providers that do not support batch operations, you must manually add `RowByRow=true` to the connect string in the **Advanced** tab of the Add Database dialog box. See [Configuration Properties for the ODBC Source](#) for an explanation of how to use the `RowByRow` property.

Using ODBC to Connect to a Source

The following topics describe what you need to use an ODBC database as a source endpoint in an Amazon RDS Migration Tool task.

- [Configuration Properties for the ODBC Source](#)
- [ODBC Source Data Types](#)
- [Configuring ODBC Databases to work as an Amazon RDS Migration Tool Source](#)
- [Using Advanced Properties when Using ODBC Databases as a Source](#)

Configuration Properties for the ODBC Source

The following table describes the configuration properties available when using ODBC databases as an Amazon RDS Migration Tool source database. For information on how to set these properties, see [Using Advanced Properties when Using ODBC Databases as a Source](#).

Table 9–1 ODBC Source Connect String Standard Properties

Property	Description/Format	Default Value
RowByRow	<p>When set to <code>true</code>, all DML queries are executed row-by-row and not with array binding.</p> <p>Possible values: <code>true</code>, <code>false</code>.</p> <p>This value must be set manually using the Override connection string parameters option in the Advanced tab.</p>	false
BulkArraySize	<p>Set the array size for all DML queries (When using array binding).</p> <p>Possible values: Any positive number.</p> <p>This value must be set manually using the Override connection string parameters option in the Advanced tab.</p>	1000
GetTableColumns	<p>Set the mode for retrieving columns information.</p> <p>Possible values:</p> <ul style="list-style-type: none"> ■ <code>Columns</code>: For using the <code>SQLColumns</code> API. ■ <code>DescribeCol</code>: For using the <code>SQLDescribeCol</code> API. ■ <code>DescribeParam</code>: For using the <code>SQLDescribeCol</code> API (for parameters only). <p>This value must be set manually using the Override connection string parameters option in the Advanced tab.</p> <p>For information on how to use this option, see Using Advanced Properties when Using ODBC Databases as a Target.</p>	Columns
username	The name of the user authorized to access the ODBC target.	

Table 9–1 (Cont.) ODBC Source Connect String Standard Properties

Property	Description/Format	Default Value
password	The password for the authorized user. Note that if you specify a password in your connection string, it will be revealed as plain text in the task log files. It is therefore recommended to specify the password in the GUI Password field.	
syntax	Specify either <code>syntax=DB2</code> , <code>syntax=GenericWithoutSchema</code> or <code>syntax=SQLMP (AIS)</code> as appropriate.	
additionalConnectionProperties	The method for connecting to the ODBC-supported database. This can either be a connection string or a specific DSN. Specify: <code>additionalConnectionProperties={my_connection_string}</code> to connect using a connection string. -OR- <code>additionalConnectionProperties={{DSN=MyDSN}}</code> to connect to an ODBC-supported database using a DSN.	

ODBC Source Data Types

The following table shows the ODBC target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped in the target, see the section for the target endpoint you are using.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 9–2 Supported ODBC Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

ODBC Data Types	Amazon RDS Migration Tool Data Types
SQL_BIT	BOOLEAN
SQL_TINYINT	INT1 UINT1 Note: SQL data types are mapped to unsigned data types when the <code>UNSIGNED_ATTRIBUTE</code> is set to <code>SQL_TRUE</code> for the data type being mapped.
SQL_SMALLINT	INT2 UINT2 Note: SQL data types are mapped to unsigned data types when the <code>UNSIGNED_ATTRIBUTE</code> is set to <code>SQL_TRUE</code> for the data type being mapped.

Table 9–2 (Cont.) Supported ODBC Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

ODBC Data Types	Amazon RDS Migration Tool Data Types
SQL_INTEGER	<p>INT4</p> <p>UINT4</p> <p>Note: SQL data types are mapped to unsigned data types when the UNSIGNED_ATTRIBUTE is set to SQL_TRUE for the data type being mapped.</p>
SQL_BIGINT	<p>INT8</p> <p>UINT8</p> <p>Note: SQL data types are mapped to unsigned data types when the UNSIGNED_ATTRIBUTE is set to SQL_TRUE for the data type being mapped.</p>
SQL_DOUBLE	REAL8
SQL_FLOAT	REAL8
SQL_REAL	REAL8
SQL_NUMERIC (P,S)	<p>NUMERIC (P,S)</p> <p>REAL8</p> <p>The SQL_NUMERIC data type is mapped to REAL8 when <i>at least one</i> of the following is true:</p> <ul style="list-style-type: none"> ■ Precision > 38 ■ Scale < 0 ■ Scale > 38 ■ Scale > Precision
SQL_DECIMAL (P,S)	<p>NUMERIC (P,S)</p> <p>REAL 8</p> <p>The SQL_NUMERIC data type is mapped to REAL8 when <i>at least one</i> of the following is true:</p> <ul style="list-style-type: none"> ■ Precision > 38 ■ Scale < 0 ■ Scale > 38 ■ Scale > Precision
SQL_DATE	DATE
SQL_TYPE_DATE	
SQL_TIME	TIME
SQL_TYPE_TIME	
SQL_TIMESTAMP	DATETIME
SQL_TYPE_TIMESTAMP	
SQL_CHAR	STRING
SQL_VARCHAR	
SQL_WCHAR	WSTRING
SQL_WVARCHAR	

Table 9–2 (Cont.) Supported ODBC Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

ODBC Data Types	Amazon RDS Migration Tool Data Types
SQL_LONGVARCHAR	CLOB
<p>To use this data type with Amazon RDS Migration Tool, you must enable the use of CLOBs for a specific task.</p> <p>During CDC, CLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	
SQL_WLONGVARCHAR	NCLOB
<p>To use this data type with Amazon RDS Migration Tool, you must enable the use of NCLOBs for a specific task.</p> <p>During CDC, NCLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	
SQL_BINARY	BYTES
SQL_LONGVARBINARY	BLOB
<p>To use this data type with Amazon RDS Migration Tool, you must enable the use of BLOBs for a specific task.</p> <p>BLOB data types are supported only in tables that include a primary key.</p> <p>For more information, see LOB support in Task Settings/Metadata.</p>	
SQL_GUID	STRING
SQL_INTERVAL_YEAR	STRING _w
SQL_INTERVAL_MONTH	
SQL_INTERVAL_DAY	
SQL_INTERVAL_MINUTE	
SQL_INTERVAL_HOUR	
SQL_INTERVAL_SECOND	
SQL_INTERVAL_YEAR_TO_MONTH	
SQL_INTERVAL_DAY_TO_HOUR	
SQL_INTERVAL_DAY_TO_MINUTE	
SQL_INTERVAL_DAY_TO_SECOND	
SQL_INTERVAL_HOUR_TO_MINUTE	
SQL_INTERVAL_HOUR_TO_SECOND	
SQL_INTERVAL_MINUTE_TO_SECOND	

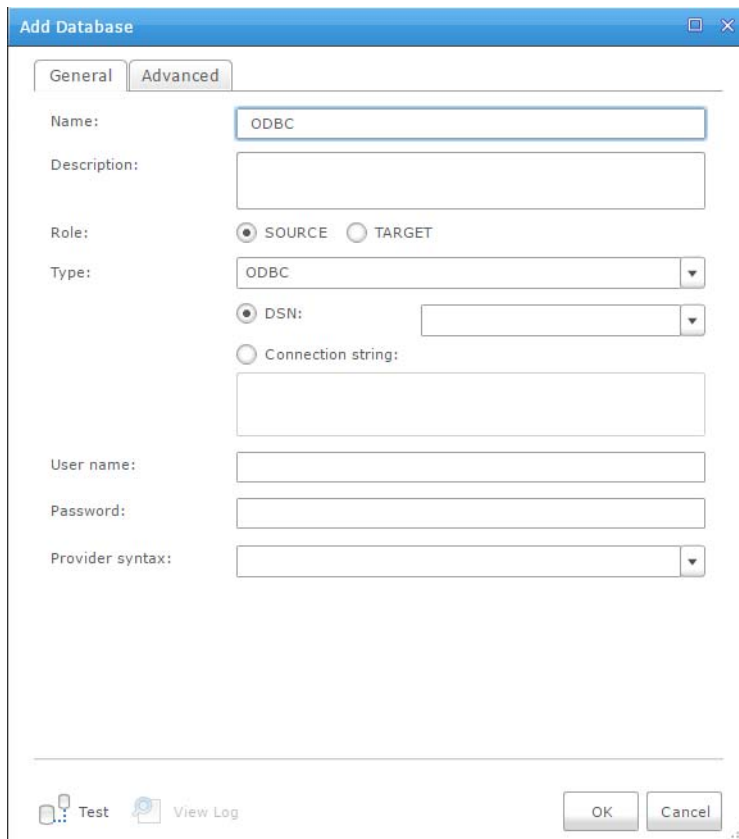
Table 9–2 (Cont.) Supported ODBC Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

ODBC Data Types	Amazon RDS Migration Tool Data Types
Provider specific data types	If column length is < or = 4000:
Note: If column length is 0 or > 4000 then:	BYTES
To use this data type with Amazon RDS Migration Tool, you must enable the use of BLOBs for a specific task.	If column length is 0 or > 4000:
BLOB data types are supported only in tables that include a primary key.	BLOB
For more information, see LOB support in Task Settings/Metadata .	

Configuring ODBC Databases to work as an Amazon RDS Migration Tool Source

You can add a database to Amazon RDS Migration Tool as a source using ODBC connectivity. For information on how to add databases, see [Working with Databases](#). When you select **SOURCE** as the role and then **ODBC** as the database type the following dialog box is displayed:

Figure 9–1 ODBC Source Database



To add an ODBC source database to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).

2. In the **Name** field, type a name for your ODBC database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the ODBC database. This is optional.
4. Select **SOURCE** as the database **Role**.
5. Select **ODBC** as the database **Type**.

You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database Role.

6. Select *one* of the following:
 - **DSN**: Select this to connect to an ODBC-supported database using a DSN. When you select DSN you must select the DSN you are using from the list.

If the DSN you want to use is not included in the list, make sure that the database client is installed on the computer with Amazon RDS Migration Tool and that the DSN is defined. Note that the ODBC provider client must be 64-bit. For more information, see [Prerequisites](#).

Note: If you are using an AIS CDC Agent as the source in a Amazon RDS Migration Tool task, you cannot select the DSN for the Amazon ODBC driver as the target. In this case, to use Amazon ODBC as a source, you must enter the connection string manually by selecting **Connection String** and following the directions for that option in this procedure.

- **Connection String**: Select this to connect to an ODBC-supported database using a connection string then type a valid connection string in the field below. For information on how to create a connection string, see the documentation for the ODBC database provider you are using.

Note that if you specify a password in your connection string, it will be revealed as plain text in the task log files. It is therefore recommended to specify the password in the GUI **Password** field.

Note:

- You can use the **Advanced** tab to add specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties when Using ODBC Databases as a Source](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

7. Type the authentication information (**User Name**, **Password**) for the authorized user for the ODBC database being used. For example, the IBM DB2 system administrator if you are using a IBM DB2 provider. If you do not know this information, see your ODBC Database System Administrator.

Notes:

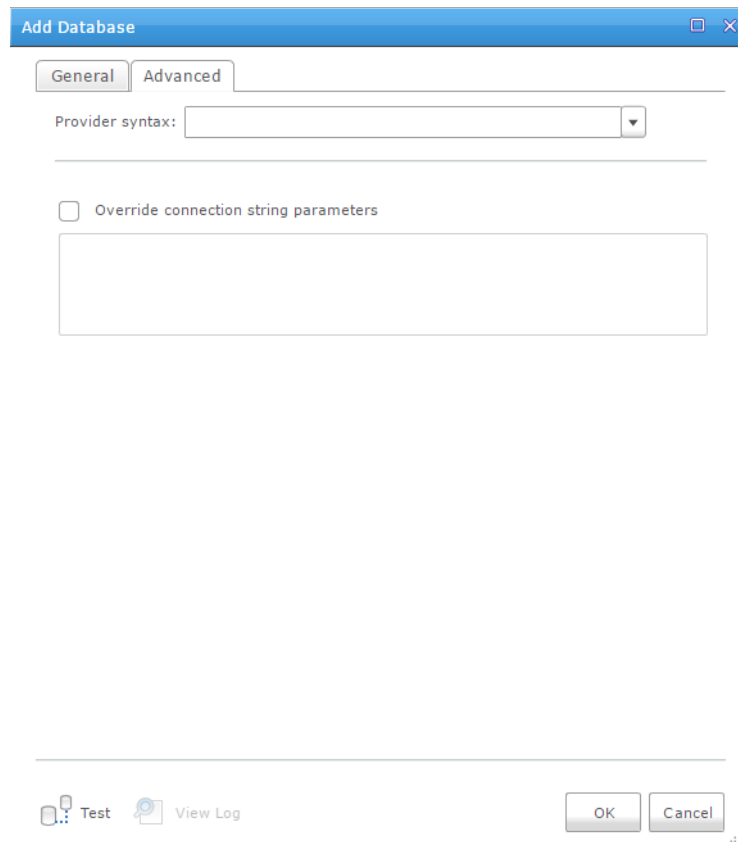
- When you select **Connection String** be sure to include **User name/password** information in the connection string that you type in the box.

If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** and **Password** properties. For more information, see [Using Advanced Properties when Using ODBC Databases as a Source](#).
 - This information is case sensitive.
 - **Important:** Make sure that the ODBC database user has the correct access privileges for the ODBC provider being used. For information on how to provide the required privileges, see.
 - You can set custom properties in the **Advanced** tab. For more information, see [Using Advanced Properties when Using ODBC Databases as a Source](#).
-
-

Using Advanced Properties when Using ODBC Databases as a Source

You can create a custom string using the **Advanced** tab. The following figure shows the ODBC **Advanced** tab:

Figure 9–2 ODBC Source Advanced Tab



Provider syntax: Select the name of the provider syntax if you are using an alternate provider syntax.

Override connection string parameters: Select this to use a custom connect string. The following is an example of a string you can enter in this dialog box:

```
DSN=PERSON;uid=PP1
```

For information on the properties you can use in the connect string, see [Configuration Properties for the ODBC Source](#).

Using Amazon Redshift as a Target

This chapter describes how to set up and use Amazon Redshift as a target in a replication task. Amazon Redshift is located in the cloud and is accessed through an Amazon Web Services (AWS) account.

This chapter contains the following topics:

- [Introducing the Amazon Redshift Target Endpoint](#)
- [Amazon Redshift Endpoint Prerequisites](#)
- [Amazon Redshift Data Types](#)
- [Setting up Amazon Redshift as a Target in Amazon RDS Migration Tool](#)

Introducing the Amazon Redshift Target Endpoint

The Amazon Redshift endpoint leverages the Amazon RDS Migration Transfer Service to convert the data from the source database tables into CSV files, which are then uploaded to an Amazon S3 bucket. The files are then copied from the Amazon S3 bucket to the proper tables in Amazon Redshift.

The Amazon Redshift endpoint provides full automation for:

- Schema generation and data type mapping
- Full load of source database tables
- Incremental load of changes made to source tables
- Application of schema changes (DDL) made to the source tables.
- Synchronization between full load and CDC processes.

Manual control is also available if needed.

Amazon Redshift Endpoint Prerequisites

The following sections describe the prerequisites necessary for working with the Amazon Redshift endpoint.

- [Get Started with Amazon Redshift](#)
- [Sign up for an Amazon S3 Bucket](#)
- [Open the Required Firewall Ports](#)

Get Started with Amazon Redshift

The following section describes how to get started using Amazon Redshift as an Amazon RDS Migration Tool target endpoint.

- Sign up for an Amazon Web Services account. Then use the AWS Management Console to launch an Amazon Redshift cluster. You should note the basic information about your AWS account and your Amazon Redshift cluster, such as your password and user name. You will need this information to configure Amazon RDS Migration Tool to work with the Amazon Redshift endpoint. For more information, see [Setting up Amazon Redshift as a Target in Amazon RDS Migration Tool](#).
- Download and install the appropriate Amazon Redshift ODBC Driver. Amazon RDS Migration Tool requires a 64-bit driver.

For a list of drivers supported by Amazon Redshift, see:

<http://docs.aws.amazon.com/redshift/latest/mgmt/configure-odbc-connection.html>

- Amazon RDS Migration Tool **Server for Windows:** Download and install the Amazon Redshift (x64) driver for Windows.
- Amazon RDS Migration Tool **Server for Linux:** To avoid conflicts, Amazon RDS Migration Tool must be installed *before* you install the driver. Install the Amazon Redshift ODBC driver with `--force` in the command, as in the following example:

```
rpm -ivh AmazonRedshiftODBC-64bit-1.2.6.1006-1.x86_64.rpm --force
```

Once the driver is installed, edit the `amazon.redshiftdbc.ini` file as follows:

- `DriverManagerEncoding=UTF-16`
- `ODBCInstLib=libodbcinst.so`

For information on signing up for an Amazon Web Services account, launching an Amazon Redshift cluster, and installing the client tools, see the Amazon Redshift Getting Started page at <http://docs.aws.amazon.com>.

Sign up for an Amazon S3 Bucket

You need to have an Amazon S3 bucket, preferably (for best performance) located in your Amazon Redshift cluster region.

You must be able to access your Amazon S3 bucket directly from the Amazon RDS Migration Tool machine.

For information on signing up for Amazon S3, see <http://aws.amazon.com/s3/>.

- **Bucket access credentials:** Make a note of the bucket name, region, access key and secret access key - you will need to provide them in the Amazon RDS Migration Tool Amazon Redshift target settings.
- **Bucket access permissions:** Amazon RDS Migration Tool requires read/write/delete permissions to the Amazon S3 bucket.

Open the Required Firewall Ports

Firewall ports 5746 (Amazon RDS Migration Transfer Service) and 5439 (Amazon Redshift Cluster) need to be opened for outbound communication.

Amazon Redshift Data Types

The Amazon Redshift endpoint for Amazon RDS Migration Tool supports most Amazon Redshift data types. The following table shows the Amazon Redshift target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types.

For information on how to view the data type that is mapped from the source, see the section for the source endpoint you are using. For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 10–1 Supported Amazon Redshift Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	Amazon Redshift Data Types
BOOLEAN	BOOL
BYTES	VARCHAR (Length)
DATE	DATE
TIME	VARCHAR(20)
DATETIME	If scale is => 0 and =< 6, then: TIMESTAMP (s) If scale is => 7 and =< 9, then: VARCHAR (37)
INT1	INT2
INT2	INT2
INT4	INT4
INT8	INT8
NUMERIC	If scale is => 0 and =< 37, then: NUMERIC (p,s) If scale is => 38 and =< 127, then: VARCHAR (Length)
REAL4	FLOAT4
REAL8	FLOAT8
STRING	If length is => 1 and =< 65535, then: VARCHAR (Length in Bytes) If length is => 65536 and =< 2147483647, then: VARCHAR (65535)
UINT1	INT2
UINT2	INT2
UINT4	INT4
UINT8	NUMERIC (20,0)

Table 10–1 (Cont.) Supported Amazon Redshift Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	Amazon Redshift Data Types
WSTRING	<p>If length is => 1 and =< 65535, then: NVARCHAR (Length in Bytes)</p> <p>If length is => 65536 and =< 2147483647, then: NVARCHAR (65535)</p>
	<p>Note about Amazon Redshift LOB support: Full LOB data types are not supported. For information on including Limited-size LOB data types in the replication, see the Metadata tab description in Customizing Tasks.</p>
BLOB	<p>VARCHAR (Max LOB Size *2)</p> <p>Note: The maximum LOB size in the Metadata tab cannot exceed 31 KB.</p>
NCLOB	<p>NVARCHAR (Max LOB Size)</p> <p>Note: The maximum LOB size in the Metadata tab cannot exceed 63 KB.</p>
CLOB	<p>VARCHAR (Max LOB Size)</p> <p>Note: The maximum LOB size in the Metadata tab cannot exceed 63 KB.</p>

Setting up Amazon Redshift as a Target in Amazon RDS Migration Tool

You can add Amazon Redshift to Amazon RDS Migration Tool to use as a target database. For information on how to add databases, see [Working with Databases](#).

To add an Amazon Redshift Target to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Manage Databases** to open the **Manage Databases** dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Manage Databases** dialog box, click **Add Database**.
3. In the **Name** field, type a name for your Amazon Redshift data warehouse [service]. This can be any name that will help to identify your Amazon Redshift endpoint.
4. In the **Description** field, type a description that helps to identify the Amazon Redshift target endpoint. This is optional.
5. Select **TARGET** as the **Role**.

You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the **Role**.

6. Select **Amazon Redshift** as the **Type**.
7. Enter the following **Amazon Redshift target** information (see [Figure 0-2](#)).
 - **Redshift cluster:** Enter the name of the Amazon Redshift cluster you are using.
 - **Port:** Enter the port number for Amazon Redshift.
 - **User name:** Enter the user name of a registered Amazon Redshift user.
 - **Password:** Enter the password for the user entered in the **User name** field.
 - **Database name:** Enter the **Database name** or select one from the list of available Amazon Redshift data warehouse [services].

The information for these properties is available from the account page for Amazon Web Services (AWS) with the Amazon Redshift cluster. If you do not have these values, refer to your AWS account or the Amazon Redshift System Administrator for your enterprise.

8. Enter the following **Amazon S3 staging** information. You may need to click the **Amazon S3 staging** header to see the information.
 - **Bucket name:** Type the name of the Amazon S3 bucket where you are copying files to.
 - **Bucket region:** Select the Amazon S3 region where the S3 buckets and folders you are using are hosted. The default value is **US East (N. Virginia)**.

Note: The bucket region specified must be the same region where your Amazon Redshift database is located.
 - **Access key:** Type the access key information for Amazon S3.
 - **Secret key:** Type the secret key information for Amazon S3.
 - **Folder:** Type or browse to the S3 folder where you are copying files to.

The information for these properties is available from your Amazon Web Services (AWS) account. If you do not have these values, refer to your AWS account or the Amazon Redshift System Administrator for your enterprise

Notes:

- If you are using the **Advanced** tab to create a custom string, make sure to include the **User Name** property. A **Password** can also be included but is not required. See [Using Advanced Properties for an Amazon Redshift Target](#) for more information.
- This information is case sensitive.
- If you want to set custom properties, see [Using Advanced Properties for an Amazon Redshift Target](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

Using Advanced Properties for an Amazon Redshift Target

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

You can set the following parameters:

- **Max file size (MB):** Select or type the maximum size of any CSV file used to transfer data to Amazon Redshift. The default value is 1024.
- **Number of threads used to upload a file:** Select the number of threads used to upload a single file. The minimum number of threads is 1. The maximum value is 64. The default value is 10.
- **ODBC driver:** The name of the ODBC driver you are using to connect to Amazon Redshift. The default driver is **Amazon Redshift (x64)**. If you are using a different driver, you must specify its name here.
- **Additional ODBC connection properties:** Type any additional ODBC connection properties if required
- **Override connection string parameters:** Select this to create a custom connect string.

The following is an example of an Amazon Redshift connection string:

```
server=server_name;maxFileSize=499;username=name;database=ere;  
port=5439;
```

Using PostgreSQL Database as a Source or Target

This chapter describes how to set up and use PostgreSQL database as a source or target in a replication task. It has the following topics:

- [Using PostgreSQL Database as a Source](#)
- [Source Prerequisites](#)
- [Security Requirements](#)
- [PostgreSQL Source Data Types](#)
- [Setting up PostgreSQL database as a Source in Amazon RDS Migration Tool](#)
- [Using PostgreSQL as a Target](#)
- [Target Prerequisites](#)
- [Security Requirements](#)
- [PostgreSQL Database Target Data Types](#)
- [Setting up PostgreSQL database as a Target in Amazon RDS Migration Tool](#)

Using PostgreSQL Database as a Source

The following topics describe what you need to use a PostgreSQL database as a source endpoint in an Amazon RDS Migration Tool task.

- [Source Prerequisites](#)
- [PostgreSQL Source Data Types](#)
- [Setting up PostgreSQL database as a Source in Amazon RDS Migration Tool](#)

Source Prerequisites

The following section lists the prerequisites for working with Amazon RDS Migration Tool and a PostgreSQL database source.

Client Side

- **Amazon RDS Migration Server for Windows:** The PostgreSQL ODBC Driver `psqlodbc_09_03_0300-x64-1` or above must be installed on the Amazon RDS Migration Tool machine.
- **Amazon RDS Migration Server for Linux:** On the Amazon RDS Migration Tool machine:

1. Install the ODBC driver `postgresql94-odbc-09.03.0400-1PGDG.<OS version>.x86_64` or above for Linux, where `<OS version>` is the OS of the Amazon RDS Migration Server machine.

For example, `postgresql94-odbc-09.03.0400-1PGDG.<rhel7>.x86_64` is the client required for Red Hat 7.

2. Make sure that the `/etc/odbcinst.ini` file contains an entry for PostgreSQL, as in the following example:

```
[PostgreSQL]
Description = PostgreSQL ODBC driver
Driver = /usr/pgsql-9.4/lib/psqlodbc.so
Setup = /usr/pgsql-9.4/lib/psqlodbcw.so
Debug = 0
CommLog = 1
UsageCount = 2
```

- When the **Apply Changes** task option is enabled, the user specified in the PostgreSQL source endpoint's **General** tab must be granted super-user permissions.

Server Side

- The IP address of the Amazon RDS Migration Tool machine must be added to the `pg_hba.conf` configuration file.
- The following parameters and values must be set in the `postgresql.conf` configuration file.

```
wal_level = logical

max_replication_slots >=1
```

The `max_replication_slots` value should be set according to the number of tasks that you want to run. For example, to run 5 tasks you need to set a minimum of 5 slots. Slots open automatically as soon as a task starts and remain open, even when task is no longer running. Note that open slots need to be manually deleted.

```
max_wal_senders >=1
```

The `max_wal_senders` parameter sets the number of concurrent tasks that can run.

```
wal_sender_timeout =0
```

The `wal_sender_timeout` parameter terminates replication connections that are inactive longer than the specified number of milliseconds. Although the default is

60 seconds, it is recommended to set this parameter to zero which will disable the timeout mechanism.

For more information on the configuration parameters, see:

<http://www.postgresql.org/docs/9.4/static/runtime-config-replication.html>

Security Requirements

The user specified in the **General** tab when [Setting up PostgreSQL database as a Source in Amazon RDS Migration Tool](#) must be a registered user in the PostgreSQL database.

Source Limitations

The following Change Processing limitations apply when using PostgreSQL as a source:

- A captured table must have a Primary Key. In the event that a table does not have a Primary Key, DELETE record operations will be ignored.
- Updating a Primary Key segment is ignored. In such cases, applying such an update will be identified by the target as an update that did not update any rows and will result in a record written to the exceptions table.
- The “Start Process Changes from Timestamp” run option is not supported.
- Change processing is not supported on Amazon RDS for PostgreSQL.
- Replication of multiple tables with the same name but a different case (e.g. table1, TABLE1 and Table1) may cause unpredictable behavior and is therefore not supported.
- Change processing of [CREATE | ALTER | DROP] table DDLs are supported unless they are held in an inner function/procedure body block or in other nested constructs.

For example, the following change will not be captured:

```
CREATE OR REPLACE FUNCTION aws.create_distributors1() RETURNS void
LANGUAGE plpgsql
AS $$
BEGIN
    create table aws.distributors1(did    serial PRIMARY KEY,name
varchar(40) NOT NULL);
END;
$$;
```

- Change processing of TRUNCATE operations is not supported.
- Replication of partitioned tables is not supported. When a partitioned table is detected, the following occurs:
 - a. The endpoint will report a list of parent and child tables.
 - b. The table will be created on the target as a regular table with the same properties as the selected tables.
 - c. If the parent table in the source database has the same Primary Key value as its child tables, a “duplicate key” error will be generated.

Note: In order to replicate partitioned tables from a PostgreSQL source to a PostgreSQL target, you first need to manually create the parent and child tables on the target. Then define a separate task to replicate to those tables. In such a case, the task settings should be configured to “Truncate before loading”. For more information on the “Truncate before loading” option, see [Full Load Settings](#).

PostgreSQL Source Data Types

The following table shows the PostgreSQL target data types that are supported when using Amazon RDS Migration Tool and the default mapping to the Amazon RDS Migration Tool data types.

For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 11–1 Supported PostgreSQL Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

PostgreSQL Data Types	Amazon RDS Migration Tool Data Types
INTEGER	INT4
SMALLINT	INT2
BIGINT	INT8
NUMERIC(P,S)	If precision is => 0 and =< 38, then: NUMERIC If precision is => 39, then: STRING
DECIMAL(P,S)	If precision is => 0 and =< 38, then: NUMERIC If precision is => 39, then: STRING
REAL	REAL4
DOUBLE	REAL8
SMALLSERIAL	INT2
SERIAL	INT4
BIGSERIAL	INT8
MONEY	NUMERIC(38,4) Note: The MONEY data type is mapped to FLOAT in Microsoft SQL Server.
CHAR	WSTRING (1)
CHAR(N)	WSTRING (n)
VARCHAR(N)	WSTRING (n)
TEXT	NCLOB
BYTEA	BLOB
TIMESTAMP	TIMESTAMP

Table 11–1 (Cont.) Supported PostgreSQL Source Data Types with Mapping to Amazon RDS Migration Tool Data Types

PostgreSQL Data Types	Amazon RDS Migration Tool Data Types
TIMESTAMP (z)	TIMESTAMP
DATE	DATE
TIME	TIME
TIME (z)	TIME
INTERVAL	STRING (128) - 1 YEAR, 2 MONTHS, 3 DAYS, 4 HOURS, 5 MINUTES, 6 SECONDS
BOOLEAN	STRING (1) F or T
ENUM	STRING (64)
CIDR	STRING (50)
INET	STRING (50)
MACADDR	STRING (18)
BIT (n)	STRING (n)
BIT VARYING (n)	STRING (n)
UUID	STRING
TSVECTOR	CLOB
TSQUERY	CLOB
XML	CLOB
POINT	STRING (255) "(x,y)"
LINE	STRING (255) "(x,y,z)"
LSEG	STRING (255) "((x1,y1),(x2,y2))"
BOX	STRING (255) "((x1,y1),(x2,y2))"
PATH	CLOB "((x1,y1),(xn,yn))"
POLYGON	CLOB "((x1,y1),(xn,yn))"
CIRCLE	STRING (255) "(x,y,r)"
JSON	NCLOB
ARRAY	NCLOB
COMPOSITE	NCLOB
INT4RANGE	STRING (255)
INT8RANGE	STRING (255)
NUMRANGE	STRING (255)
STRRANGE	STRING (255)

Setting up PostgreSQL database as a Source in Amazon RDS Migration Tool

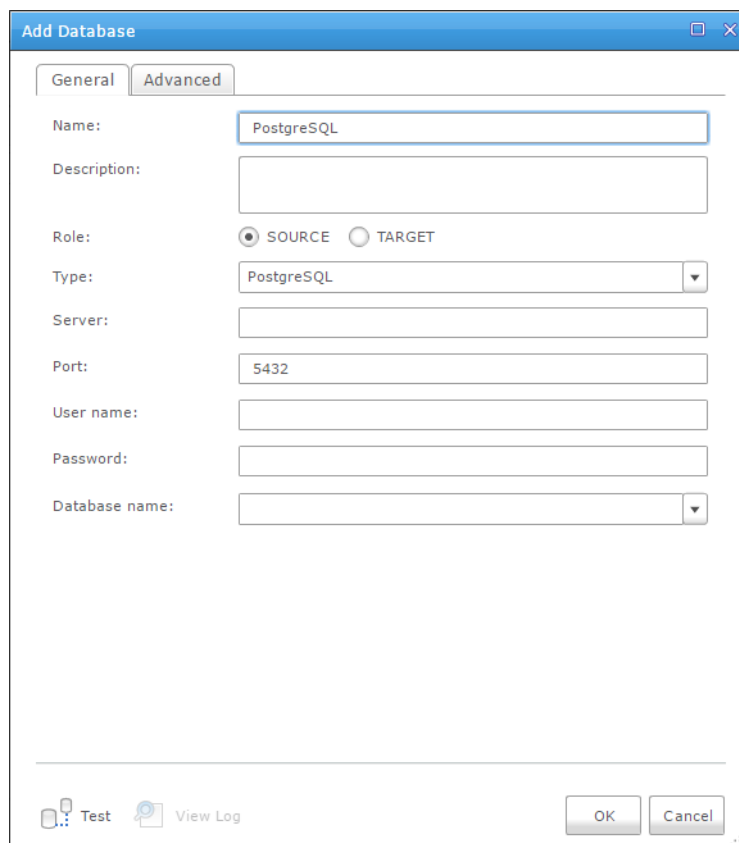
You can add an PostgreSQL source endpoint to Amazon RDS Migration Tool, which can then be used in a replication task.

To add a PostgreSQL database source endpoint to Amazon RDS Migration Tool

1. In the Amazon RDS Migration Tool console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your PostgreSQL database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the PostgreSQL database. This is optional.
4. Select **SOURCE** as the database **Role**.
5. Select **PostgreSQL** as the database **Type**.

The following settings are displayed.

Figure 11–1 PostgreSQL Source Database - General Tab



6. Type the **Server** name. This is the name or IP address of the computer with the PostgreSQL database that you want to access.
7. Optionally, change the default port (5432).
8. Enter the PostgreSQL database authentication information (**User name**, **Password**) of an authorized PostgreSQL user. If you do not know this information, see your PostgreSQL database system manager.

Notes:

- This information is case sensitive.
- **Important:** Make sure that the PostgreSQL database user entered in the PostgreSQL database Authentication section has the correct access privileges.
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

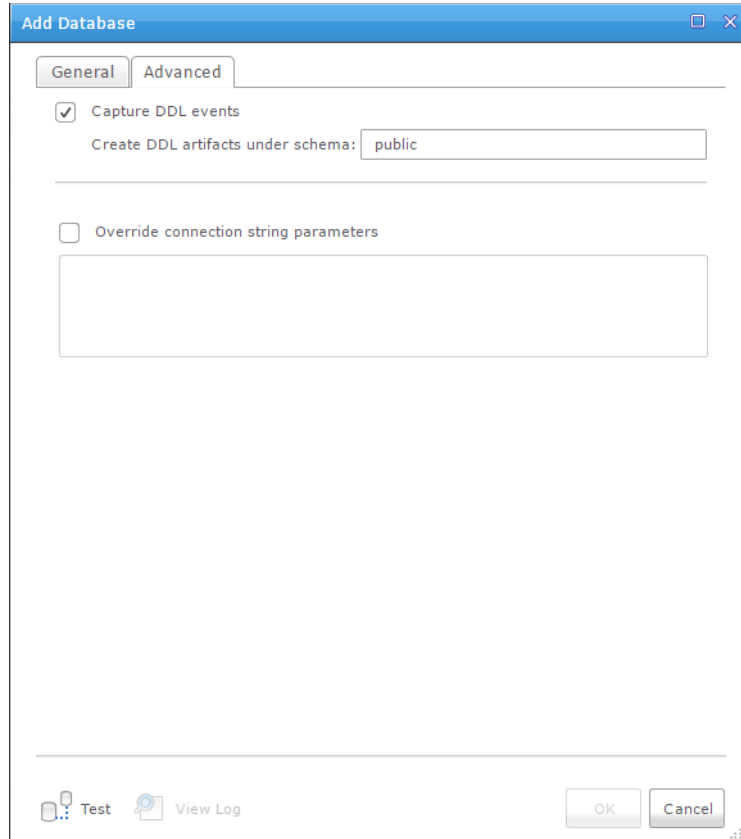
To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

9. Type the **Database name** or select one from the list of available databases. This is the name of the PostgreSQL database from which you are replicating data.
10. Click **OK** to save your settings and close the dialog box.

Using Advanced Properties for a PostgreSQL Source

In the **Advanced** tab of the **Add Database** dialog box, you can enable capture of DDL events and override the connection parameters specified in the dialog box.

Figure 11–2 PostgreSQL Source Advanced Tab



- **Capture DDLs** - When this option is selected, the following actions occur:
 - Operational artifacts are created (by RDS Migration Tool) in the database when the task starts. In order to capture DDL events, Amazon RDS Migration Tool creates various artifacts in the PostgreSQL database when the task starts. You can later remove these artifacts as described in [Removing RDS Migration Tool Artifacts from the Source Database](#).
 - Streamed DDL events are captured.
- **Create DDL artifacts in schema** - The schema in which the operational DDL database artifacts will be created. The default value is "Public".
- **Override connection string parameters:** Select this option if you need to:
 - Specify internal RDS Migration Tool parameters. Such parameters are rarely required and are therefore not exposed in the UI.
 - Specify pass-through (passthru) values for the specific database client. The pass-through parameter will be included in the connection sting passed (by RDS Migration Tool) to the database client.

Note that, apart from the password (which is *never* revealed in plain text), any parameters already set in the UI will also be displayed in the edit box.

Syntax:

```
username=<user>;database=<name>;server=<IP address or hostname>;<internal_
parameter>=<value>;<passthru_parameter>=<value>
```

Removing RDS Migration Tool Artifacts from the Source Database

In order to capture DDLs, Amazon RDS Migration Tool creates various artifacts in the PostgreSQL database when the task starts. When the task completes, you may wish to remove these artifacts.

To remove the artifacts, issue the following statements (in the order they appear below), where `public` is the default schema in which the artifacts were created:

```
drop event trigger amazon_intercept_ddl;
```

Note that the event `trigger` does not belong to a specific schema.

```
drop function public.amazon_intercept_ddl()
```

```
drop table public.amazon_ddl_audit
```

```
drop schema public
```

IMPORTANT: Dropping a schema should be done with extreme caution, if at all. Never drop an operational schema, especially not `public`.

Using PostgreSQL as a Target

This section contains information on how to configure the Amazon RDS Migration Tool PostgreSQL target endpoint.

- [Target Prerequisites](#)
- [Security Requirements](#)
- [Setting up PostgreSQL database as a Target in Amazon RDS Migration Tool](#)

Target Prerequisites

The following section describes the client prerequisites when replicating to a PostgreSQL target.

- **Amazon RDS Migration Server for Windows:**
 - **The PostgreSQL ODBC Driver:** The PostgreSQL ODBC Driver `psqlodbc_09_03_0300-x64-1` or above must be installed on the Amazon RDS Migration Tool machine.
 - **pgAdmin:** The pgAdmin Open Source administration and development platform for PostgreSQL must be installed on the Amazon RDS Migration Tool machine.

- **Amazon RDS Migration Server for Linux:** On the Amazon RDS Migration Tool machine:

1. Install unixODBC driver `postgresql94-odbc-09.03.0400-1PGDG.<OS version>.x86_64` or above for Linux, where `<OS version>` is the OS of the RDS Migration Tool Server machine.

For example, `postgresql94-odbc-09.03.0400-1PGDG.<rhel7>.x86_64` is the client required for Red Hat 7.

2. Makes sure that the `/etc/odbcinst.ini` file contains an entry for PostgreSQL, as in the following example:

```
[PostgreSQL]
Description = PostgreSQL ODBC driver
Driver = /usr/lib/odbc/psqlodbc.so
Setup = /usr/lib/odbc/libodbcpsqlS.so
Debug = 0
CommLog = 1
UsageCount = 2
```

Security Requirements

The user specified in the **General** tab when [Setting up PostgreSQL database as a Target in Amazon RDS Migration Tool](#) must be a registered user in the PostgreSQL database.

PostgreSQL Database Target Data Types

The PostgreSQL database endpoint for Amazon RDS Migration Tool supports most PostgreSQL database data types. The following table shows the PostgreSQL database target data types that are supported when using Amazon RDS Migration Tool and the default mapping from Amazon RDS Migration Tool data types. Unsupported data types are listed below the table.

For information on how to view the data type that is mapped from the source, see the section for the source endpoint you are using. For additional information about Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

Table 11–2 Supported PostgreSQL database Data Types with Mapping from Amazon RDS Migration Tool Data Types

Amazon RDS Migration Tool Data Types	PostgreSQL database Data Types
BOOL	BOOL
BYTES	BYTEA
DATE	DATE
TIME	TIME
TIMESTAMP	If scale is => 0 and =< 6, then: TIMESTAMP If scale is => 7 and =< 9, then: VARCHAR (37)
INT1	SMALLINT
INT2	SMALLINT
INT4	INTEGER
INT8	BIGINT
NUMERIC	DECIMAL (P, S)
REAL4	FLOAT4
REAL8	FLOAT8
STRING	If length is 1 - 21845, then: VARCHAR (Length in Bytes) If length is 21846 - 2147483647, then: VARCHAR (65535)
UINT1	SMALLINT
UINT2	INTEGER
UINT4	BIGINT
UINT8	BIGINT
WSTRING	If length is 1 - 21845, then: VARCHAR (Length in Bytes) If length is 21846 - 2147483647, then: VARCHAR (65535)
BLOB	BYTEA
NCLOB	TEXT
CLOB	TEXT

Data Types when Replicating from a PostgreSQL Source

When replicating from a PostgreSQL source, the target table will be created with the same data types for all columns, apart from columns with user-defined data types. In such cases, the data type will be created as "character varying" in the target.

Setting up PostgreSQL database as a Target in Amazon RDS Migration Tool

You can designate PostgreSQL database as the target endpoint in an Amazon RDS Migration Tool task.

To add a PostgreSQL target database to Amazon RDS Migration Tool:

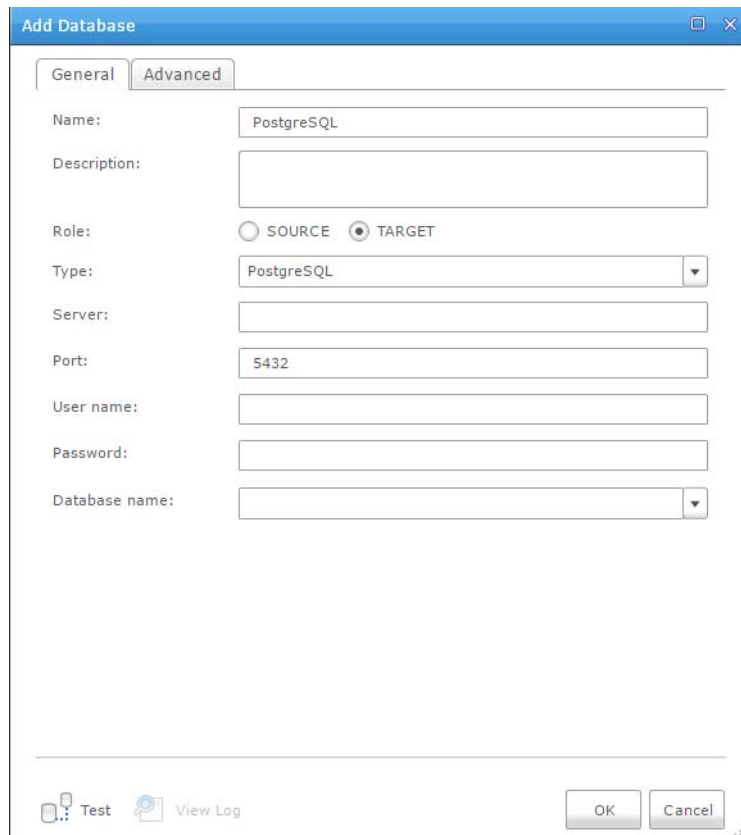
1. In the Amazon RDS Migration Tool console, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
2. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
3. In the **Description** field, type a description that helps to identify the PostgreSQL database. This is optional.
4. Select **TARGET** as the database **Role**.

You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.

5. Select **PostgreSQL** as the database **Type**.

The following dialog box is displayed.

Figure 11–3 PostgreSQL Target Database



6. Type the **Server** name. This is the name or IP address of the computer with the PostgreSQL database that you want to access.

7. Optionally, change the default port (5432).
8. Enter the PostgreSQL database authentication information (**User name**, **Password**) of an authorized PostgreSQL user. If you do not know this information, see your PostgreSQL database system manager.

Notes:

- This information is case sensitive.
- **Important:** Make sure that the specified PostgreSQL database user has the correct access privileges.
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

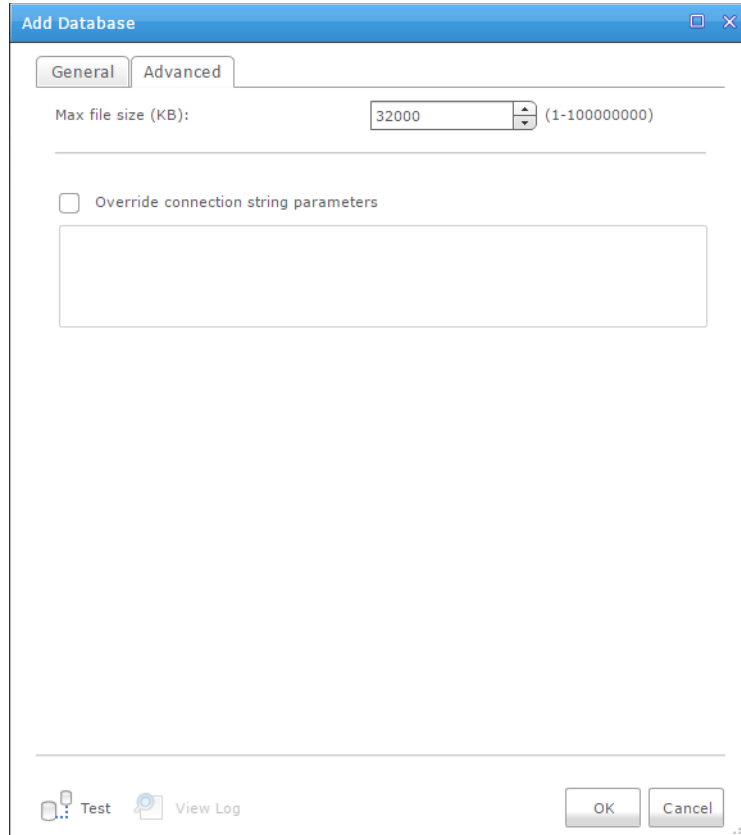
To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

9. Type the **Database name** or select one from the list of available databases. This is the name of the PostgreSQL database to which you are replicating data.
10. Click **OK** to save your settings and close the dialog box.

Using Advanced Properties for a PostgreSQL database Target

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

Figure 11–4 PostgreSQL Target Advanced Tab



You can set the following parameters:

- **Max file size (KB):** Select or type the maximum size (in KB) of a CSV file before the file is loaded into the PostgreSQL target database. The default value is 32000 KB.
- **Override connection string parameters:** Select this option if you need to:
 - Specify internal RDS Migration Tool parameters. Such parameters are rarely required and are therefore not exposed in the UI.
 - Specify pass-through (passthru) values for the specific database client. The pass-through parameter will be included in the connection sting passed (by RDS Migration Tool) to the database client.

Note that, apart from the password (which is *never* revealed in plain text), any parameters already set in the UI will also be displayed in the edit box.

Syntax:

```
username=<user>;database=<name>;server=<IP address or hostname>;<internal_
parameter>=<value>;<passthru_parameter>=<value>
```

Using the Amazon RDS Migration Tool File Channel

This chapter describes how to use the Amazon RDS Migration Tool file channel as a source or target in a replication task.

It covers the following topics:

- [Setting Up Amazon RDS Migration Tool File Channel Tasks](#)
- [Working with the File Channel Data Files](#)
- [Amazon RDS Migration Tool Installation Requirements for the File Channel](#)
- [Security](#)
- [Limitations](#)
- [Using the File Channel as a Source](#)
- [Using the File Channel as a Target](#)

Setting Up Amazon RDS Migration Tool File Channel Tasks

To replicate data using the file channel, you must set up two tasks of the following type:

- [Local Task](#)
- [Remote Task](#)

Note: When using file channel, Change Tables can be enabled for the remote task *but not for the local task* (enabling Change Tables for the local task will result in remote task failure).

Local Task

You set up the local task using the File-Channel endpoint as a target. The binary file created in this task is used as the source for one or more remote tasks using the File-Channel source endpoint.

The local task replicates data from an Amazon RDS Migration Tool supported database to the file channel. If you changed the default folder for storing data files (during the installation), then you must specify the location of the binary file created by the file channel. This location can be anywhere in your system. For more information on setting up a local task, see [Using the File Channel as a Target](#).

Remote Task

Remote tasks use the File Channel as a source endpoint. You use the file created by the local task for this source. You can replicate the data to any database that is supported by Amazon RDS Migration Tool. You define the location of the File-Channel file as the remote location where the file was created. The data is pushed over the network to the defined location anywhere in your system. You can also define more than one location for the replicated data. In this case, define a separate remote task for each location.

If you want to push the data to a database that is not in your LAN, use the [File Transfer Service](#) to send the files created in the local task to the remote location.

When you run the remote task, data is sent to the target in the following instances:

- The first time you run the task as a full load.
- Each time changes are made to the file. In this case, change processing takes place.

When the remote task runs, it will continuously look for the source file until the task is stopped. When the file is found, the data is replicated to the target database. If no source file is found, an error is displayed; however, the task will continue to check for the correct file. Therefore, it is recommended that you run the local task first to ensure that the file exists.

Note: To replicate tables that were added to the local file channel task after the initial full load, you need to reload both the local and the remote file channel tasks.

For more information on setting up a remote task, see [Using the File Channel as a Source](#).

Replicating to Multiple Targets (Distribution)

You can use the File Channel to distribute from a single source database to multiple targets, either of the same type (e.g. Microsoft SQL Server to Microsoft SQL Server) or of different types (e.g. Microsoft SQL Server to Oracle and SAP Sybase ASE).

To do this:

1. For each of the target databases, define a separate (remote) task that replicates from the File Channel source to the target database. In the **Advanced** tab of the File Channel source settings, make sure to clear the **Delete processed files** check box. This ensures that the File Channel files will be available for distribution as required.
2. Define a local task that replicates from the source database to a File Channel target.
3. Run the local task (this will create the File Channel files required by the remote task).
4. For each of the remote tasks, select which tables to replicate (from the File Channel source) and optionally apply Filters and Transformations to them.
5. Run the remote tasks.

For more information on defining tasks, see [Designing Tasks](#).

For information on Filters and Transformations, see [Customizing Tasks](#).

Note: By default, all the metadata for the selected source tables is replicated from the [Local Task](#) to the [Remote Task](#). This allows you to remove, add and transform tables in the remote task as needed. However, if you want the tables in the source and target databases to be identical, you can prevent replication of the metadata (and thereby shorten the processing time) by specifying `providerremotemetadata=N` in the **Override connection string parameters** field of the File Channel target's **Advanced** tab.

Adding Tables to a Running Remote Task

When distributing to multiple targets, it is possible to replicate a different subset of tables to each target if necessary. Before starting the task, you can select which tables to replicate using the standard procedure described in [Adding Tables to a Task](#). However, if the task is already running, you need to perform the following procedure:

1. Stop the remote task.
2. Add the desired tables (as described in [Adding Tables to a Task](#)).
3. Resume the remote task. The newly added tables will be marked as “Queued”.
4. Reload the newly added tables in the local task (by selecting the tables and clicking the **Reload** icon in Monitor view).

For information on removing specific tables from a replication task, see [Removing Specific Tables from a Replication Task](#).

Note: Adding tables to the remote task is not supported in Apply Changes (CDC-only) replication tasks. For more information on the available replication options, see [Setting up Tasks](#).

Working with the File Channel Data Files

The File Channel stream data files are encoded in an internal binary format. For full-load operations, the File Channel binary files contain packed data records for each of the table records and an end-of-file (EOF) record. For change-processing operations, the file contains:

- A packed data record for each DDL and/or DML change.
- A **begin-load-table** record with the stream name that marks the beginning of table loading.
- A packed table-definition record with the table metadata. These records come before each DDL and **begin-load-table** record.

You do not need to work directly with the file-channel files, however if you find it necessary to work with them they are located in the [File-Channel Directory](#).

File-Channel Directory Structure

The file-channel directory contains the following files and folders:

s_msgs: This folder contains messages sent from the source side to the replication server on the remote target side.

Messages are removed from this folder at the source side when an acknowledgment message is received stating that the file was transferred successfully or possibly with a timeout.

Messages are removed from this folder at the target side after they are read.

This folder contains the following files:

- **s_msgs/xxxxxxx.fcm:** This file contains a JSON message from the source side to the target side.
- **yyyymmddhhMMsss.mtd:** This file contains the captured tables list.

s_status: This folder contains status updates from the source side to the target side. Status updates appear as a fixed name file that is periodically updated. This file lists

the last processed target status file. It receives the `t_status/ccccccc.fcs` file. These files are deleted when the file-channel source endpoint finishes reading the file. You can configure the file-channel source to keep the files, if necessary. See [Using Advanced Properties for a File-Channel Source](#) for more information.

t_status: This folder contains status updates from the target side to the source side. Status updates appear as an infinite set of data files that are created according to a specific schedule. These files are sent from the target by the source. The folder contains also a fixed name file that is updated with the last created status file name. It contains the following file:

- **t_status/ccccccc.fcs:** This is a file channel status file (.fcs) where the file name is a hexadecimal counter of length 8. These files will be transferred in order with the lower numbers transferred first. If you need to view them, you should order them by timestamp because alphabetical ordering will not be consistent with the hexadecimal name.

File channel status files are deleted by the source after being read and by the target when source status file indicates that this file was already processed.

You can configure the maximum amount of time that the files are kept before a new file is created as well as the maximum file size for each file. The minimum file size is 50 MB.

For more information, see [Using Advanced Properties for a File-Channel Target](#).

- **streams/<stream-name>:** This folder contains stream sub-folder, one sub-folder per stream. A stream represents a finite or infinite set of data files being sent from the source to the target. The file channel allows creating and destroying named streams dynamically. For example, there can be a fixed-named stream `cdc` (`streams/cdc`) and there could be a dynamically created stream `loadXXXXXXXX` that can be removed at the source side when a status update from the target is received (for example, when processing completed) in the **t_status** folder.

You can configure the maximum number of streams and the maximum disc space for each stream. For more information, see [Change Processing](#).

This folder contains the following file:

- **streams/<stream-name>/ccccccc.fcd:** This is a file channel data file (.fcd) where the file name is a hexadecimal counter of length 8. These files are processed at the target in order or in parallel depending on the case. However, the files are transferred in order with the lower numbers transferred first.

File channel data files are deleted by the source when transferred successfully and by the target when processed.

You can configure the maximum amount of time that the files are kept before being creating a new file and the maximum file size for each file. The minimum file size is 10 MB and the minimum time that a file is kept is 5 seconds.

Amazon RDS Migration Tool Installation Requirements for the File Channel

To work with the file-channel endpoint, you must install Amazon RDS Migration Tool anywhere on the network for each LAN that you are working with.

Security

When using the File Transfer Service, file-channel files are always transferred over an encrypted session.

The session is encrypted as follows:

The client and server create an AES-256 session key using the Diffie-Hellman key exchange protocol (using the OpenSSL library). After the key is created, all file transfers between the client and the server will take place over a secure and encrypted communication channel.

However, even though the session is encrypted, communication between the client and the server may still be susceptible to man-in-the-middle attacks. A man-in-the-middle in possession of the session key would be able to intercept any data transferred between the client and the server.

To eliminate man-in-the-middle attacks, a "shared password" needs to be provided when configuring the local and remote file channel endpoints. Once the session is established, both the client and the server use the shared password to re-key the session key during the next packet exchange, thereby preventing the original session key from being used for man-in-the-middle attacks.

To sum up:

1. Strong encryption is used regardless of whether a password was provided.
2. Providing a password eliminates the risk of a man-in-the-middle attack.

For more information about the File Transfer Service, see [File Transfer Service](#).

Limitations

- The File-Channel endpoint does not support the use of BLOBs.
- You cannot use the full-load resume function if you are using the file-channel endpoint. To resume a full-load operation, you must delete the original data and then run the task again.
- You must delete the file-channel folder before re-starting a task for change processing.
- Control tables defined for the Local File Channel task but not for the Remote File Channel task will not be created on the remote task's target endpoint.

For information on defining Control Tables, see [Control Tables](#).

Using the File Channel as a Source

The File Channel source endpoint is an Amazon RDS Migration Tool endpoint that consumes and applies the contents of a file channel directory structure that was produced by a corresponding File Channel target endpoint.

This section contains the following topics:

- [File Channel Source Properties](#)
- [Setting up a File Channel Source using Amazon RDS Migration Tool](#)

File Channel Source Properties

The following table describes the configuration properties available for the File-Channel source endpoint. For information on how to set these properties, see [Using Advanced Properties for a File-Channel Source](#).

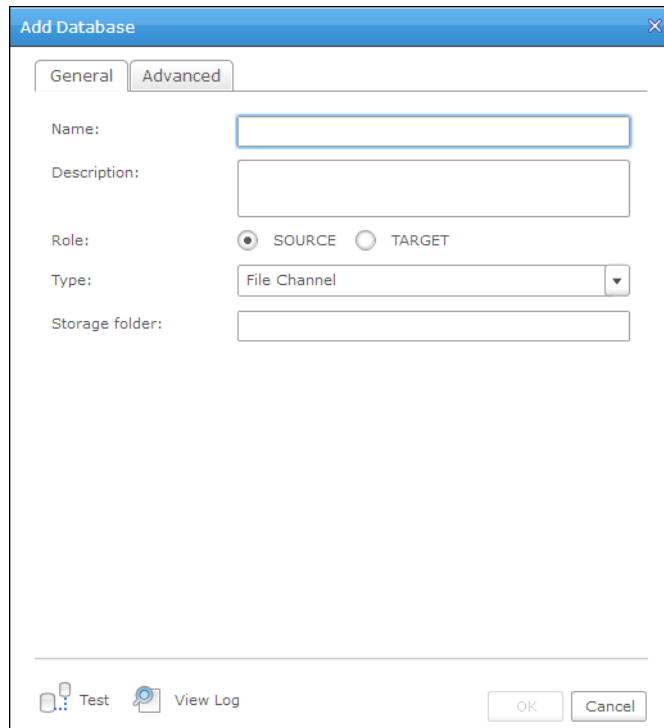
Table 12–1 File-Channel Source Connect String Standard Properties

Property	Description/Format	Default Value
path	The full path to the location where the binary file is saved. It should be in a location that is accessible from anywhere in the WAN you are working with.	
DeleteProcessedFiles	When <code>true</code> , the File-Channel binary file is deleted when the source endpoint finishes reading the file. Set this to <code>false</code> if more than one task must read the same file.	<code>true</code>
ReceiveTransferredFiles	Set to <code>Y</code> to receive the source input files using the RDS Migration Tool File Transfer Service. This value should only be enabled if the File Channel Target is also configured to use the RDS Migration Tool File Transfer Service. For more information on the RDS Migration Tool File Transfer Service, see File Transfer Service .	<code>N</code>

Setting up a File Channel Source using Amazon RDS Migration Tool

You use the File-Channel source as a database in an Amazon RDS Migration Tool task. The Add Database dialog box opens when you click **Add Database** from the Manage Database wizard. For information on how to add databases, see [Working with Databases](#).

Figure 12–1 File-Channel Source Endpoint



To use the File-Channel source endpoint in Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Manage Databases** to open the Manage Databases wizard.
2. In the Manage Databases wizard, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
3. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
4. In the **Description** field, type a description that helps to identify the information being replicated to the file. This is optional.
5. Select **SOURCE** as the database **Role**.

You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.

6. Select **File Channel** as the database **Type**.
7. Type the full path to the **Storage Folder** where the File Channel files will be created. The default path when *not* using the File Transfer Service is:

```
C:\Program Files\Amazon\RDS Migration Tool\data\tasks\.
```

If you are using the File Transfer Service, the default path is:

```
C:\Program Files\Amazon\RDS Migration Tool\data\databases\  
<file-channel_db_name>\fc
```

Note: The RDS Migration Tool File Transfer Service always transfers the local file channel task's files *to the default directory* on the remote system (C:\Program Files\Amazon\RDS Migration Tool\data\databases\
<remote_file-channel_db_name>\fc). Consequently, if you are using the File Transfer Service, ensure that the default directory always has enough space for the incoming files.

For more information on using the File Transfer Service, see [File Transfer Service](#) and [Using Advanced Properties for a File-Channel Source](#).

This folder should be in a location that is accessible from anywhere in the WAN you are working with.

Note:

- You can use the **Advanced** tab to define specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a File-Channel Source](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

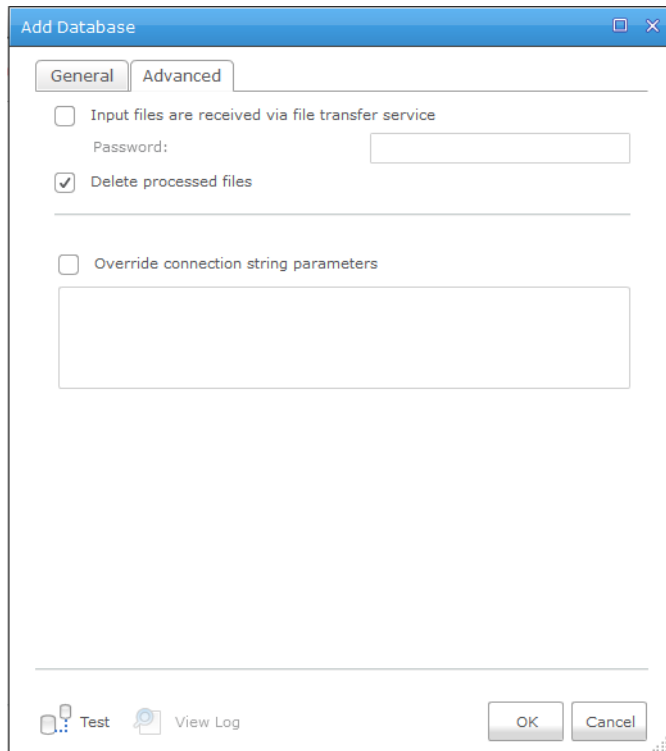
To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

8. Click **OK** to finish the setup and save the changes.

Using Advanced Properties for a File-Channel Source

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

Figure 12–2 File-Channel Source Advanced Tab



You can set the following properties in the **Advanced** tab:

- **Input files are received via file transfer service:** Select this check box to receive the source input files using the RDS Migration Tool File Transfer Service.
 - **Password:** The password that will be used to establish a secure connection with the File Channel Target.

Important: When using the File Transfer Service, an agreed upon password is required in order to establish a secure connection between the File Channel Source and the File Channel Target. Accordingly, the password specified in the File Channel Source settings and the password specified in the File Channel Target settings *must be identical*.

For more information about the File Transfer Service, see [File Transfer Service](#).

- **Delete processed files:** Select this check box to delete the File Channel files after the data has been replicated to the target database.

You should clear this check box if other tasks need to use the files.
- **Override connection string parameters:** Select this to use a custom connect string. The following is an example of a string you can enter in this dialog box:

```
DeleteProcessedFiles=Y;BufferSize=1000;path=\\folder\file_channel;
```

For a description of the supported connection string parameters, see [File Channel Source Properties](#).

Using the File Channel as a Target

The File-Channel target endpoint is an Amazon RDS Migration Tool endpoint that creates and maintains a file-based directory structure containing replication artifacts (task definitions, metadata, full load data, CDC data and status updates). This file channel directory structure is consumed by a corresponding File-Channel source endpoint in a different task and possibly in a remote location.

This section contains the following topics:

- [File Channel Target Properties](#)
- [Setting up a File Channel Target using Amazon RDS Migration Tool](#)

File Channel Target Properties

The following table describes the configuration properties available for the File-Channel target endpoint. For information on how to set these properties, see [Using Advanced Properties for a File-Channel Target](#).

Table 12–2 File Channel Target Connect String Standard Properties

Property	Description/Format	Default Value
path	The full path to the location where the binary file is saved. It should be in a location that is accessible from anywhere in the WAN you are working with.	
MaxFileSize	The maximum file size (in kilobytes) allowed for the files created in the target. You can use a value from 1 - 100000000.	32000
MaxStorageSize	The amount of disk space to set aside for the File Channel files. When the limit is reached, Amazon RDS Migration Tool will stop writing the files to the designated storage. You can use a value from 0 - 1000000. 0 means that there is no limit.	1000000
MaxBatchingTimeInterval	The maximum time (in seconds) for files to be batched before being written in a single operation. You can use a value from 1 - 86400.	120
TransferFiles	Whether to use the RDS Migration Tool File Transfer Service to transfer files to the source File Channel.	N
RemoteDatabases	The information required to access the remote File Channel when TransferFiles=Y. Example: RemoteDatabases=MyDB@123.123.123.1:9090	9191
fileTransferUploadStreams	The maximum number of streams to use to transfer the files when TransferFiles=Y. You can adjust the number of streams to achieve optimal transfer speeds.	10

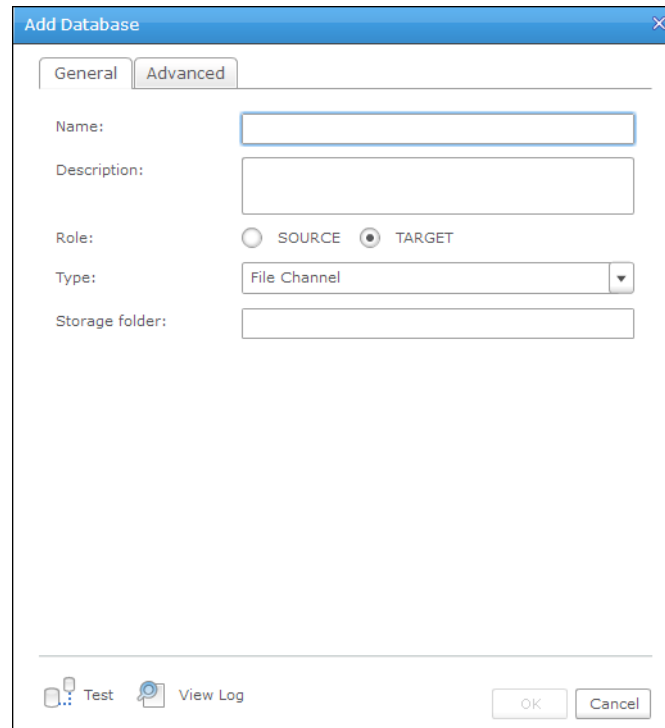
Table 12-2 (Cont.) File Channel Target Connect String Standard Properties

Property	Description/Format	Default Value
providerremotemetadata	Whether to replicate the target File Channel metadata to the source File Channel. If you do not need to add, remove or transform tables in the remote task, you can shorten processing time by specifying providerremotemetadata=N.	Y

Setting up a File Channel Target using Amazon RDS Migration Tool

You use the File-Channel target as a database in an Amazon RDS Migration Tool task. The Add Database dialog box opens when you click **Add Database** from the Manage Database wizard. For information on how to add databases, see [Working with Databases](#).

Figure 12–3 File-Channel Target Endpoint



Notes:

- The **Type** is different depending on the type of file you are creating, however the information you enter is the same for all file types.
 - All files are used as targets, however you can use an Amazon RDS Migration Tool file as a source only after you created the file by loading data into it as a target.
-

To use the File-Channel target endpoint in Amazon RDS Migration Tool

1. In the Amazon RDS Migration Console, click **Manage Databases** to open the Manage Databases wizard.
2. In the Manage Databases wizard, click **Add Database** to open the Add Databases dialog box. For more information on adding a database to Amazon RDS Migration Tool, see [Working with Databases](#).
3. In the **Name** field, type a name for your database. This can be any name that will help to identify the database being used.
4. In the **Description** field, type a description that helps to identify the information being replicated to the file. This is optional.

5. Select **TARGET** as the database **Role**.

You can do this step before any of the other steps if you want, however before you can continue with the next step in this process, you must select the database **Role**.

6. Select **File Channel** as the database **Type**.
7. If you changed the default data folder during installation, type the full path to the **Storage Folder** (e.g. D:\data\tasks\) where the file is being created. Otherwise, you can leave this field empty. Note that this field will be ignored when the **Transfer files to remote file channel** option is enabled in the **Advanced** tab.

Note:

- You can use the **Advanced** tab to define specific properties and create a custom connect string. In this case, you do not need to enter information in this tab. For more information on using the **Advanced** tab, see [Using Advanced Properties for a File-Channel Target](#).
- To determine if you are connected to the database you want to use or if the connection information you entered is correct, click **Test**.

If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

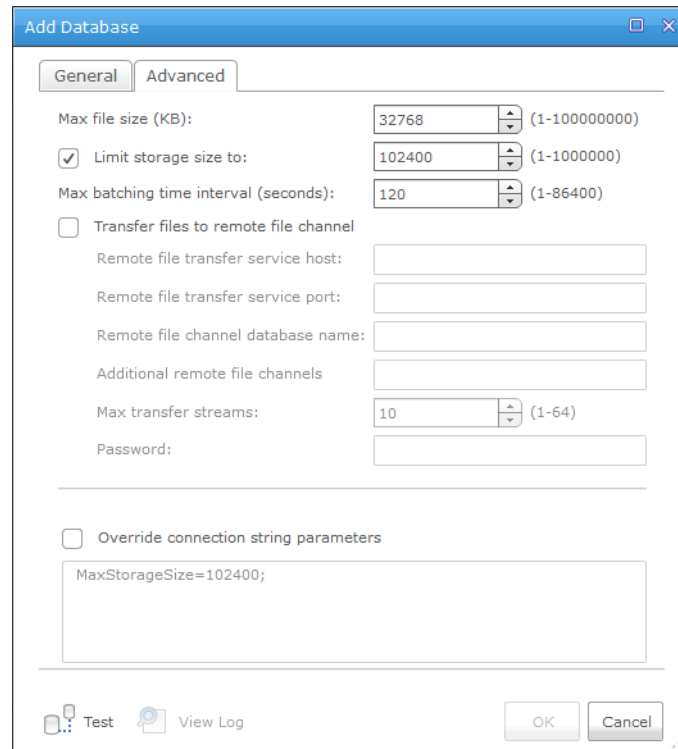
To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

8. Click **OK** to finish the setup and save the changes.

Using Advanced Properties for a File-Channel Target

You can set custom properties or change the default settings for various parameters by adding them to a custom connect string in the **Advanced** tab of the Add Database dialog box.

Figure 12–4 File-Channel Target Advanced Tab



You can set the following properties in the **Advanced** tab:

- **Max file size (KB):** Click the arrows to select, or type the maximum file size (in kilobytes) allowed for the files created in the target.
- **Limit storage size to (MB):** To allocate a specific amount of disk space to the File Channel files, enable this option and then specify the amount of disk space to set aside (using the arrows or by typing). When the limit is reached, Amazon RDS Migration Tool will stop writing the files to the designated storage.
- **Max batching time interval (seconds):** Click the arrows to select, or type the maximum time (in seconds) for files to be batched before being written in a single operation.
- **Transfer files to remote file channel:** Select this check box to transfer files to the File Channel Source (on the remote Amazon RDS Migration Server) using the Amazon RDS Migration File Transfer Service. This can dramatically improve transfer speeds when the source database and the target database are located on different LANs. For more information about the Amazon RDS Migration File Transfer Service, see [File Transfer Service](#).
 - **Remote file transfer service host:** The host name or IP address of the computer on which the Amazon RDS Migration File Transfer Service is running.

- **Remote file transfer service port:** The port on the remote computer through which the files will be transferred (from the storage folder to the remote file channel).
- **Remote file channel database name:** The name of the File Channel Source database on the remote machine.
- **Additional remote file channels:** When sending to multiple File Channel Source databases, specify the target destinations using the following format:
`file_channel_db_name@host:port,file_channel_db_name@host:port`
- **Max transfer streams:** The maximum number of streams to use when transferring the files. Adjust the number of streams as required to optimize transfer speeds.
- **Password:** The password that will be used to establish a secure connection with the File Channel Source.

Important: When using the File Transfer Service, an agreed upon password is required in order to establish a secure connection between the File Channel Source and the File Channel Target. Accordingly, the password specified in the File Channel Target settings and the password specified in the File Channel Source(s') settings *must be identical*.

- **Override connection string parameters:** Select this to use a custom connect string. The following is an example of a string you can enter in this dialog box:

```
BufferSize=1000;MaxStorageSize=100000;MaxBatchingTimeInterval=120;MaxFileSize=32000;path=\\folder\file_channel;
```

For a description of the supported connection string parameters, see [File Channel Target Properties](#).

Using the Amazon RDS Migration Console

This chapter describes the elements of the Amazon RDS Migration Console and how to control access to the console by taking advantage of Amazon RDS Migration Tool's built-in security roles. The Console is a Web-based interface that runs in most modern browsers. (For information on supported browsers, see [Additional Required Software](#).) You can connect from any computer to the Amazon RDS Migration Server. The following topics describe the Amazon RDS Migration Console.

- [The Tasks View](#)
- [Viewing Specific Tasks](#)
- [The Server View](#)
- [Security Roles](#)

The Tasks View

The tasks view provides you with a list of the tasks you defined. Use this view to navigate through the tasks and view them.

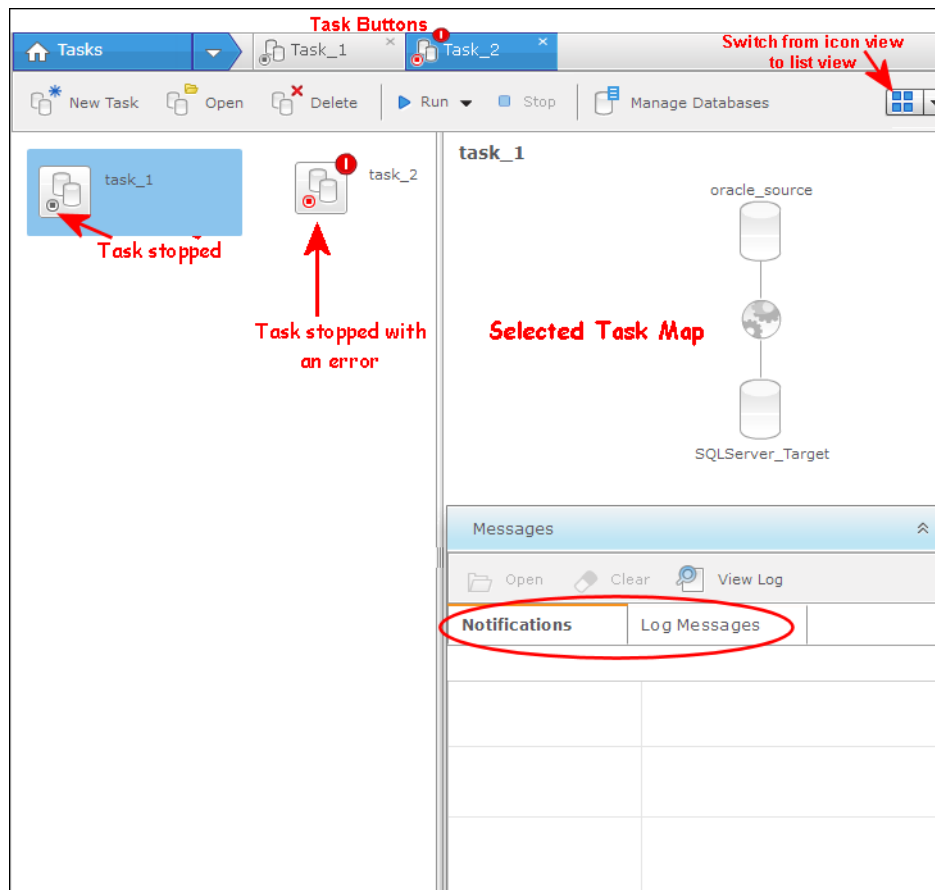
To access the Tasks view

Select **Tasks** from drop-down list at the top left corner under the Amazon RDS Migration Tool logo.

Note: The task view is the default view when you open Amazon RDS Migration Tool for the first time.

The following is an example of the Amazon RDS Migration Console in the Tasks view.

Figure 13–1 Task View



This view shows the following:

- **New Task button:** Click this to set up a new task. See [Setting up Tasks](#).
- **Task buttons:** Click on the Task buttons to view information about the selected task. For information about viewing tasks, see [Viewing Specific Tasks](#). These are available in the Task view only.
- **Database buttons** (not shown). Click on the database button to view information about the source and target databases that you add to Amazon RDS Migration Tool. This is available in the Settings view only. For more information, see [The Server View](#).

- **Task List:** This shows all of the tasks that were created. You can show the tasks in icon view (as shown above) or in list view. Note that each icon shows the current state of the task. See [Task Icons](#) for more information.

To change from icon view to list view, click the view from the drop-down at the top right of the screen, as shown in the figure above.

For information about creating a task, see [Designing Tasks](#).

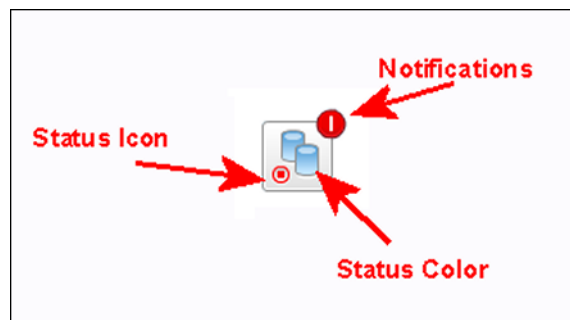
- **Task Map:** The task map shows a graphic view of the selected task. Any notices about the task as it is run is displayed at the bottom. For a more detailed view of a task, select the task from the button bar to view it in.
- **Messages (Notifications and Log Messages):** The bottom of the Task Map in the Tasks pane displays various notifications and log messages about the task. The message should give you an idea of the problem encountered. For information on viewing messages, see [Reading Messages about a Task](#).

The messages display the time of a notification and a description of the notification. You can define the descriptions displayed for a notification. For more information, see [Define the Notification Message](#).

Task Icons

Icons in the task view show the current status for each task. The following figure describes the information shown in the icon.

Figure 13–2 Task Icon Information






The figure above shows the following:

- **Notifications:** This indicates whether there are any errors or warnings when the task is running. The following information describes the Notifications icon:
 - The number indicates the amount of notifications. If there are more than nine notifications, an asterisk (*) is displayed.
 - The shape and color indicate the severity of the most severe error. Red circle indicates an error, yellow triangle indicates a warning, and a blue circle indicates information only.
- **Status Color:** If the database icons in the middle are white, the task is stopped. If they are blue, the task is running. A running task may still be paused or suspended due to an error or other situation.
- **Status Icon:** This icon indicates whether the task is running or stopped or stopped due to an error. See the table below for a detailed explanation.

The following shows two examples of task icons.

Table 13–1 Task Icons

Task Icon	Description
	This icon has a status icon that shows that the task is stopped. This is also indicated by the color of the database icons.
	This status icon shows that the task is stopped due to an error. The red stopped icon shows this and the blue database icons indicate the task process is still active. You can stop the task from the task monitor in the Amazon RDS Migration Console. The red circle shows that there is at least one error notification.
	This icon indicates that the task is running. The status icons and the database color (blue) indicate this.

Viewing Specific Tasks

The task view, in addition to listing the tasks you have configured, lets you view each task in detail. It is in the task view that where you design each task, then run and monitor the task.

To view a specific task

1. From the left side of the Task view, select the task you want to work with.

The task diagram is displayed on the right side of the page.

Note: If there are no tasks listed in the task view, you must create a task. See [Designing Tasks](#) for more information.

2. from the button bar at the top of the right pane, click **Open**.

The task view has two modes:

- [Design Mode](#)
- [Monitor Mode](#)

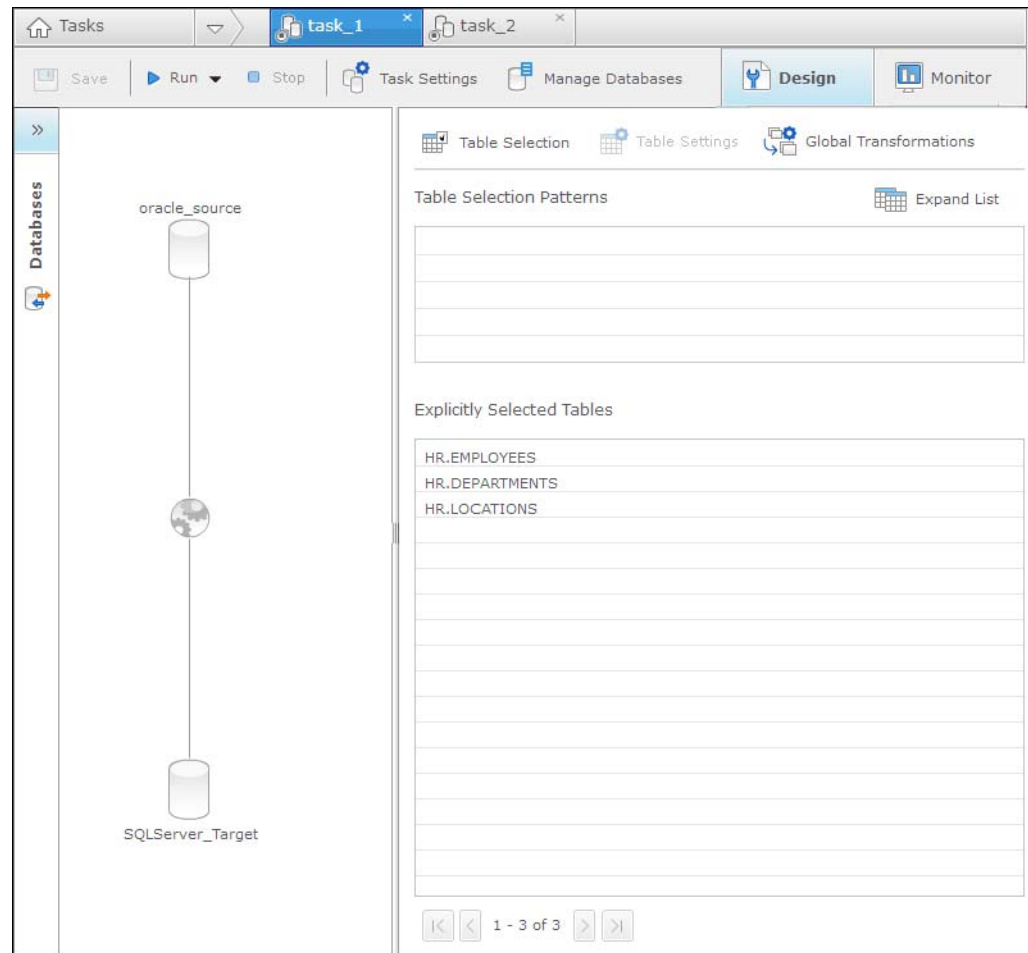
Design Mode

You use the design mode to create a task and assign tables to be replicated. To view the Design mode, click **Design** from the right side of the button bar.

Note: The Design mode is the default mode displayed when you view a task.

The following figure shows the overview of the design mode:

Figure 13–3 Viewing a Task in Design Mode



The following describes the design mode:

- **Database List:** This lists the source and target databases that you add to Amazon RDS Migration Tool. For more information, see [Working with Databases](#). The database list may be collapsed and the list not visible. In that case, click the arrows at the top to open the list.
- **Task Map:** This shows the connection between the source and target databases for the task. When creating a task, drag the databases to the correct place. For more information, see [Adding a Source and Target Database to a Task](#).
- **Change to Monitor Mode:** Click the correct button to switch between design mode and monitor mode.

See also:

[Monitor Mode](#)

[Monitoring and Controlling Replication Tasks](#)

- **Run Task buttons:** Use these buttons when you want the task to run in real time. You will not be able to view the task activity unless you switch to monitor mode. For more information, see [Reading Messages about a Task](#).
- **Task Settings Button:** Click this to open the Task Settings dialog box to make specific configurations for a task, see [Task Settings](#).

- **Manage Databases button:** Click this button to view the databases, edit them, or add new databases. For more information, see [Working with Databases](#).
- **Select and Define Tables:** This area lets you select the tables you want to include in your replication task. In addition, you can use transformation and filter operations to create new tables or to replicate parts of tables.

See also:

[Adding Tables to a Task](#)

[Using Filters](#)

[Define Transformations on a Single Table](#)

- **Global Transformation:** Click this link to create transformations for all tables in a task.

For more information, see [Defining Global Transformations](#).

- **Settings (not shown):** The settings tab in the Add Tables pane, lets you configure specific parameters for a task and view the task logs.

See also:

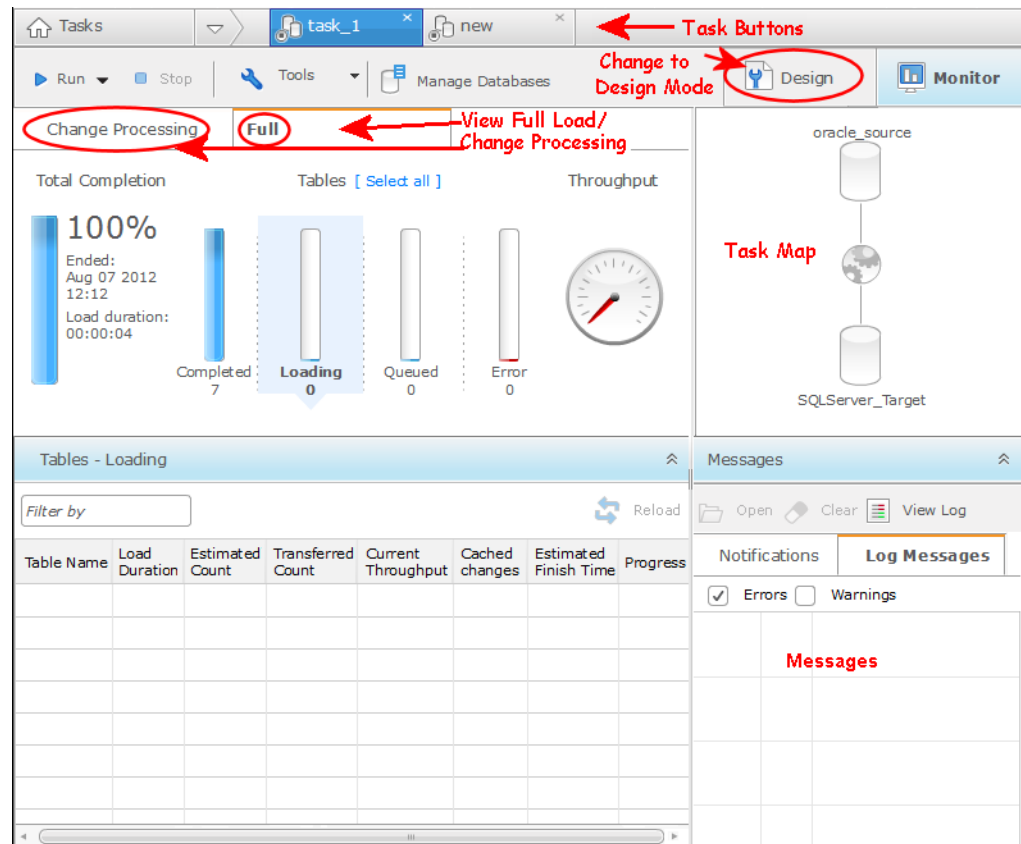
[Task Settings](#)

Monitor Mode

The monitor mode lets you view the replication task activities in real time. To view the Monitor mode, click **Monitor** from the right side of the button bar.

The figure below shows an overview of the monitor mode:

Figure 13-4 Viewing a Task in Monitor Mode



The following describes the monitor mode:

- **Task Map:** This shows the connection between the source and target databases for the task. For more information, see [Adding a Source and Target Database to a Task](#).
- **Change to Design Mode:** Click the correct button to switch between design mode and monitor mode.
See also:
[Design Mode](#)
[Designing Tasks](#)
- **Run Task buttons:** Use these buttons when you want the task to run in real time. You will can view the task activity in monitor mode. Monitor mode provides information on the full load process, change processing, and displays notifications about events that may occur during a task. For more information, see [Monitoring and Controlling Replication Tasks](#).
- **Task buttons:** Click on the Task buttons to view information about the selected task.
See also:
[Viewing Specific Tasks](#)
[Designing Tasks](#)
- **Change Processing/Full Load:** Select the monitor view that you want to see. By default, the **Full Load** view is displayed.

Note: The figure above shows the **Full Load** view.

For more information, see [Monitoring and Controlling Replication Tasks](#).

- **Messages:** This displays notifications and logging messages. For more information, see [Reading Messages about a Task](#).

The Server View

The Server view lets you view and configure options for the Amazon RDS Migration Server. The following options are available in the settings view.

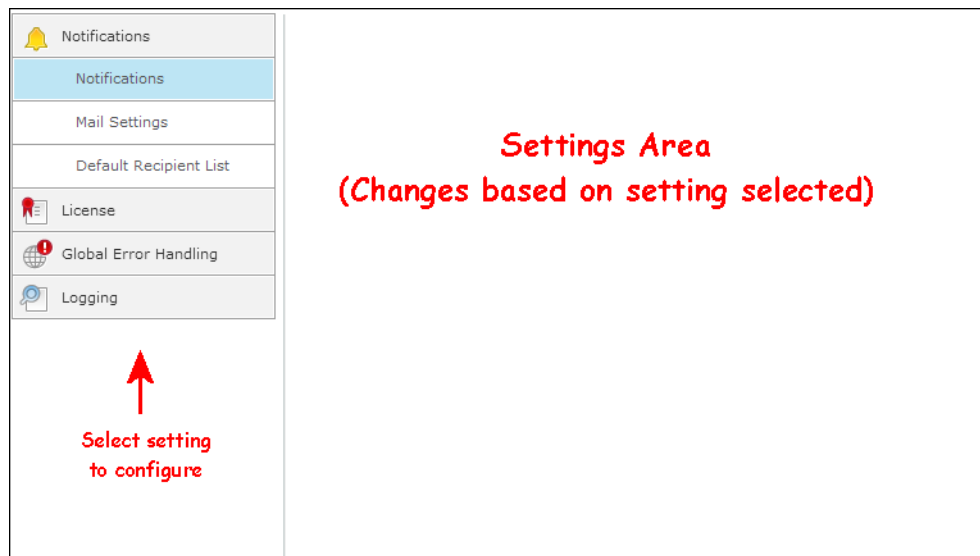
- **Notifications:** Lets you create and configure notifications, configure the mail server, and create a default recipients list for sending notifications.
See [Notifications Settings](#) for more information.
- **License:** Lets you request or register the license for Amazon RDS Migration Tool.
See [License Settings](#) for more information.
- **Global Error Handling:** Lets you set error-handling policy for the Amazon RDS Migration Server.
See [Error Handling](#) for more information.
- **Logging:** Lets you configure the logging settings for the Amazon RDS Migration Server.
See [Logging Settings \(Server\)](#) for more information.

To access the Settings view

Select **Server** from drop-down list at the top left corner under the Amazon RDS Migration Tool logo.

The following figure shows the Amazon RDS Migration Console in the Server view.

Figure 13–5 Server View



For additional information, see [Amazon RDS Migration Server Settings](#).

Security Roles

Security roles allow you to grant Amazon RDS Migration Tool users different roles according to the tasks you want them to perform. Four predefined security roles are available: Admin, Designer, Operator and Viewer. Each role has its own set of permissions, which are described in [Table 13–2](#) below.

Table 13–2 Security Role Permissions

Permission	Admin	Designer	Operator	Viewer
Monitor tasks (in Monitor view)	Yes	Yes	Yes	Yes
Perform runtime operations (e.g. start, stop, reload tasks, and so on)	Yes	Yes	Yes	No
Create and design tasks	Yes	Yes	No	No
Delete tasks	Yes	Yes	No	No
Manage databases	Yes	Yes	No	No
Edit server settings (e.g. notifications, license registration, global error handling, and so on)	Yes	No	No	No

Security roles cannot be edited nor can the permissions assigned to each of the roles be changed. The user under whose account Amazon RDS Migration Tool is installed will be associated with the Admin role by default.

You can associate users with each of the security roles by adding the user to the appropriate Active Directory group.

Creating Active Directory Groups

User roles are predefined in the UserConfiguration.xml file as **GroupRef** elements.

The default location of this file is: C:\Program Files\Amazon\RDS Migration Tool\data

The table below shows the available roles and their corresponding **GroupRef** element.

Table 13–3 Roles and their Corresponding GroupRef Element Names

Role	GroupRef Element
Administrator	AmazonRDSMigrationToolAdmins
Designer	AmazonRDSMigrationToolDesigners
Operator	AmazonRDSMigrationToolOperators
Viewer	AmazonRDSMigrationToolViewers

To enable security roles, you need to create Active Directory groups with the same names as the **GroupRef** elements listed in the table above.

Any users added to the Active Directory groups will be able to perform tasks in the Amazon RDS Migration Console according to their group's predefined [Security Role Permissions](#).

The following is an example of how these roles appear in the UserConfiguration.xml file. Make sure that the computer can access the domains where the users are located.

```
<Role name="Admin" anonymous="false">
```

```

<UserRef name="MY_DOMAIN\Dave.Smith" />
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolAdmins"
useSamAccountName="false" />
</Role>
<Role name="Designer" anonymous="false">
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolDesigners"
useSamAccountName="false" />
</Role>
<Role name="Operator" anonymous="false">
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolOperators"
useSamAccountName="false" />
</Role>
<Role name="Viewer" anonymous="false">
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolViewers"
useSamAccountName="false" />
</Role>

```

Associating Multiple Groups with the Same Role

You can associate users from different Active Directory groups with the same role by adding a new **GroupRef** element to the UserConfiguration.xml file.

For example, to associate users from an Active Directory group called **MyOperators** with the **Operator** role, you would need to edit the UserConfiguration.xml file as follows:

```


<Role name="Admin" anonymous="false">
  <UserRef name="MY_DOMAIN\Dave.Smith" />
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolAdmins"
useSamAccountName="false" />
</Role>
<Role name="Designer" anonymous="false">
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolDesigners"
useSamAccountName="false" />
</Role>
<Role name="Operator" anonymous="false">
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolOperators"
useSamAccountName="false" />
  <GroupRef name="MY_DOMAIN\MyOperators" useSamAccountName="false" />
</Role>
<Role name="Viewer" anonymous="false">
  <GroupRef name="MY_DOMAIN\AmazonRDSMigrationToolViewers"
useSamAccountName="false" />
</Role>

```


Security Roles and Console Elements

In the Amazon RDS Migration Console, users will only see menu items and buttons that are appropriate for their particular role.

Designers, Operators or Viewers will not see the **Server** entry in the Tasks/Server menu.




Operators and Viewers will not be able to see the **Design** tab or navigate to it. If an Operator or Viewer opens a task, the **Monitor** view will be shown.




In the **Tasks** view, only the following buttons will be displayed:

For Operators:



For Viewers:




In the **Monitor** view, only the following buttons will be displayed:

For Operators:



No action buttons will be available to Viewers, since Viewers are not allowed to perform any operations.

The security role associated with the Active Directory user appears in parenthesis to the right of the user label in the top right corner of the console.



Designing Tasks

This chapter describes how to design a replication task. To design a replication task, you must first be sure that you have configured at least one source and one target database to work with Amazon RDS Migration Tool. Then you can create the task and add tables to be replicated.

It is also possible to customize a task by creating new tables or columns for the target database or by selecting only some of the data from each column to be replicated. This is done using transforms and filters. For more information, see [Customizing Tasks](#).

For more information about replication tasks, see [What is a Replication Task](#).

This chapter contains the following topics:

- [Setting up Tasks](#)
- [Working with Databases](#)
- [Adding a Source and Target Database to a Task](#)
- [Adding Tables to a Task](#)
- [Editing a Replication Task](#)
- [Deleting a Replication Task](#)
- [Importing and Exporting Tasks](#)

Setting up Tasks

Before you get started with designing all of the features that you need for a task, you must define its default behavior.

To get started setting up a task

1. In the Tasks view, click **New Task**. The New Task dialog box is displayed.
2. Type a name for the task. The name can be anything that is useful in remembering what the task is doing.
3. Type a description for the task. This is optional.
4. Toggle the task options:
 - **Full Load:** Click this button to enable or disable Full Load options for this task.

When full load is enabled, Amazon RDS Migration Tool loads the initial source data to the target database. By default a full load is carried out for this

task. If you want to change this setting after you begin working with this task, you make the change in the Task Settings, [Full Load](#) tab.

- **Apply Changes:** Click this button to enable or disable Apply Changes (Change Processing).

When Apply Changes is enabled, Amazon RDS Migration Tool processes the changes. By default, change processing is carried out for this task. You can view the change processing in the Monitor view. For more information, see [Monitoring Change Processing Operations](#). If you want to change this setting after you begin working with this task, you make the change in the Task Settings, [Change Processing](#) tab.

Working with Databases

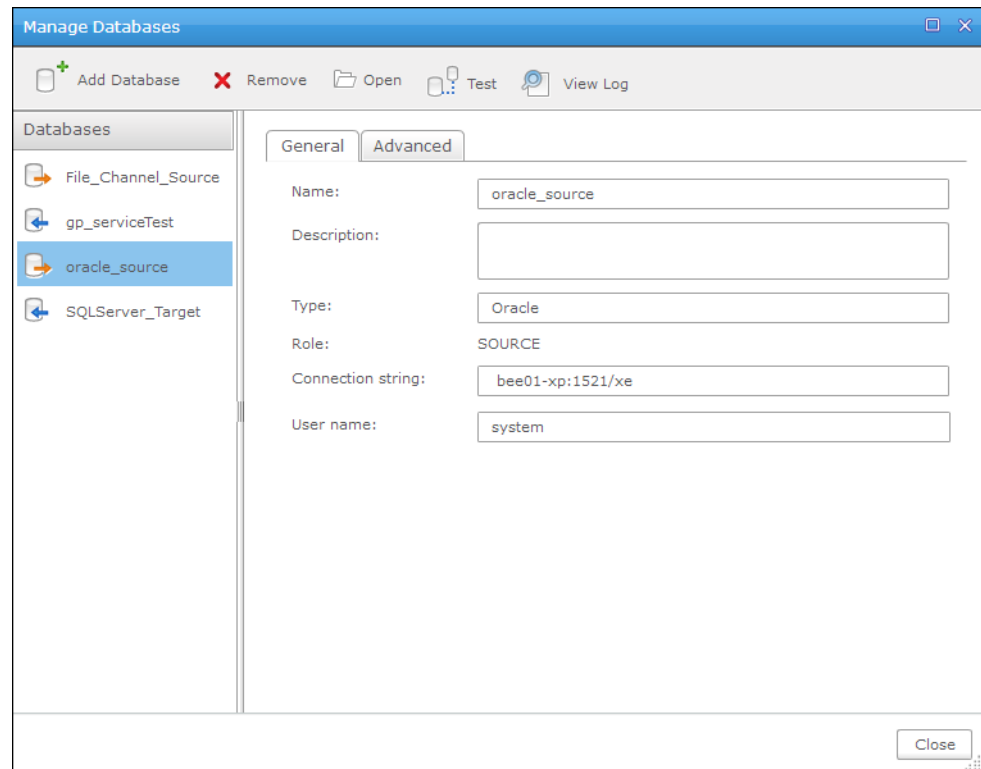
Amazon RDS Migration Tool requires information to connect to any supported database you are using as a source or target. For a list of databases you can work with in Amazon RDS Migration Tool, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

You use the Manage Databases dialog box to add databases and edit and view the database configuration information.

- [Adding a Database](#)
- [Editing Database Configuration Information](#)
- [Viewing Database Configuration Information](#)

The figure below is an example of the Manage Database dialog box. Note that the connection information displayed depends on the database type that you select.

Figure 14–1 Manage Databases



Adding a Database

Before you can begin to design a task, you must add databases to Amazon RDS Migration Tool. To use a database you must have access to it somewhere in your system. When you add the database to Amazon RDS Migration Tool, you must provide connection information and proper user credentials.

To add a database

1. In the Tasks view, click **Manage Databases** in the toolbar. The Add Database dialog box opens.
2. In the Manage Databases dialog box click **Add Database**. The Add Database dialog box opens.
3. In the Add Database dialog box, select the type of database you are using. The information that you must enter depends on which database you select. For more information, see the chapter that describes the database you are using. For a list of databases supported by Amazon RDS Migration Tool, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

Once you add databases to Amazon RDS Migration Tool, you can begin to use them to build a replication task. For information on how to add a database to a replication task, see [Adding a Source and Target Database to a Task](#).

Editing Database Configuration Information

After you add the database to Amazon RDS Migration Tool and provide the connection information, you can make changes to some of the information.

To edit database configuration information

1. In the Manage Databases dialog box, double-click the database you want to edit.

or

Select the database you want to edit and click **Open** from the toolbar at the top of the dialog box.

The database dialog box opens.

2. Make changes to the information in any of the tabs in the database dialog box. For more information about the information you can edit, see the chapter for the specific Amazon RDS Migration Tool endpoint you are using. For information which databases are supported by Amazon RDS Migration Tool, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

Note: You *cannot* change the following information in the database dialog box:

- The name you provided for the database.
 - The database **Type**, for example Oracle or Microsoft SQL Server.
 - The database **Role**, either SOURCE or TARGET.
-
-

Viewing Database Configuration Information

After you add the database to Amazon RDS Migration Tool and provide the connection information, you can view the information in the Manage Databases dialog box.

To view the database configuration information

Select a database from the database list in the left pane then click on the tabs to view all of the information. The information displayed is the same as the information provided in the database dialog box.

Testing a Database Connection

You can try to contact the database to make sure that you are connected to the database you want to work with.

To test the database connection

In the Manage Databases dialog box, select the database you want to work with.

1. From the top of the Manage Databases wizard, click **Test**.
2. If the connection is successful a message in blue is displayed. If the connection fails, a message in red that includes the log error information is displayed at the bottom of the dialog box.

To view the log entry if the connection fails, click **View Log**. The server log is displayed with the information for the connection failure. Note that this button is not available unless the test connection fails.

Adding a Source and Target Database to a Task

Once you have added the databases to Amazon RDS Migration Tool, you can begin to design the replication task. The first step in this process is to define the source


database from where your data is currently stored and the target databases where you want to replicate the data. Amazon RDS Migration Tool makes this very easy. You just drag the one of the databases you added into the task map in the Amazon RDS Migration Console.

To add source and target databases to a task

1. In the Tasks view of the Amazon RDS Migration Console, click **New Task**.

The Task map is displayed, with the available databases listed in the pane on the left as shown in the following figure.

Figure 14–2 Add Database to Task



2. Drag a source database to the top square in the task map, where it is indicated to drop the source database. If Amazon RDS Migration Tool does not let you do this, make sure that the database you are using is a source database. You can see this information in the [Manage Databases](#) dialog box.
3. Drag a target database to the bottom square in the task map, where it is indicated to drop the target database. If Amazon RDS Migration Tool does not let you do this, make sure that the database you are using is a target database. You can see this information in the [Manage Databases](#) dialog box.
4. Click **Save** to make sure that Amazon RDS Migration Tool saves this task.

Once you select the database for your task, you must select the tables from the source database to be replicated. The next step in creating a replication task is [Adding a Source and Target Database to a Task](#).

Adding Tables to a Task

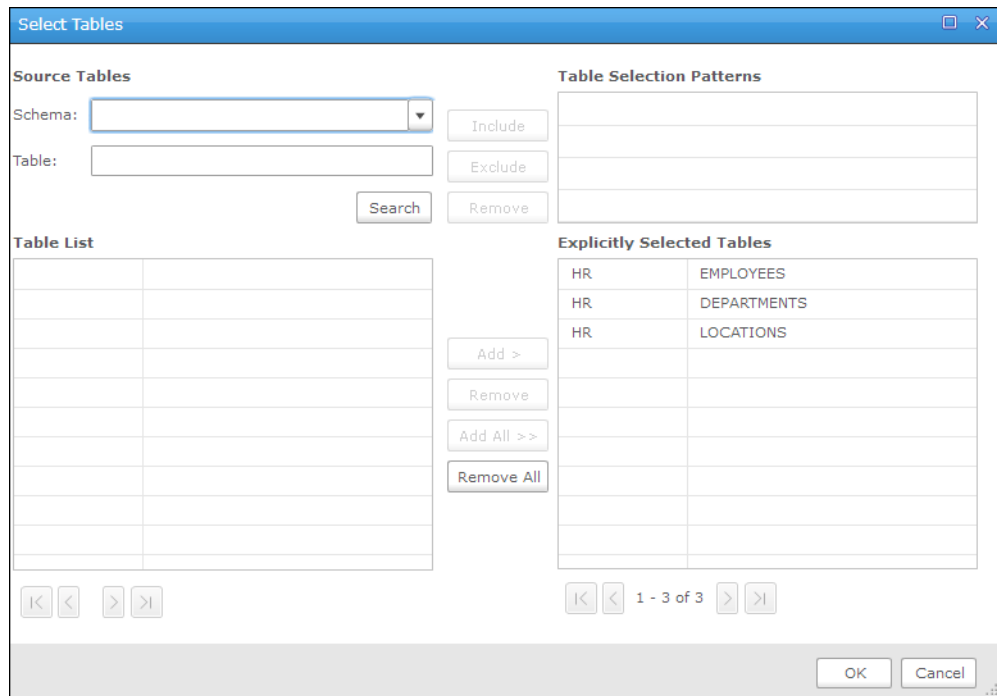
This procedure describes how to select the tables from the source database to be replicated. You select the tables using the Select Tables dialog box.

You access the Select Tables dialog box from the **Tables** tab in the replication task you are working with. Make sure you are in the Task view before you begin.

To open the Select Tables dialog box

1. Open the task you are working with if it is not displayed in the Amazon RDS Migration Console. For information on opening a task, see [Editing a Replication Task](#).
2. On the right side, click **Table Selection**. The Select Tables dialog box opens.

Figure 14–3 Select Tables Dialog Box



See the following for information on how to work with the Select Tables dialog box:

- [Search for Tables to use in a Replication Task](#)
- [Selecting Specific Tables for Replication](#)
- [Create Table Selection Patterns](#)

Once you have selected tables to replicate, you can run the replication task. However if you want to make any changes to the structure of the tables in the target database or only select specific columns, you will need to carry out one or both of the following procedures:

- [Define Transformations on a Single Table](#)
- [Using Filters](#)

Search for Tables to use in a Replication Task

To select specific tables to include in a replication task, you must first select them. In order to do this, you search for tables that match a specific criteria and then select the tables to include from the list provided. You can also carry out another search with new criteria and then add additional tables to the replication task.

To search for tables to use in a replication task

1. In the **Source Tables** section, do one or both of the following:
 - Select a table schema from the **Schema** drop-down list.
 - Type the name or partial name of a table in the **Table** field.

Note: You can also include special characters in your search string. For more information, see the [Note in Creating a Record Selection Condition for One or More Columns](#).

2. Click **Search** to display a list of tables. Any table that matches the criteria that you entered in the fields above will be displayed. You can use this table list to select tables to use for replication.

After you finish searching, you can [Selecting Specific Tables for Replication](#).

Selecting Specific Tables for Replication

This option lets you directly select tables to replicate in full.

To explicitly select tables for replication

1. Create a [Search for Tables to use in a Replication Task](#) to get a list of tables that you want to select from.
2. From the **Table List**, select a table that you want to include in the replication task and click **Add**.

The selected table is entered in the **Explicitly Selected Tables** list.

3. To select all the tables in the **Table List**, click **Add All**.
4. Click **OK** to close the Select Tables dialog box.
5. Click **Save** to make sure that Amazon RDS Migration Tool saves the table information for this task.

When you explicitly select tables, all selected tables are replicated in full unless you define transformations or filters for the table. If you need to make changes to the table structures in the target database or if you only want to select specific columns, then you will need to perform the procedures described in [Define Transformations on a Single Table](#) and [Using Filters](#) respectively.

Removing Specific Tables from a Replication Task

This option lets you remove specific tables from the replication task.

To remove tables from the Explicitly Selected Tables list

1. From the **Explicitly Selected Tables** list, select a table that you want to remove from the replication task and click **Remove**.
2. To remove all of the tables from the **Explicitly Selected Tables** list, click **Remove All**.
3. Click **OK** to close the **Select Tables** dialog box.
4. Click **Save** to make sure that Amazon RDS Migration Tool saves the table information for this task.

Create Table Selection Patterns

Using this option lets you select tables using patterns. For example, you can include all tables that belong to the HR schema except for one or two tables that you exclude. You can also only exclude one or more table schemas or tables. This will replicate the entire database except for those tables you excluded.

To select table selection patterns

1. In the **Source Tables** section, do any of the following:
 - Select a table owner from the **Schema** drop-down list. If you select this all tables that belong to that schema is included in the table selection pattern.
 - Type the name or partial name of a table in the **Table** field. Any specific table that you list here is included in the table selected pattern.

If the table that you type here is a member of the schema you selected in the **Schema** drop-down list, then you only have to type the name of the table.

If you did not select a schema or the table is a belongs to another schema, include the schema with the table name in the following format:
HR.EMPLOYEES, where HR is the owner.
2. Click **Include** to include all of the tables that match the criteria you selected in the Source Tables section.
3. Click **Exclude** to exclude any tables that math the criteria entered in the Source tables section.
4. Click **OK** to close the Select Tables dialog box.
5. Click **Save** to make sure that Amazon RDS Migration Tool saves the table information for this task.

The following example shows a pattern that will replicate all tables that are members of the HR schema except for the HR.EMPLOYEES table.

```
Include HR.%  
Exclude HR.EMPLOYEES%
```

When you explicitly select tables, all selected tables are replicated in full unless you define transformations or filters for the table. If you need to make changes to the table structures in the target database or if you only want to select specific columns, then you will need to perform the procedures described in [Define Transformations on a Single Table](#) and [Using Filters](#) respectively.

Note: To view all of the tables included when you use a table selection pattern, click **Expand List**. The **Full Table List** lists all of the tables included in any table pattern you defined and all explicitly selected tables.

Click **Collapse List** to return to the main part of the **Tables** tab.

Editing a Replication Task

You can make changes to tasks that you previously created. You must open the task and then make the changes in the same way that you did when you created the task.

To open a task

1. Make sure you are in the Task view. If you see **Settings** in the button bar at the top of the Amazon RDS Migration Console, click on the triangle and select **Tasks**.
2. Click the task you want to work with. A task map is displayed in the right pane.
3. At the top of the task map, click **Open**. The information about the task is displayed.

See the following for information on how to make the changes to a task.

- [Adding a Source and Target Database to a Task](#)
- [Adding Tables to a Task](#)
- [Define Transformations on a Single Table](#)
- [Using Filters](#)
- [Task Settings](#)

Deleting a Replication Task

You can delete tasks that you created.

To delete a task

1. Make sure you are in the Task view. If you see **Settings** in the button bar at the top of the Amazon RDS Migration Console, click on the triangle and select **Tasks**.
2. Select the task you want to work with.
3. At the top of the task map, click **Delete**. The task and all of its defined endpoints are deleted. The source and target database definitions that you created are still available to be used in other tasks.

Note: If you are using a Microsoft SQL Server database a Microsoft SQL Server system administrator must delete the Microsoft SQL Server Replication Publisher definitions for the database that was used in the task from SQL Server.

For more information see the [Limitations](#) in the Microsoft SQL Server chapter.

Importing and Exporting Tasks

You can export Amazon RDS Migration Tool tasks to a file. All exported tasks are saved to a folder or directory called **imports**, which is located in the folder or directory where Amazon RDS Migration Tool was installed.

You can import an export file (*.json) to another instance of the Amazon RDS Migration Server. This will let you use a task that you created in Amazon RDS Migration Tool in a different environment. For example, if you created tasks in a development environment and now want to use the task in a production environment.

Importing and exporting a task must be done from the command line. Therefore, you must carry out this task from the computer where Amazon RDS Migration Tool is installed;

Note: If you need to access the computer with Amazon RDS Migration Tool from a remote computer, you can use a telnet connection.

When you export a task to a different environment, you may need to edit the task information. For example, you may need to change the connect string for a database. The following section describes how to work with export files.

- [How to: Export and Import a Task](#)
- [Editing an Exported \(json\) File](#)

How to: Export and Import a Task

Use the command line to export a task and then import it into another Amazon RDS Migration Tool instance.

To export a task

1. From the Amazon RDS Migration Tool computer where you defined the task you want to import, open the Amazon RDS Migration Tool command line console by doing the following:

On Windows: Go to **All Programs** in the **Start** menu and point to Amazon RDS Migration Tool, then to **Utilities** and then select Amazon RDS Migration Tool **Command Line**.

A command-line console is displayed with the correct prompt for Amazon RDS Migration Tool.

Note: You can also open the Windows command-line console and change the directory to the following:

```
<Full path to the Amazon RDS Migration Tool root folder>\bin,  
for example to use the path to the folder or directory where Amazon RDS  
Migration Tool is installed by default type, C:\Program Files\Amazon\RDS  
Migration Tool\bin.
```

On Linux: Enter the Linux computer, then type the following before you continue:

```
source ./arep_login.sh
```

2. At the prompt in the command-line console, type the following:

```
reptl exportrepository task=<task_name>
```

A file called `task_name.json` is copied to the `<Full path to the Amazon RDS Migration Tool root folder>\data\imports` folder or directory.

You can import this file into another Amazon RDS Migration Tool instance.

To import a task

1. From the Amazon RDS Migration Tool computer where you want to export the task, open the Amazon RDS Migration Tool command line console by doing the following:

Go to **All Programs** in the **Start** menu and point to Amazon RDS Migration Tool, then to **Utilities** and then select Amazon RDS Migration Tool **Command Line**.

A command-line console is displayed with the correct prompt for Amazon RDS Migration Tool.

Note: You can also open the Windows command-line console and change the directory to the following:

```
<Full path to the Amazon RDS Migration Tool root folder>\bin
```

For example to use the path to the folder or directory where Amazon RDS Migration Tool is installed by default type, C:\Program Files\Amazon\RDS Migration Tool\bin.

2. At the prompt in the command-line console, type the following:

```
repctl importrepository json_file=<Full path to the exported *.json file>
```

The exported *.json file will be located in the import folder or directory on the original computer where the task was created and exported. (for example, c:\finance\replication\comp\data\imports\many_tables.json). If you are importing this task into a different environment, you should copy the file to a location on the second Amazon RDS Migration Tool computer and then import the file from there.

In many cases, when you import the task into a different environment, you will need to make changes to the task. For example, you may need to change the connect strings for the databases in the task or change the user password. In that case, you will need to edit the *.json file. See [Editing an Exported \(json\) File](#) for more information.

Editing an Exported (json) File

You can open the *.json file in any plain text editor. It is possible to make changes to any of the sections in the file; however, be sure that you only change the data and not the field names. For example, the entry "name" : :DB_Name" displays the name field for a source table in a defined database. In this case, you can change the data "DB_Name" but not the included metadata ("name").

Notes:

- Make any changes to the *.json file *before* you carry out the import operation.
- Information about the databases, tables, tasks, task settings, and logger settings should be changed using the Amazon RDS Migration Console after the file is imported.

To be able to use the new task, you will need to make changes to the database password and connection strings by editing the *.json file. See [Making Changes to the Database Connection Information](#) for more information.

Making Changes to the Database Connection Information

In the "databases" section you can make changes to the connection string and the password. The following is an example of the "databases" section of the *.json file.

```
"databases" : [{
```

```
"name": "Oracle_Source",
"type": "Oracle",
"connection_string": "server= bee01-xp:1521/xe;username=SYSTEM",
"authenticator":
"{01000000D08C9DDF0115D1118C7A00C04FC297EB010000003EA495B32CAAE14CB9777B96B3CC00B3
000000002000000000003660000A8000000100000002765A3287AB56447DA31508F71CE6270000000
0004800000A00000001000000088D5C1BBD615BEEEF5FAC1B9B0E20800800000075D89177A9C6F11B
1400000047B3110B80920DD9EB0A5FABA05679979B78DDD0}",
"role": "SOURCE"
}, {
"name": "SQLSERVER_Target",
"type": "SQLServer",
"connection_string":
"server=bee01-xp;database=tempdb;WindowsAuthentication=Y;CDCBCP=Y;FullloadBCP=Y;BC
PPacketSize=16384",
"role": "TARGET"
```

To change the connect string

1. In the *.json file, under "databases", find "connection string".

For example, "connection_string": "server= bee01:1521/xe;username=SYSTEM".

2. Change the information after the equal (=) signs as necessary.

For example, if the database in the new environment is on a computer called B2, change server=bee01 to server=B2.

Important: Make sure that the connect string remains between the quotation marks ("").

To change a database password

1. In the *.json file, under "databases", find "authenticator".

For example, "authenticator":

```
"{01000000D08C9DDF0115D1118C7A00C04FC297EB010000003EA495B32CA
AE14CB9777B96B3CC00B3000000000200000000003660000A80000001000
00002765A3287AB56447DA31508F71CE62700000000004800000A0000001
000000088D5C1BBD615BEEEF5FAC1B9B0E20800800000075D89177A9C6F1
1B1400000047B3110B80920DD9EB0A5FABA05679979B78DDD0}"
```

Note: The password is presented as an encrypted string.

2. Change the password string to the relevant password for the database in the new environment. Type the new password using plain text exactly as it is configured in the database. For example, 8yTkLMt.

When you save the file and then import it to the new environment, the password is encrypted automatically.

Customizing Tasks

This chapter describes how to customize a replication task. Customizing tasks lets you create new tables or columns for the target database or select only some of the data from each column to be replicated. This is done using transforms and filters.

For more information about replication tasks, see [What is a Replication Task](#).

The following topics describe all of the procedures that can be included in customizing a replication task.

- [Table Settings](#)
- [Defining Global Transformations](#)
- [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#)
- [Task Settings](#)

Table Settings

Table settings lets you define how the data for each individual table is replicated to the target. To carry out these tasks, use the Table Settings dialog box.

To open the Table Settings dialog box

1. Open the task you are working with if it is not displayed in the Amazon RDS Migration Console. For information on opening a task, see [Editing a Replication Task](#).
2. From the right side of the Amazon RDS Migration Console, double-click the table where you want to perform the transforms or select the table then click **Table Settings**. If the table you want to perform the transformation on was defined by creating a table selection pattern, click **Expand List** to view all of the tables.

For information on how to define table selection patterns, see [Create Table Selection Patterns](#).

3. You can do the following in the Table Settings window:
 - [Carry out General Tasks for a Single Table](#)
 - [Define Transformations on a Single Table](#)
 - [Using Filters](#)
4. When you are finished entering the information, click **OK** to close the Table Settings window.
5. Click **Save** to make sure that Amazon RDS Migration Tool saves the table and column information for this task.

6. If you want to revert all of the changes you make to tables to their default values, click **Reset table to default**. This option is available at the bottom of all the tabs in this dialog box.

Note that this changes the data for all columns in the table to their default and removes any calculated columns that were added.

Note: This option is available for tables where changes were made. When you make changes to a table the table is displayed in the table list with the word **(changed)**.

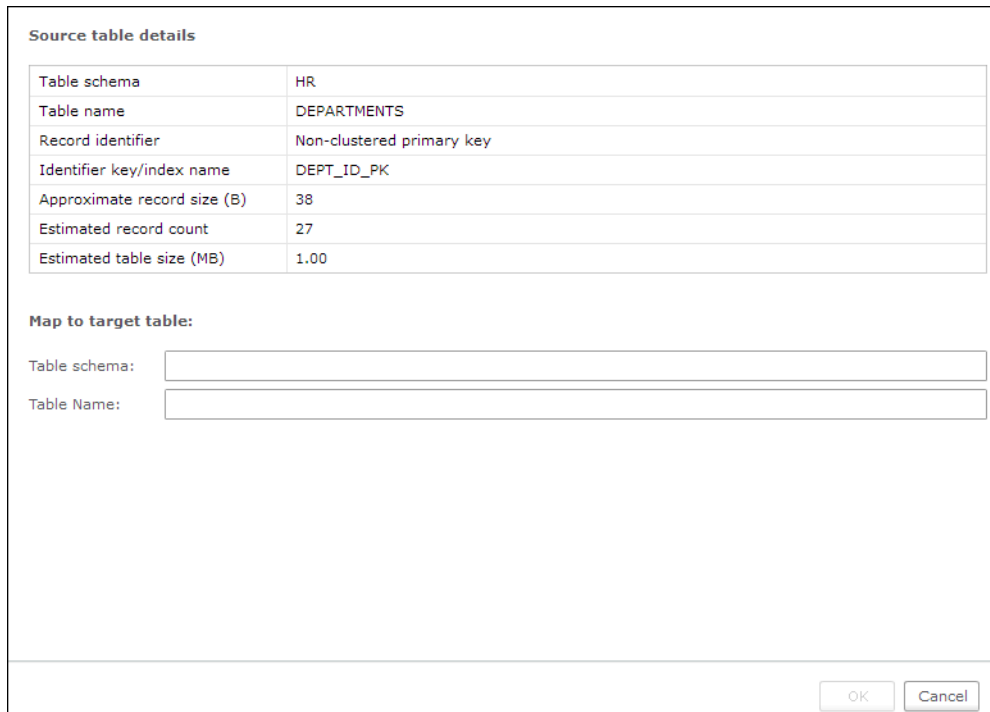
Carry out General Tasks for a Single Table

The **General** tab in the Table Settings window displays basic information about the selected table. See [Table 15-1](#) to view an example of this information. In addition you can use this tab to rename the target table or schema.

To edit the General table settings

Select the table you want to return to default and open the [Table Settings](#) window. Click **General** on the left side of the window. The following figure shows the information in the **General** tab of the Table Settings window.

Figure 15-1 Table Settings: General Tab



Source table details	
Table schema	HR
Table name	DEPARTMENTS
Record identifier	Non-clustered primary key
Identifier key/index name	DEPT_ID_PK
Approximate record size (B)	38
Estimated record count	27
Estimated table size (MB)	1.00

Map to target table:

Table schema:

Table Name:

The **General** tab lets you rename the table or table schema:

To Rename the table or table schema

In the **Map to target table** section, type one or both of the following:

- Type a new schema in the **Table Schema** field. This will add the schema name to the table. For example, if your table is `HR.Department`, type `PERS` to change the table name to `PERS.Department`.
- Type a new table name in the **Table Name** field. For example, if your table is `HR.Department`, type `DEPT2` to change the table name to `DEPT2`.

Note: If your table name includes a Schema, for example, `HR.DEPT`, when you add a value to both the Table Schema and Table Name fields, then the new table name in the example above is: `PERS.DEPT`.

Define Transformations on a Single Table

This section describes how define data transformations. Amazon RDS Migration Tool carries out these transformations when the task is run. Data transformation is optional. If you do not define any transformations, the data is replicated "as is" from the source to the target.

Amazon RDS Migration Tool lets you make the following changes to the tables and columns:

- Rename any column for the target table
- Delete a target column
- Change the data type and/or the length of any target column
- Add additional target columns
- Designate which target columns (i.e. segments) will comprise the Unique Index
- Recalculate the data

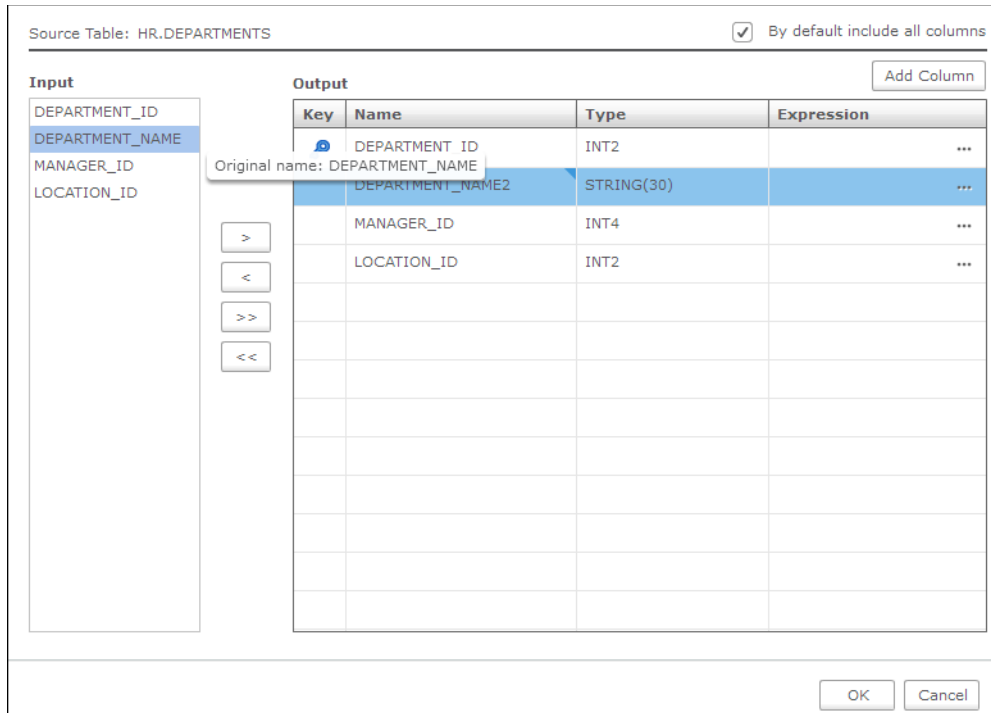
Limitations

- Transformations are not supported for calculating columns of Right-to-Left languages.
- The only supported transformation for LOB/CLOB data types is to drop the column on the target.

To define a data transformation for a single table

Select the table you want to return to default and open the [Table Settings](#) window. Click **Transform** on the left side of the window. The following figure shows the information in the **Transform** tab of the Table Settings window.

Figure 15–2 Table Settings: Transforms



Use this method if you need to make transformations that are specific to only one or a few tables in your task. For an explanation of how to configure transforms, see [Using the Transform Tab](#).

To make a similar change over multiple tables, see [Defining Global Transformations](#).

Using the Transform Tab

The **Transform** Tab in the Table Settings window has the following components:

- **Input list:** This lists the available columns from the where you are performing the transforms. You can drag the columns into the Expression column or the Calculate Expression field to create expressions. See the [Transformation Actions](#) table below for more information.
- **Output table:** This table shows the defined output for the columns in the table where you are performing the transforms. It has the following columns:
 - **Key:** This indicates whether the column is a segment of the Unique Index. A key icon is displayed next to columns that are segments of the Unique Index.
 - **Name:** The name of the column. To change the name of the column, double-click in the field with the column name you want to change and type a new name in this column if you want to change the name of the column or if the column is calculated (added to the table). See the [Transformation Actions](#) table for more information.
 - **Type:** The data type for the column. To change the data type for the column, double-click in the field with the data type you want to change and select a new data type. See the [Transformation Actions](#) table for more information.
 - **Expression:** An expression using SQLite operators to define the data in the column. For information on how to create an expression, see the [Transformation Actions](#) table below.

The following table describes the actions you can carry out in the Transform Table window.

Table 15–1 Transformation Actions

Action	Description
Rename a column	<p>Double-click the Name column for the table column you want to change. Type in a new name.</p> <p>The top right corner turns blue when the name is changed. To view the original name, hover the mouse pointer over the field and the original name is displayed.</p>
Change the data type for a column	<p>Double-click the Type column for the table column you want to change and select a new Amazon RDS Migration Tool data type from the drop-down list. Make sure that the data type you select is compatible with the data in that column. For a description of Amazon RDS Migration Tool data types, see Amazon RDS Migration Tool Data Types.</p> <p>For information about data-type mapping from the native database to Amazon RDS Migration Tool data types, see the chapter for the database you are using. For a list of databases supported by Amazon RDS Migration Tool, see Databases Supported as Amazon RDS Migration Tool Endpoints.</p>
Add a new column	<p>Click Add Column to add a new column. When you add a column, the Name is blank and the Type is listed as <code>string(50)</code>.</p> <p>Type a name for the new column in the Name column.</p> <p>Click in the Type column and select a data type from the list.</p>
Add an existing column	<p>From the Input pane, select one or more columns and click the right facing triangle button.</p> <p>To add all of the columns, click the right-facing double triangle.</p> <p>Note: By default all tables columns are included in the Output list. To include only some of the columns clear the By default include all columns check box at the top of the Transform Table window. This removes all of the columns from the list. You can then add back any existing column.</p>
Delete a column	<p>From the Output list, select the row with the column you want to delete and click the left-facing triangle button.</p> <p>To remove all columns, click the left-facing double triangle. Note that all the columns except for columns defined as a primary key are deleted.</p>
Add/Remove a Unique Index segment to/from a target column	<p>A key icon indicates which target columns segments of the Unique Index.</p> <p>To add a Unique Index segment, click in the Key column to the left of target column to which you want to add the segment. A key icon will appear.</p> <p>To remove a Unique Index segment, click the key icon to the left of the target column from which you want to remove the segment. The key icon will disappear.</p>

Table 15–1 (Cont.) Transformation Actions

Action	Description
Re-calculate the data for a column in the target database	<p>Click in the Expression column in the row with the table column you want to change the data for. Enter an expression using SQLite syntax.</p> <p>See Creating an Expression for Transforms and Using SQLite Syntax with Transforms for information on creating expressions.</p> <p>Once you add a calculated expression, you can test the expression. See Using the Expression Builder (for Filters, Transforms, and Global Transformations).</p>

Creating an Expression for Transforms

You use an expression to define the contents of a new or re-calculated column.

To create an expression

- In the Transform window, select the row with the column where you want to create a new expression.
or
Click **Add Column** to add a new column.
- Drag column inputs into the Calculation Expression field at the bottom of the window. This will create an expression that includes the value of the column.
- Add strings or number values to the Calculation Expression. These will affect the value of the column after the transformation.
- Add SQLite operators between the inputs and or strings/number values, for example `||`, which concatenates the strings for the two inputs you add.

The following are some examples of expressions you can use in the transform window.

```
FIRST_NAME || LAST_NAME
```

This combines the string values from each of the columns.

```
$$SALARY+0.1*SALARY
```

This adds 10 per-cent to each salary value. The salary column should be a numeric data type.

For information on the SQLite syntax operators you can use, see [Using SQLite Syntax with Transforms](#).

Using SQLite Syntax with Transforms

Amazon RDS Migration Tool supports the following SQLite operators.

Table 15–2 SQLite Operators used by Amazon RDS Migration Tool

Operator	Description
<code> </code>	<p>Concatenate strings.</p> <pre>FIRST_NAME LAST_NAME</pre> <p><code>PHONE_NUMBER <Office Only></code> (adds the string Office Only to the telephone number).</p>
<code>+</code>	<p>Adds two values together.</p> <pre>DEPARTMENT_ID+100</pre> <p>(adds 100 to each ID number). Any column used in an expression with this operator must be a numeric data type.</p>

Table 15–2 (Cont.) SQLite Operators used by Amazon RDS Migration Tool

Operator	Description
-	Subtracts a value from another value. MANAGER_ID-100 (subtracts 100 from each ID number). Any column used in an expression with this operator must be a numeric data type.
%	Uses the remainder of a division expression as the value. %SALARY/7 (Divides the value of the Salary column by 7 and uses any remainder from the expression as the column value).
/	Divides one value into another. SALARY/.16 (Divides the value of the Salary column by .16. Note: If the two values in the division expression are integers (two NUMERIC columns with no digits after the decimal) and the result is a fractional value, the result returned will be 0.
*	SALARY*.16 (Multiplies the value of the Salary column by .16. This could be used to calculate taxes that are subtracted from a salary).

For more information about SQLite syntax, see the SQLite documentation.

Using Filters

The filtering operation lets you create filters that define the information from a column to include in/exclude from a replication task. This lets you replicate only the specific data that you need.

This section covers the following topics:

- [Limitations](#)
- [Opening the Filter Table Window](#)
- [Creating a Filter Condition for a Specified Column](#)
- [Creating a Record Selection Condition for One or More Columns](#)
- [Using the Range Builder](#)
- [Using SQLite Syntax with Filtering](#)

Limitations

When creating a filter, the following limitations apply:

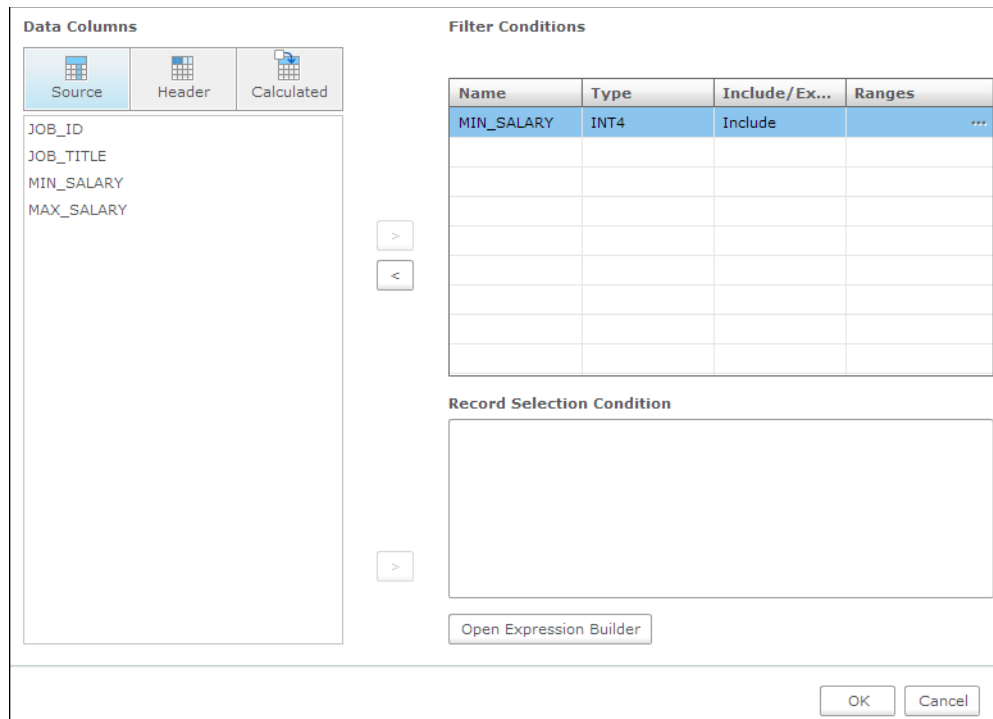
- Filters are not supported for calculating columns of Right-to-Left languages.
- When a filter is created to exclude specific rows in a column, the specified rows will always be excluded, even if the rows that were initially excluded are later changed. For example, if you chose to exclude rows "1-10" in a column named "Age" and those rows were later changed to "11-20", the rows will continue to be excluded, even though the data is no longer the same.

Opening the Filter Table Window

1. Select the table you want to return to default and open the [Table Settings](#) window.
2. Click **Filter** on the left side of the window.

The following figure shows the information in the **Filter** tab of the Table Settings window.

Figure 15–3 Table Settings: Filter



The Filter Table tab has the following information:

- **Data Columns** list: This list contains a list of the columns for the table where you filtering data. You can use these to select the columns to use in the filtering operations.

This list has the following tabs:

- **Source:** This tab lists the original source columns in the table.
- **Header:** This tab lists the available header columns. You can create filters using these columns and include them in expressions. For information on these header columns, see [Header Columns](#).
- **Calculated:** This tab lists the columns added to the table. You add columns through transformations. For more information, see [Define Transformations on a Single Table](#).
- **Filter Conditions** table: This table has the following columns:
 - **Name:** The name of the column where you are filtering the data.
 - **Type:** The data type for the column.
 - **Include/Exclude:** Indicate whether to include or exclude the filtered data for this column.

- **Ranges:** Click the button on the right of the Ranges field to open the Range Builder or type a string(s) with a specific value or a range(s) of values to be replicated. For more information on creating filter ranges, see [Creating a Filter Condition for a Specified Column](#).

For information on creating a value or ranges with the Range Builder, see [Using the Range Builder](#).

For more information on typing in the filter ranges manually, see [Using SQLite Syntax with Filtering](#).

- **Record Selection Condition:** Enter a complex condition that can include multiple columns. The condition must evaluate to TRUE to be accepted. You create a condition using SQLite operators or by [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#). For information on using the SQLite operators, see [Creating a Record Selection Condition for One or More Columns](#).

Creating a Filter Condition for a Specified Column

You can create a simple condition for a single column in the table you are working with. You can include any combination of ranges or specific values in the filter and determine whether to include or exclude the defined data.

To create a filter condition

1. Select a column from the data columns list and then click the right-facing triangle next to the **Filter Conditions** table.

To remove the column, click on it in the **Filter Conditions** table and then click the left-facing triangle. Any data entered for this column in the **Include/Exclude** or **Values** columns is also deleted.

2. Click in the **Include/Exclude** column to select whether to include or exclude the data that meets this condition.

3. Click in the **Ranges** column and enter a filter value. You can do any of the following:

- Click the button on the right of the **Ranges** field to open the Range Builder. For information on entering a value using the Range Builder, see [Using the Range Builder](#).
- Type a specific value, for example 32 or book. Any data of this value will be included or excluded.
- Type multiple strings of data separated by a comma. For example, APPLE, ORANGE, PEAR, PEACH. Any data that matches any of these values will be included or excluded. You can use < or > to find names with higher or lower alphabetical values. For example, Fruit List > A will list only ORANGE, PEAR, and PEACH.

- Type a range of values. A range of values is written with . . to determine the range. Some examples:

23 . . 92 indicates all data from 23 to 92. Therefore if you have ID numbers and you want to exclude all the numbers from 23-92, enter this value and select **Exclude**.

. . 50 indicates all values less than 50.

50 . . indicates all values more than 50.

- Type multiple ranges separated by commas. For example, 25..30, 35..40, 45..50
- Enter any combination of string values and ranges separated by commas. For example, 1, 5, 10, 25..30

Creating a Record Selection Condition for One or More Columns

You can create a record selection condition using SQLite operators by using drag and drop or by using the Expression Editor.

To create a record selection condition using drag and drop

1. Drag a column into the Record Selection Condition field at the bottom of the Filter Table window.

or

Type in a string value to use in the condition.

2. Use SQLite operators, such as < or = to create the condition. Use any amount of strings or columns as you need to create a condition.

For example `$EMPLOYEE_ID < 100 AND $SALARY > 100,000`

In this case only rows that satisfy both of these conditions are replicated in the replication task.

The following example provides an example using SQL search pattern strings. Only rows that satisfy this condition are replicated.

```
$EMPLOYEE_NAME IS 'Sm_th'
```

Note: You can use the following special characters when entering a string:

- %: Matches any string of zero or more characters. For example, `Mc%` searches for every name that begins with **Mc** or `%bob%` includes every name that contains **bob**.
- _: Matches a single character (as a wildcard). For example: `'Sm_th'` includes names that begin with **Sm** and end with **th**, such as **Smith** or **Smyth**. To search for an underscore character, use `[_]`
- `[. .]`: Includes a range or set of characters. For example, `[CK]ars[eo]` includes names **Carsen**, **Karsen**, **Carson**, and **Karson** or `[M-Z]inger` includes all words that end in **inger** with the first letter between **M** and **Z**, such as **Ringer**, **Singer**, or **Zinger**.

For more information see documentation on how to use Transact-SQL.

For information on what SQLite operators can be used to create Record Selection Condition filters, see [Using SQLite Syntax with Filtering](#).

To create a record selection condition using the Expression Builder

Click **Open Expression Builder**. This button is directly under the record selection condition box. Follow the directions for creating an expression in the section [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#).

Using the Range Builder

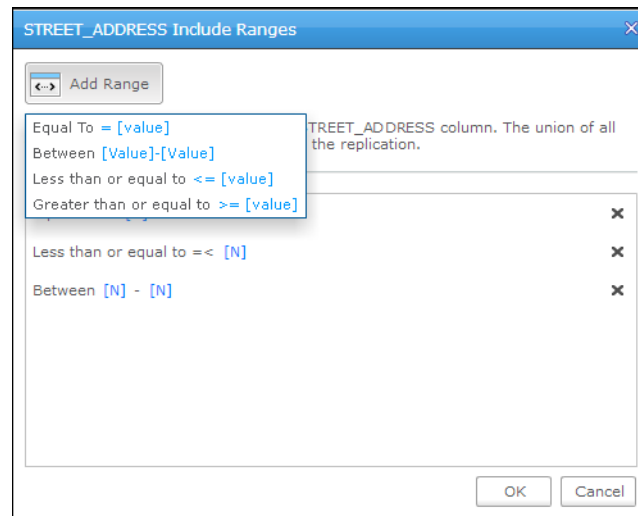
You can add values to the Ranges column using the Range Builder. You can enter one or more ranges in this column using the Range Builder.

Note that values that match *any* of the ranges in the list are included in the replication.

To use the range builder

1. In the Filter tab of the [Table Settings](#) window, select a column to filter. For more information, see [Using Filters](#).
2. Click the button to the right of the **Ranges** column. The Ranges Builder opens.

Figure 15–4 Range Builder



3. Click **Add Range**. Select any of the following from the drop-down list displayed.

- **Equal to:** Select **Equal to** to enter a single value. The following is displayed in the range list.

Equal to = [N]

Click the [N] and type a value in the field that is displayed.

When the value in the selected column equals the value you enter, the result is included or excluded in the replication task depending on the option selected in the **Include/Exclude** column.

- **Between:** Click **Between** to enter a range of values. The following is displayed in the range list.

Between [N] - [N]

Click each [N] and type a value in the fields that are displayed.

When the column contains the values between the two values entered, the result is included or excluded in the replication task depending on the option selected in the **Include/Exclude** column.

- **Less than or equal to:** Select **Less than or equal to** and enter a maximum value. The following is displayed in the range list.

Less than or Equal to =< [N]

Click the [N] and type a value in the field that is displayed.

When the value in the selected column is equal to or less than the value you enter, the result is included or excluded in the replication task depending on the option selected in the **Include/Exclude** column.

- Greater than or equal to:** Select **Greater than or equal to** and enter a minimum value. The following is displayed in the range list.

Greater than or Equal to => [N]

Click the [N] and type a value in the field that is displayed.

When the value in the selected column is equal to or more than the value you enter, the result is included or excluded in the replication task depending on the option selected in the **Include/Exclude** column.

To delete a filter range from the Range Builder

- In the Filter tab of the [Table Settings](#) window, select the column with the filter condition you want to delete.
- Click the button to the right of the **Ranges** column. The Ranges Builder opens.
- Click the red X next to the range you want to delete. The deleted range is removed from the list.

Note: Filter ranges that you enter manually are also displayed in the Filter Builder. You can use the Filter Builder to delete them.

Using SQLite Syntax with Filtering

Amazon RDS Migration Tool supports the following SQLite operators when creating Record Selection Condition filters.

Note: You must put the (\$) in front of each input as shown below.

Table 15–3 *SQLite Operators used by Amazon RDS Migration Tool for Filtering*

Operator	Description
<	Is less than. \$SALARY<100000
<=	Is less than or equal to \$SALARY<=100000
>	Is greater than \$SALARY>100000
>=	Is more than or equal to \$SALARY>=100000
=	Is equal to \$SALARY=100000
!= or <>	Is not equal to \$SALARY!=100000
IS	Is the same as \$HIRE_DATE IS 2014-09-29 IS functions the same as = unless one or both of the operands are NULL. In this case, if both operands are NULL, then the IS operator evaluates to 1 (true). If one operand is NULL and the other is not, then the IS operator evaluates to 0 (false).

Table 15–3 (Cont.) SQLite Operators used by Amazon RDS Migration Tool for Filtering

Operator	Description
IS NOT	<p>Is not the same as</p> <pre>\$HIRE_DATE IS NOT 2014-09-29</pre> <p>IS NOT functions the same as != unless one or both of the operands are NULL. In this case, if both operands are NULL, the IS NOT operator evaluates to 0 (false). If one operand is NULL and the other is not, then the IS NOT operator evaluates to 1 (true).</p>
AND	<p>Both operands are true.</p> <pre>\$MANAGER_ID AND EMPLOYEE ID >100</pre>
OR	<p>Either operand is true.</p> <pre>\$MANAGER_ID OR EMPLOYEE ID >100</pre>

For more information on how to use the SQLite syntax, see the SQLite documentation.

Defining Global Transformations

Use Global Transformations to make similar changes to multiple tables, owners, and columns in the same task.

You may need to use this option when you want to change the names of all tables. You can change the names using wild cards and patterns. For example, you may want to change the names of the tables from `account_*` to `ac_*`. This is helpful when replicating data from an Microsoft SQL Server database to an Oracle database where the Microsoft SQL Server database has a limit of 128 characters for a table name and the Oracle database has a limit of 31 characters.

You may also need to change a specific data type in the source to a different data type in the target for many or all of the tables in the task. Global transformation will accomplish this without having to define a transformation for each table individually.

This section includes the following topics:

- [Limitations](#)
- [Starting the New Transformation Rule Wizard](#)
- [Selecting the Transformation Type](#)
- [Defining What to Transform](#)
- [Defining the Transformation Rule](#)
- [Using the Global Transformation Rules List](#)

Limitations

- Transformations are not supported for calculating columns of Right-to-Left languages.
- The only supported transformation for LOB/CLOB data types is to drop the column on the target.

Starting the New Transformation Rule Wizard

You define a rule for global transformation using the New Transformation Rule wizard. The transformation affects all of the tables in the task as you define them using the wizard.

To start the New Transformation Rule wizard

1. From the Task view, select the task where you want to create a global transformation, and click **Open**.

You can click **Open** from the top of the Tasks list or you can double-click the Task.

2. In the Single-Task view, make sure you are in the Design mode. If you are not in the Design Mode, click **Design** at the top right of the screen. For more information on the Task View and the Design Mode, see [Design Mode](#) in [The Tasks View](#).

3. In the Design mode, click **Global Settings**.

The [Global Transformation Rules](#) window opens.

4. From the top of the Global Transformation Rules window, click **New Global Transformation**.

The New Transformation Rules wizard opens.

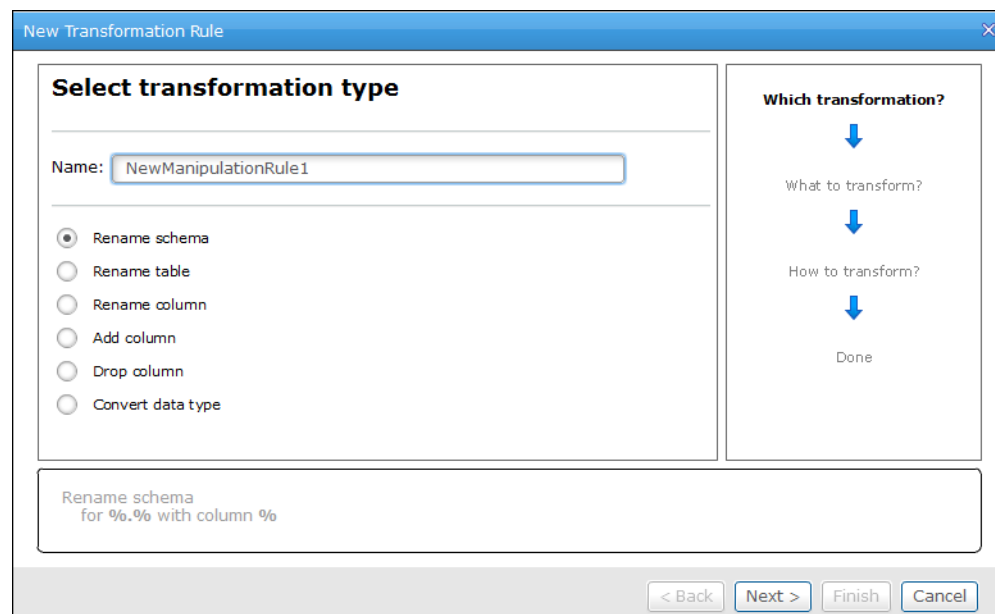
5. Enter the information to define a global transformation rule. The first step is [Selecting the Transformation Type](#).

Selecting the Transformation Type

In the **Select transformation type** step, you define the type of transformation you are carrying out.

The following figure shows the **Select transformation type** page:

Figure 15–5 Select Transformation Type Page



In this page, select one of the following:

- **Rename Schema:** Select this if you want to change the schema name for multiple tables. For example, if you want all HR tables to be renamed PERS.
- **Rename Table:** Select this if you want to change the name of multiple tables. For example, if you want all tables named SALARY to be called WAGES.
- **Rename Column:** Select this if you want to change the name of multiple columns. For example, if you want to change all columns with word MINIMUM to MIN.
- **Add Column:** Select this if you want to add a column with a similar name to multiple tables.
- **Drop Column:** Select this if you want to drop a column with a similar name from multiple tables.
- **Convert Data Type:** Select this if you want to change a specific data type to a different one across multiple tables. For example, if you want to change all Integer data types to a string.

After you complete the Selecting the Transformation Type step, click **Next** to go to the [Defining What to Transform](#) step.

Note: You can only create one rule for each transformation type. If you create multiple rules for a single transformation type, only the last rule you create will be valid.

For example, if you create the following rules (in order) to rename a schema:

Rename Schema: Add Prefix

Rename Schema: Add Suffix

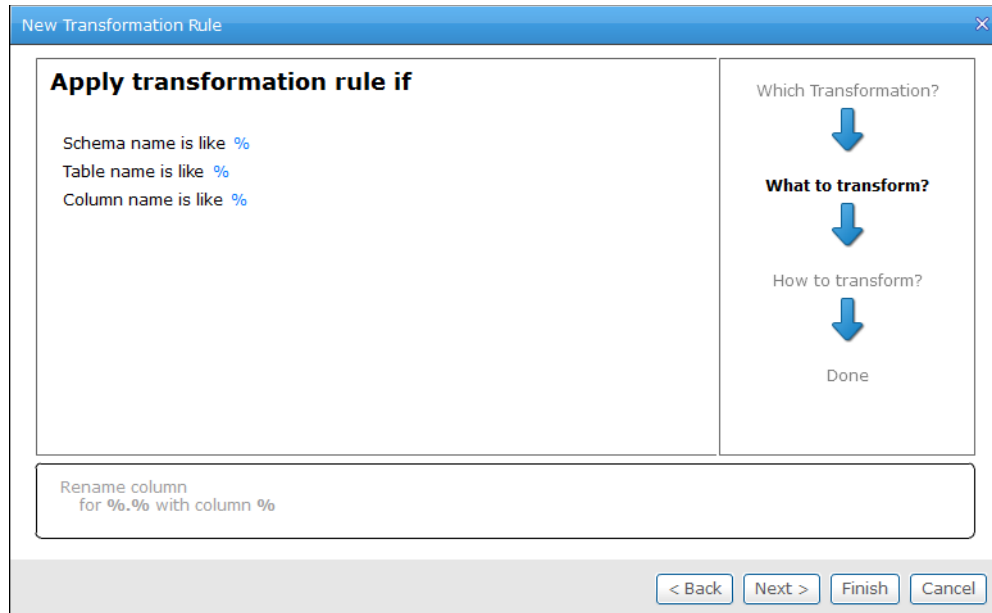
Only the second rule (adding a suffix to the schema name) will be executed.

Defining What to Transform

In the **Apply transformation rule** step, you define to which tables the transformation rule is applied. For example, you can apply the rule to all tables that contain the word SALARY as part of its name.

The following figure shows the **Apply transformation rule** page:

Figure 15–6 Apply Transformation Rule Page



Note: The above figure is for example purposes only. The options displayed on this page depend on the [Transformation Type](#) selected.

The following table describes the options available on this page:

Table 15–4 Apply Transformation Rule if...

Option	Available when transformation type is:	Description
Schema name is like %	Always	<p>Leave the % sign to include all schemas in your global transformation.</p> <p>Click the % sign to add a filter. In this case you can enter any name combination to include only that schema in your global transformation rule.</p> <p>For example, enter HR to include only tables that have the schema HR.</p> <p>You can use the % sign as a wild card. For example, H% includes all tables with a schema that begins with the letter H, such as HR, HELLO, or HQ.</p> <p>The % wildcard can be used in any position. For example, if you use it at the beginning, %H, then all table names that end in H are included in the transformation rule. The % can also be used in a middle position.</p> <p>Note: If you are using an Oracle target, you must enter a schema that exists on the target database. Amazon RDS Migration Tool does not create new schemas on an Oracle database. If you want to use a new schema for the target, create the schema on the Oracle database before running the task. For more information, see <link></p>
Table name is like %	Always	<p>Leave the % sign to include all table names in your global transformation rule.</p> <p>Click the % sign to add a filter. In this case you can enter any name combination to include only tables with that specific name in your global transformation rule.</p> <p>You can use the % sign as a wild card. For example, J% includes all tables with a name that begins with the letter J, such as JOBS, JOBS_HISTORY, or JACKSONVILLE.</p> <p>The % wildcard can be used in any position. For example, if you use it at the beginning, %H, then all table names that end in H are included in the transformation rule. The % can also be used in a middle position.</p>

Table 15–4 (Cont.) Apply Transformation Rule if...

Option	Available when transformation type is:	Description
Column name is like %	Rename Column Drop Column Convert Data Type	<p>Leave the % sign to include all column names in your global transformation rule.</p> <p>Click the % sign to add a filter. In this case you can enter any name combination to include only columns with that specific name in your global transformation rule.</p> <p>You can use the % sign as a wild card. For example, N% includes all columns with a name that begins with the letter N, such as NAME, NAME_FIRST, or NAME_LAST .</p> <p>The % wildcard can be used in any position. For example, if you use it at the beginning, %IES, then all column names that end in with the string "IES" are included in the transformation rule. The % can also be used in a middle position.</p>
Data type is	Convert Data Type	<p>Select a new Amazon RDS Migration Tool data type from the drop-down list. Make sure that the data type you select is compatible with the data in that column. For a description of Amazon RDS Migration Tool data types, see Amazon RDS Migration Tool Data Types.</p> <p>For information about data type mapping from the native database to Amazon RDS Migration Tool data types, see the chapter for the database you are using. For a list of databases supported by Amazon RDS Migration Tool, see Databases Supported as Amazon RDS Migration Tool Endpoints.</p>

After you complete the Selecting the Transformation Type step, click **Next** to go to the [Defining the Transformation Rule](#) step.

Note: If the global transformation type you are defining is Drop Column, you do not need to create a [Transformation Rule](#). In this case, click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Defining the Transformation Rule

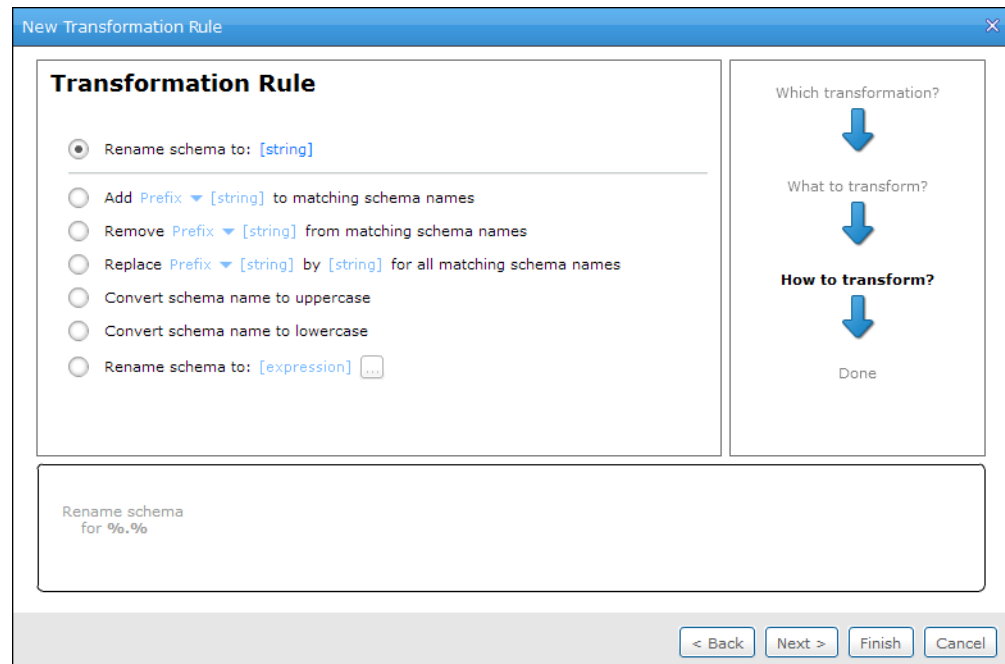
In the **Transformation rule** step, you define what happens to the tables that the transformation rule is applied to. For example, you can define a new name for the affected tables or add a prefix to the table names. For more information on defining the affected tables, see [Defining What to Transform](#).

Limitations

- Transformations are not supported for calculating columns of Right-to-Left languages.
- The only supported transformation for LOB/CLOB data types is to drop the column on the target.

The following figure shows the **Transformation rule** page:

Figure 15–7 Transformation Rule Page



Note: The above figure is for example purposes only. The options displayed on this page depend on the [Transformation Type](#) selected.

Depending on the [Transformation Type](#) selected, you define the rule to be carried out using the options on this page. See the section for any of the following transformation types you are using.

- [Rename Schema](#)
- [Rename Table](#)
- [Rename Column](#)
- [Add Column](#)
- [Drop Column](#)
- [Convert Data Type](#)

Rename Schema

If your transformation type is **Rename Schema**, you can do the following:

- [Rename schema to \(string\):](#)
- [Add a Prefix or Suffix](#)
- [Remove a Prefix or Suffix](#)
- [Replace a Prefix or Suffix with Different Characters](#)
- [Convert schema name to uppercase](#)
- [Convert schema name to lowercase](#)
- [Rename schema \(expression\)](#)

Rename schema to (string): Use the **Rename schema to: [string]** option to change the name of all table schemas that you defined in the [Apply transformation rule](#) step to a different name. For example, if you have a schema called `Human_Resources` and want to change all instances of this name to `HR` then enter the string `HR`. You can enter any string in this field.

Add a Prefix or Suffix Use the **Add a prefix or suffix** option to add additional characters to the beginning or end of the schema name for all schemas that fit the definition you created in the [Apply transformation rule](#) step. For example, if the schema name is `HR`, you can add a suffix, such as `TAR` or `_TAR` to the schema name for all tables with that schema name. In this case, the resulting schema name will be `HRTAR` or `HR_TAR`.

To globally add a prefix or suffix

1. Select **Add <Prefix/Suffix> Insert Characters to matching schema names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click **[string]** to activate the field.
4. Type the characters you want as the prefix or suffix. If you want to include an underscore or other legal character to separate the prefix/suffix from the original name, you must add it as part of the character string.
5. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Note: If you are using Oracle as your target database, Amazon RDS Migration Tool does not create a new schema. Therefore, the schema name that is the result of adding a prefix or suffix must exist in the Oracle target database. If the resulting schema name does not exist, you must create the schema in the Oracle database before carrying out this task.

For more information, see [Limitations](#) in [Chapter 5, "Using an Oracle Database as a Source or Target"](#).

Remove a Prefix or Suffix Use the **Remove a prefix or suffix** option to remove a string of characters from the beginning or end of a schema name for all schema that fit the definition you created in the [Apply transformation rule](#) step.

For example, you can use this option to remove the letters `_REV` from the schema name for all tables in the schema `HR_REV`. In this case the schema name in the target will be `HR`.

To globally remove a prefix or suffix

1. Select **Remove <Prefix/Suffix> Insert Characters from matching schema names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click **[string]** to activate the field.
4. Type the characters you want to remove. If you want to remove an underscore or other legal character from the original name, you must add it as part of the character string.
5. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Note: If you are using Oracle as your target database, Amazon RDS Migration Tool does not create a new schema. Therefore, the schema name that is the result of removing a prefix or suffix must exist in the Oracle target database. If the resulting schema name does not exist, you must create the schema in the Oracle database before carrying out this task.

For more information, see Oracle [Limitations](#).

Replace a Prefix or Suffix with Different Characters Use the **Replace a prefix or suffix** option to replace a string of characters with a different string of characters. You determine whether to replace the characters at the beginning or end of a schema name for all schema that fit the definition you created in the [Apply transformation rule](#) step.

For example, you can use this option to replace the letters `_ORIG` with `_REPL` in the schema name for all tables in the schema `HR_ORIG`. In this case the schema name in the target will be `HR_REPL`.

To globally replace a prefix or suffix

1. Select **Replace <Prefix/Suffix> Insert Characters by Insert Characters for all matching schema names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click the first **[string]** to activate the field.
4. Type the characters from the existing (source) schema that you want to replace. If you want to include an underscore or other legal character from the original name in the string that you want to replace, you must add it as part of the character string.
5. Click the second **[string]** to activate the field.
6. Type the characters you want to use in the target. These characters replace the original (source) characters in the target.
7. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Note: If you are using Oracle as your target database, Amazon RDS Migration Tool does not create a new schema. Therefore, the schema name that is the result of replacing a prefix or suffix with a different string of characters must exist in the Oracle target database. If the resulting schema name does not exist, you must create the schema in the Oracle database before carrying out this task.

Convert schema name to uppercase Use the **convert to uppercase** option to convert all of the letters in a schema name to upper case. For example:

`Schema_cat`, becomes `SCHEMA_CAT`

`schema_cat`, becomes `SCHEMA_CAT`

`sChEMa_Cat`, becomes `SCHEMA_CAT`

To globally change the schema name to all uppercase

1. Select **Convert schema name to uppercase**.
2. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Convert schema name to lowercase Use the **convert to lowercase** option to convert all of the letters in a schema name to lower case. For example:

Schema_cat, becomes schema_cat

SCHEMA_CAT, becomes schema_cat

sChEMa_Cat, becomes schema_cat

To globally change the schema name to all uppercase

1. Select **Convert schema name to lowercase**.
2. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Rename schema (expression) Use the **Rename schema to [expression]** option to change the name of all table schemas that you defined in the [Apply transformation rule](#) step to a different name. For example, if you have a schema called Human_Resources and want to change all instances of this name to HR.

To globally change a schema name

1. Select **Rename schema to [expression]**
2. Click the button to the right of the **Rename schema** option to open the Expression Editor. For information on how to use the Expression Editor, see [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#). Then go to step 4.

or

Click **[expression]** to activate the field and continue with step 3.

3. Type an SQLite expression or a string (in quotes) to rename the schema. For example:

- "New_Schema"
- 'PREF_' || \$SCHEMA_NAME_VAR || '_SUFF'

You can use the following variables in the SQLite expression:

- \$SCHEMA_NAME_VAR
- \$TABLE_NAME_VAR
- \$COLUMN_NAME_VAR
- \$COLUMN_DATATYPE_VAR

4. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Note: If you are using Oracle as your target database, Amazon RDS Migration Tool does not create a new schema. Therefore, you must enter a schema name that exists on the database or you must create the schema in the Oracle database before carrying out this task.

For more information, see Oracle [Limitations](#).

Rename Table

If your transformation type is **Rename Table**, you can do the following:

- [Rename table to \(string\)](#):
- [Add a Prefix or Suffix](#)

- [Remove a Prefix or Suffix](#)
- [Replace a Prefix or Suffix with Different Characters](#)
- [Convert table name to uppercase](#)
- [Convert table name to lowercase](#)
- [Rename table \(expression\)](#)

Rename table to (string): Use the **Rename table to: [string]** option to change the name of all tables that you defined in the [Apply transformation rule](#) step to a different name. For example, if you have a table called EMPLOYEE and want to change all instances of this name to EMP then enter the string EMP. You can enter any string in this field.

Add a Prefix or Suffix Use the **Add a prefix or suffix** option to add additional characters to the beginning or end of the table name for all tables that fit the definition you created in the [Apply transformation rule](#) step. For example, if the table name is EMPLOYEEES, you can add a suffix, such as TAR or _TAR to the table name for all tables with that table name. In this case, the resulting table name will be EMPLOYEEESTAR or EMPLOYEEES_TAR.

To globally add a prefix or suffix

1. Select **Add <Prefix/Suffix> Insert Characters to matching table names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click **[string]** to activate the field.
4. Type the characters you want as the prefix or suffix. If you want to include an underscore or other legal character to separate the prefix/suffix from the original name, you must add it as part of the character string.
5. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Remove a Prefix or Suffix Use the **Remove a prefix or suffix** option to remove a string of characters from the beginning or end of a table name for all tables that fit the definition you created in the [Apply transformation rule](#) step.

For example, you can use this option to remove the letters _REV from the table name for all tables with the name EMPLOYEEES. In this case the table name in the target will be EMPLOYEEES.

To globally remove a prefix or suffix

1. Select **Remove <Prefix/Suffix> Insert Characters from matching table names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click **[string]** to activate the field.
4. Type the characters you want to remove. If you want to remove an underscore or other legal character from the original name, you must add it as part of the character string.
5. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Replace a Prefix or Suffix with Different Characters Use the **Replace a prefix or suffix** option to replace a string of characters with a different string of characters. You determine whether to replace the characters at the beginning or end of a table name for all tables that fit the definition you created in the [Apply transformation rule](#) step.

For example, you can use this option to replace the letters `_ORIG` with `_REPL` in the table names for all tables called `EMPLOYEE_ORIG`. In this case the table name in the target will be `EMPLOYEE_REPL`.

To globally replace a prefix or suffix

1. Select **Replace <Prefix/Suffix> Insert Characters by Insert Characters for all matching schema names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click the first **[string]** to activate the field.
4. Type the characters from the existing (source) schema that you want to replace. If you want to include an underscore or other legal character from the original name in the string that you want to replace, you must add it as part of the character string.
5. Click the second **[string]** to activate the field.
6. Type the characters you want to use in the target. These characters replace the original (source) characters in the target.
7. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Convert table name to uppercase Use the convert to uppercase option to convert a table name to all upper case. For example:

Table_cat, becomes TABLE_CAT
table_cat, becomes TABLE_CAT
taBlE_Cat, becomes TABLE_CAT

To globally change the table name to all uppercase

1. Select **Convert table name to uppercase**.
2. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Convert table name to lowercase Use the convert to lowercase option to convert a table name to all lower case. For example:

Table_cat, becomes table_cat
TABLE_CAT, becomes table_cat
taBlE_Cat, becomes table_cat

To globally change the table name to all lowercase

1. Select **Convert table name to lowercase**.
2. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Rename table (expression) Use the **Rename table to [expression]** option to change the name of all tables that fit the definition you created in the [Apply transformation rule](#) step. For example, if you have a table called `EMPLOYEE` and want to change all instances of this name as defined in the previous step it to `EMP`.

To change the table name

1. Select **Rename table to: [expression]**

2. Click the button to the right of the **Rename table** option to open the Expression Editor. For information on how to use the Expression Editor, see [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#). Then go to step 4.

or

Click **[expression]** to activate the field and continue with step 3.

3. Type an SQLite expression or a string (in quotes) to rename the table. For example:
 - "New_Table"
 - 'PREFIX'|'\$TABLE_NAME_VAR'|'_SUFFIX'

You can use the following variables in the SQLite expression:

- \$SCHEMA_NAME_VAR
- \$TABLE_NAME_VAR
- \$COLUMN_NAME_VAR
- \$COLUMN_DATATYPE_VAR

Rename Column

If your transformation type is **Rename Column**, you can do the following:

- [Rename column to \(string\)](#):
- [Add a Prefix or Suffix](#)
- [Remove a Prefix or Suffix](#)
- [Replace a Prefix or Suffix with Different Characters](#)
- [Convert column name to uppercase](#)
- [Convert column name to lowercase](#)
- [Rename Column \(expression\)](#)

Rename column to (string): Use the **Rename column to: [string]** option to change the name of all columns that you defined in the [Apply transformation rule](#) step to a different name. For example, if you have a table called `SALARY` and want to change all instances of this name to `EMP` then enter the string `SAL`. You can enter any string in this field.

Add a Prefix or Suffix Use the **Add a prefix or suffix** option to add additional characters to the beginning or end of the column name for all columns that fit the definition you created in the [Apply transformation rule](#) step. For example, if the column name is `SALARY`, you can add a suffix, such as `TAR` or `_TAR` to the table name for all tables with that table name. In this case, the resulting table name will be `SALARYTAR` or `SALARY_TAR`.

To globally add a prefix or suffix

1. Select **Add <Prefix/Suffix> Insert Characters to matching column names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click the **[string]** to activate the field.

4. Type the characters you want as the prefix or suffix. If you want to include an underscore or other legal character to separate the prefix/suffix from the original name, you must add it as part of the character string.
5. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Remove a Prefix or Suffix Use the **Remove a prefix or suffix** option to remove a string of characters from the beginning or end of a column name for all columns that fit the definition you created in the [Apply transformation rule](#) step.

For example, you can use this option to remove the letters `_REV` from the column name for all columns with the name `SALARY`. In this case the column name in the target will be `SALARY`.

To globally remove a prefix or suffix

1. Select **Remove <Prefix/Suffix> Insert Characters from matching column names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click **[string]** to activate the field.
4. Type the characters you want to remove. If you want to remove an underscore or other legal character from the original name, you must add it as part of the character string.
5. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Replace a Prefix or Suffix with Different Characters Use the **Replace a prefix or suffix** option to replace a string of characters with a different string of characters. You determine whether to replace the characters at the beginning or end of a column name for all columns that fit the definition you created in the [Apply transformation rule](#) step.

For example, you can use this option to replace the letters `_ORIG` with `_REPL` in the column names for all columns called `SALARY_ORIG`. In this case the column name in the target will be `SALARY_REPL`.

To globally replace a prefix or suffix

1. Select **Replace <Prefix/Suffix> Insert Characters by Insert Characters for all matching schema names**.
2. Click the word **Prefix** or **Suffix** and select one of these two from the list.
3. Click the first **[string]** to activate the field.
4. Type the characters from the existing (source) column that you want to replace. If you want to include an underscore or other legal character from the original name in the string that you want to replace, you must add it as part of the character string.
5. Click the second **[string]** to activate the field.
6. Type the characters you want to use in the target. These characters replace the original (source) characters in the target.
7. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Convert column name to uppercase Use the **convert to uppercase** option to convert a column name to all upper case. For example:

`Column_cat`, becomes `COLUMN_CAT`

column_cat, becomes COLUMN_CAT

coLUMnM_Cat, becomes COLUMN_CAT

To globally change the table name to all uppercase

1. Select **Convert column name to uppercase**.
2. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Convert column name to lowercase Use the convert to lowercase option to convert a column name to all lower case. For example:

Column_cat, becomes column_cat

column_cat, becomes column_cat

coLUMnM_Cat, becomes column_cat

To globally change the column name to all lowercase

1. Select **Convert column name to lowercase**.
2. Click **Finish** to add the rule to the [Global Transformation Rules](#) list.

Rename Column (expression) Use the **Rename column to [expression]** option to change the name of all tables that fit the definition you created in the [Apply transformation rule](#) step. For example, if you have a column called SALARY and want to change it to SAL.

To change the column name

1. Select **Rename column to: [expression]**
2. Click the button to the right of the **Rename column** option to open the Expression Editor. For information on how to use the Expression Editor, see [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#). Then go to step 4.

or

Click **[expression]** to activate the field and continue with step 3.

3. Type an SQLite expression or a string (in quotes) to rename the column. For example:

- "New_Column"
- 'PREF_' || \$COLUMN_NAME_VAR || '_SUFF'

You can use the following variables in the SQLite expression:

- \$SCHEMA_NAME_VAR
- \$TABLE_NAME_VAR
- \$COLUMN_NAME_VAR
- \$COLUMN_DATATYPE_VAR

Add Column

When you add a column to multiple tables, you must provide a name, define the data type for the column and define the data that the column contains. The column that you define here is added to all tables that fit the definition you created in the [Apply transformation rule](#) step.

The following describes the information you must enter in the Transformation rule page for adding a column.

- **Column name:** Click the **[string]** to activate the field. Type the name for the column in the field. A column with this name is added to all tables that fit the definition you created in the [Apply transformation rule](#) step.
- **Column data type:** Click the drop-down for a list of data types and select a new Amazon RDS Migration Tool data type from the drop-down list. Make sure that the data type you select is compatible with the data in that column. For a description of Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

For information about data type mapping from the native database to Amazon RDS Migration Tool data types, see the chapter for the database you are using. For a list of databases supported by Amazon RDS Migration Tool, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

- **Computation expression:** Click the button to the right of this field to open the Expression Editor or type an expression using SQLite operators to define the data in the column.

For information on how to use the Expression Editor to create an expression, see [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#).

For more information on creating expressions, see [Creating an Expression for Transforms](#) and [Using SQLite Syntax with Transforms](#).

Drop Column

This option does not require a transformation rule. For this option you complete the Global Transformation Rule after the [Defining What to Transform](#) step.

Convert Data Type

When you convert the data type for a column, use this page to select the data type you want to convert to. The data type that you define in this step is applied to all columns and tables that fit the definition you created in the [Apply transformation rule](#) step. Make sure that the data type you select is compatible with the data in columns you defined.

To select a converted data type

Select an Amazon RDS Migration Tool data type from the drop-down list. For a description of Amazon RDS Migration Tool data types, see [Amazon RDS Migration Tool Data Types](#).

For information about data type mapping from the native database to Amazon RDS Migration Tool data types, see the chapter for the database you are using. For a list of databases supported by Amazon RDS Migration Tool, see [Databases Supported as Amazon RDS Migration Tool Endpoints](#).

Using the Global Transformation Rules List

The Global Transformation Rules List lists all of the notification rules that are defined for the Amazon RDS Migration Tool instance you are working with. It has the following information:

- **Name:** The name of the Global Transformation rule.

- **Description:** The description of the rule.

You can do the following in this page:

- [Edit a Global Transformation Rule](#)
- [Delete a Global Transformation Rule](#)

Edit a Global Transformation Rule

You can make changes to any global transformation rule.

To edit a global transformation rule

1. From the [Global Transformation Rules](#) list, select the transformation you want to edit.
2. Click **Open** (at the top of the list).
or
Double-click the notification you want to edit.
The Edit Existing Transformation Rule wizard opens.
3. You can make changes when in the [Apply transformation rule](#) step and in the [Transformation rule](#) step.

Delete a Global Transformation Rule

You can delete a Global Transformation rule.

To edit a global transformation rule

1. From the [Global Transformation Rules](#) list, select the transformation you want to edit.
2. Click **Delete** (at the top of the list).
The transformation is deleted from the system and no longer is available in the Transformation Rule list.

Editing Global Transformation Rules

You can make changes to any Transformation rule.

To edit a transformation rule

1. From the [Global Transformation Rules](#) list select the transformation rule you want to edit.
2. Click **Open** (at the top of the list).
or
Double-click the transformation rule you want to edit.
The Edit Transformation Rule wizard opens.
3. Make any changes you need in the wizard. For information on how to work with each of the pages in the New Transformation Rule wizard, see [Defining Global Transformations](#).

Notes:

- You cannot change name of the transformation rule.
 - You cannot change the [Transformation Type](#).
-
-

Using the Expression Builder (for Filters, Transforms, and Global Transformations)

The Amazon RDS Migration Tool Expression Builder provides an easy way to build an expression. It provides you with easy access to the required elements for your expression without having to type out any information manually. You access the Expression Builder through the dialog boxes where you define [Filters](#), [Transformations](#), and global transformations when you do any of the following:

- [Rename Schema](#)
- [Rename Table](#)
- [Rename Column](#)

The following topics describe the Expression Builder:

- [Overview of the Expression Builder](#)
- [Build an Expression](#)
- [Evaluate an Expression](#)
- [Test an Expression](#)
- [Using Elements in the Expression Builder](#)

Overview of the Expression Builder

The following is an example of the Expression Builder with its four main parts shown. The Expression Builder you are working with may look different depending on whether you want to build an expression for a filter, a transformation, or a global transformation.

add it to the expression. You can also add operators from the Element Pane, **Operators** tab.

- **D. Test Expression Panel:** This panel displays the results of a test that you can run after you provide values to each of the parameters in your expression. For more information, see [Test an Expression](#).

Build an Expression

The first step in using the expression builder is to build an expression. The expression that you build is displayed in the top section of the right pane. You can open the Expression when:

- You define [Transformations](#) for a single table.
- You define [Filters](#) for a single table.
- Use the Global Transformations dialog box to [Rename Schema](#), [Rename Table](#), [Rename Column](#), or [Add Column](#).

To build an expression

1. In the Elements Pane, select any element you want to include in your expression. For information on the elements you can use in an expression, see [Functions](#).
2. Move the selected element into the Build Expression box by doing one of the following:
 - Select the element and then click the arrow at the right of the element.
 - Double-click the element.
 - Drag the element into the Build Expression box.
3. Continue to add elements as needed.

Note: To add operators to your expression, you can use the **Operator** tab in the Element pane or the Operator buttons at the top of the Build Expression panel or any combination of these. See [Operators](#) and [Operator Button Bar](#).

For example, to create an expression that will combine the first name and last name, do the following:

1. In the **Input Columns** tab drag the FIRST_NAME column into the **Build Expression** box.
2. Click the concatenate (| |) operator from the **Operator** bar at the top of the Build Expression box.
3. In the Input Columns tab drag the LAST_NAME column into the **Build Expression** box.

Operator Button Bar

The Operator Button bar is above the Build Expression box. It contains the most common operators so you can easily add them to an expression.

To use the Operator Button Bar

1. Click the space in the Build Expression box where you want to add the operator.
2. Click the operator you want to add. It is added to the expression.

The following operators are available in the Operator Button bar:

+ - * / != == || AND OR

For information on these operators, see [Operators](#).

Evaluate an Expression


You can evaluate an expression to determine its parameters and to determine whether the expression is valid.

To evaluate an expression

1. From the Expression Builder window, click [Build an Expression](#).
2. Click **Evaluate**.

If the expression is *not* valid, an error message is written in red at the bottom of the Expression Builder window. The following figure shows an example of an error message.

Figure 15–9 Error Message



If the expression is valid, the expression parameters are displayed in the **Parameter** column in the **Evaluate Expression** section. See [Figure 15–10, "Test Expression"](#).

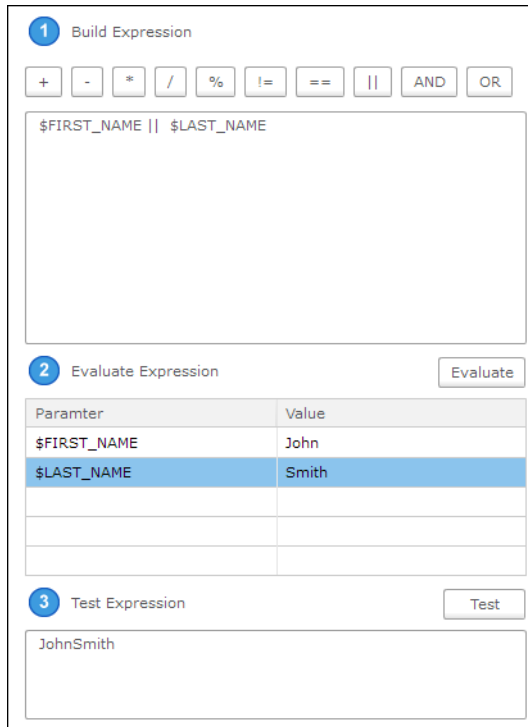
Type a valid value for each of the parameters in the **Value** column to [Test an Expression](#).

For example, type `John` for the `FIRST_NAME` and `Smith` for the `LAST_NAME` in the **Value** column. Once you type in values, you can [Test an Expression](#).

Test an Expression

You can use the Amazon RDS Migration Tool Test procedure to display the results of a test expression. The following figure is an example of a built expression that is evaluated and contains a test result.

Figure 15–10 Test Expression



The screenshot shows the Expression Builder interface with three main sections:

- 1 Build Expression:** A toolbar with operators (+, -, *, /, %, !=, ==, ||, AND, OR) and a text input field containing the expression `$FIRST_NAME || $LAST_NAME`.
- 2 Evaluate Expression:** An **Evaluate** button and a table showing parameter values.

Parameter	Value
<code>\$FIRST_NAME</code>	John
<code>\$LAST_NAME</code>	Smith
- 3 Test Expression:** A **Test** button and a text input field displaying the result `JohnSmith`.

To test an expression

1. From the Expression Builder window, [Build an Expression](#).
2. Click **Evaluate**. See [Evaluate an Expression](#) for more information.
3. View the parameters that are displayed. If your expression is not valid, an error message is displayed. See [Evaluate an Expression](#).
4. Type values for each parameter then click **Evaluate** to see the calculated expression.

For example, type `John` for `FIRST_NAME` and `Smith` for `LAST_NAME`. The result displayed is `JohnSmith`. If you want a space between the words add it to the end of the `FIRST_NAME` value or the beginning of the `LAST_NAME` value.

Using Elements in the Expression Builder

You can use the following types of elements to build expressions for transforms, filters, and global transformations. Select the appropriate tab to select the elements.

- [Input Columns \(Transforms and Filters only\)](#)
- [Metadata \(Global Transformations only\)](#)
- [Operators](#)
- [Functions](#)
- [Header Columns](#)

Input Columns (Transforms and Filters only)

This section lists the columns for the table you are working with. The table you are working with is the table you selected when you opened the Table Settings dialog box.

Metadata (Global Transformations only)

The Metadata tab contains the following variables that you can use in an expression:

AR_M_SOURCE_SCHEMA

AR_M_SOURCE_TABLE_NAME

AR_M_SOURCE_COLUMN_NAME

AR_M_SOURCE_COLUMN_DATATYPE

Operators

The sections below describe the SQLite operators you can use to build an expression with the Expression builder. The Expression builder divides the operators into the following categories:

- [Strings](#)
- [Logical](#)
- [Mathematical](#)

Strings

You can use the following string:

||

Name: Concatenate strings.

Examples:

FIRST_NAME || LAST_NAME

PHONE_NUMBER || <Office Only> (adds the string Office Only to the telephone number).

Logical

The following table describes the logical SQLite operators used by the Amazon RDS Migration Tool Expression Builder:

Table 15–5 Logical SQLite Operators used by Amazon RDS Migration Tool Expression Builder

Operator	Description
!= or <>	Is not equal to <code>\$\$SALARY!=100000</code>
IS	Is the same as <code>\$\$HIRE_DATE IS 2014-09-29</code> IS functions the same as = unless one or both of the operands are NULL. In this case, if both operands are NULL, then the IS operator evaluates to 1 (true). If one operand is NULL and the other is not, then the IS operator evaluates to 0 (false).
IS NOT	Is not the same as <code>\$\$HIRE_DATE IS NOT 2014-09-29</code> IS NOT functions the same as != unless one or both of the operands are NULL. In this case, if both operands are NULL, the IS NOT operator evaluates to 0 (false). If one operand is NULL and the other is not, then the IS NOT operator evaluates to 1 (true).

Table 15–5 (Cont.) Logical SQLite Operators used by Amazon RDS Migration Tool Expression Builder

Operator	Description
IN	<p>The IN operator takes a single scalar operand on the left and a vector operand on the right formed by an explicit list of zero or more scalars or by a single subquery. When the right operand of an IN operator is a subquery, the subquery must have a single result column. When the right operand is an empty set, the result of IN is false regardless of the left operand and even if the left operand is NULL.</p> <p>Note: SQLite allows the parenthesized list of scalar values on the right-hand side of an IN operator to be an empty list but most other SQL database engines and the SQL92 standard require the list to contain at least one element.</p>
LIKE	<p>The LIKE operator does a pattern matching comparison. The operand to the right of the LIKE operator contains the pattern and the left operand contains the string to match against the pattern. A percent symbol ("%") in the LIKE pattern matches any sequence of zero or more characters in the string. An underscore ("_") in the LIKE pattern matches any single character in the string. Any other character matches itself or its lower/upper case equivalent. (By default SQLite only understands upper/lower case for ASCII characters. The LIKE operator is case sensitive by default for unicode characters that are beyond the ASCII range.</p> <p>For example, the expression 'a' LIKE 'A' is TRUE but 'æ' LIKE 'Æ' is FALSE.)</p> <p>LIKE can be preceded by the NOT keyword.</p>
GLOB	<p>The GLOB operator acts in the same way as the LIKE operator but uses the UNIX file globbing syntax for its wildcards. GLOB is case sensitive.</p> <p>GLOB can be preceded by the NOT keyword to invert the sense of the test. The infix GLOB operator is implemented by calling the function glob(Y,X) and can be modified by overriding that function.</p>
MATCH	<p>The MATCH operator is a special syntax for the match() application-defined function. The default match() function implementation raises an exception and is not really useful for anything. But extensions can override the match() function with more helpful logic.</p>
REGEXP	<p>The REGEXP operator is a special syntax for the regexp() user function. No regexp() user function is defined by default and so use of the REGEXP operator will normally result in an error message.</p>
AND	<p>Both operands are true.</p> <p><code>\$MANAGER_ID AND EMPLOYEE ID >100</code></p>
OR	<p>Either operand is true.</p> <p><code>\$MANAGER_ID OR EMPLOYEE ID >100</code></p>
<<	<p>Bitwise shift left.</p> <p><code>x << n</code></p> <p>A bitwise shift to the left of x by n bits.</p>
>>	<p>Bitwise shift right.</p> <p><code>x >> n</code></p> <p>A bitwise shift to the right of x by n bits.</p>
&	<p>Unary and</p>
	<p>Unary or</p>
<	<p>Is less than.</p> <p><code>\$SALARY<100000</code></p>
<=	<p>Is less than or equal to</p> <p><code>\$SALARY<=100000</code></p>

Table 15–5 (Cont.) Logical SQLite Operators used by Amazon RDS Migration Tool Expression Builder

Operator	Description
>	Is greater than \$SALARY>100000
>=	Is more than or equal to \$SALARY>=100000
= or ==	Is equal to \$SALARY=100000

Mathematical

The following table describes the mathematical SQLite operators used by the Amazon RDS Migration Tool Expression Builder:

Table 15–6 SQLite Mathematical Operators used by the Amazon RDS Migration Tool Expression Builder

Operator	Description
+	Adds two values together. DEPARTMENT_ID+100 (adds 100 to each ID number). Any column used in an expression with this operator must be a numeric data type.
-	Subtracts a value from another value. MANAGER_ID-100 (subtracts 100 from each ID number). Any column used in an expression with this operator must be a numeric data type.
%	Uses the remainder of a division expression as the value. %SALARY/7 (Divides the value of the Salary column by 7 and uses any remainder from the expression as the column value).
/	Divides one value into another. SALARY/.16 (Divides the value of the Salary column by .16. Note: If the two values in the division expression are integers (two NUMERIC columns with no digits after the decimal) and the result is a fractional value, the result returned will be 0.
*	SALARY*.16 (Multiplies the value of the Salary column by .16. This could be used to calculate taxes that are subtracted from a salary).

Functions

The sections below describe the SQLite functions you can use to build an expression with the Expression builder. The Expression builder divides the functions into the following categories:

- [Strings](#)
- [LOBs](#)
- [Numeric](#)
- [NULL check](#)
- [Date and Time](#)
- [Operation](#)
- [Data Enrichment](#)

- Other Functions

Strings

The following table describes the string functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–7 SQLite String Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
lower(x)	The lower(x) function returns a copy of string x with all characters converted to lower case. The default built-in lower() function works for ASCII characters only.
ltrim(x,y)	The ltrim(x,y) function returns a string formed by removing all characters that appear in y from the left side of x. If there is no value for y, ltrim(x) removes spaces from the left side of x.
replace(x,y,z)	The replace(x,y,z) function returns a string formed by substituting string z for every occurrence of string y in string x.
rtrim(x,y)	The rtrim(x,y) function returns a string formed by removing all characters that appear in y from the right side of x. If there is no value for y, rtrim(x) removes spaces from the right side of x.
substr(x,y,z)	The substr(x,y,z) function returns a substring of input string x that begins with the y-th character and which is z characters long. If z is omitted then substr(x,y) returns all characters through the end of the string x beginning with the y-th. The left-most character of x is number 1. If y is negative then the first character of the substring is found by counting from the right rather than the left. If z is negative then the abs(z) characters preceding the y-th character are returned. If x is a string then characters indices refer to actual UTF-8 characters. If x is a BLOB then the indices refer to bytes.
trim(x,y)	The trim(x,y) function returns a string formed by removing all characters that appear in y from both sides of x. If there is no value for y, trim(x) removes spaces from both sides of x.
upper(x)	The upper(x) function returns a copy of string x with all characters converted to upper case.

LOBs

The following table describes the LOB functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–8 SQLite Lob Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
hex(x)	The hex() function receives an argument as a BLOB and returns an upper-case hexadecimal string version of the BLOB content.
randomblob(N)	The randomblob(N) function returns an N-byte BLOB that contains pseudo-random bytes. If N is less than 1 then a 1-byte random BLOB is returned.
zeroblob(N)	The zeroblob(N) function returns a BLOB that consists of N bytes of 0x00.

Numeric

The following table describes the numeric functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–9 SQLite Numeric Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
abs(x)	The abs(x) function returns the absolute value of the numeric argument X. Abs(x) returns NULL if x is NULL. Abs(x) returns 0.0 if x is a string or BLOB that cannot be converted to a numeric value.
random()	The random() function returns a pseudo-random integer between -9223372036854775808 and +9223372036854775807.
round(x,y)	The round(x,y) function returns a floating-point value x rounded to y digits to the right of the decimal point. If there is no value for y, it is assumed to be 0.
max(x,y...)	The multi-argument max() function returns the argument with the maximum value, or returns NULL if any argument is NULL. The multi-argument max() function searches its arguments from left to right for an argument that defines a collating function and uses that collating function for all string comparisons. If none of the arguments to max() define a collating function, then the BINARY collating function is used. Note that max() is a simple function when it has two or more arguments but operates as an aggregate function if it has a single argument.
min(x,y...)	The multi-argument min() function returns the argument with the minimum value. The multi-argument min() function searches its arguments from left to right for an argument that defines a collating function and uses that collating function for all string comparisons. If none of the arguments to min() define a collating function, then the BINARY collating function is used. Note that min() is a simple function when it has two or more arguments but operates as an aggregate function if it has a single argument

NULL check

The following table describes the NULL check functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–10 SQLite NULL Check Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
coalesce(x,y...)	The coalesce() function returns a copy of its first non-NULL argument, it returns NULL if all arguments are NULL. Coalesce() have at least two arguments.
ifnull(x,y)	The ifnull() function returns a copy of its first non-NULL argument, it returns NULL if both arguments are NULL. Ifnull() must have exactly two arguments. The ifnull() function is the same as coalesce() with two arguments.
nullif(x,y)	The nullif(x,y) function returns a copy of its first argument if the arguments are different and returns NULL if the arguments are the same. The nullif(x,y) function searches its arguments from left to right for an argument that defines a collating function and uses that collating function for all string comparisons. If neither argument to nullif() defines a collating function then the BINARY is used.

Date and Time

The following table describes the Date and Time functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–11 SQLite Date and Time Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
<code>date(timestring, modifier, modifier...)</code>	Returns the date in the format YYYY-MM-DD.
<code>time(timestring, modifier, modifier...)</code>	Returns the time in the format HH:MM:SS.
<code>datetime(timestring, modifier, modifier...)</code>	Returns the date and time in the format YYYY-MM-DD HH:MM:SS.
<code>julianday(timestring, modifier, modifier...)</code>	The <code>julianday()</code> function returns the number of days since noon in Greenwich on November 24, 4714 B.C.
<code>strftime(format, timestring, modifier, modifier...)</code>	The <code>strftime()</code> routine returns the date formatted according to the format string specified as the first argument. It supports the following variables: %d: day of month %H: hour 00-24 %f: ** fractional seconds SS.SSS %j: day of year 001-366 %J: ** Julian day number %m: month 01-12 %M: minute 00-59 %s: seconds since 1970-01-01 %S: seconds 00-59 %w: day of week 0-6 sunday==0 %W: week of year 00-53 %Y: year 0000-9999 %?: %

Time strings can be in the following formats:

- YYYY-MM-DD
- YYYY-MM-DD HH:MM
- YYYY-MM-DD HH:MM:SS
- YYYY-MM-DD HH:MM:SS.SSS
- YYYY-MM-DDTHH:MM (T is a literal character that separates the date and time)
- YYYY-MM-DDTHH:MM:SS (T is a literal character that separates the date and time)
- YYYY-MM-DDTHH:MM:SS.SSS (T is a literal character that separates the date and time)
- HH:MM
- HH:MM:SS
- HH:MM:SS.SSS
- now (Converted to current date and time using UTC)
- DDDD.DDDD (The Julian day number expressed as a floating point value).

Data Enrichment

Data Enrichment functions allow the selected source tables to be augmented with data from other records located in either the source or target databases. Practical applications of data enrichment functions include code lookup or master record lookup (e.g. social security number lookup to find a person's name).

You can enrich the target tables with supplemental data retrieved from the source or target database by defining a transformation on the table. For more information about defining transformations on a single table, see [Define Transformations on a Single Table](#).

Limitations

Amazon Redshift is not supported.

Data Enrichment Functions

The table below describes the source and target lookup functions, which can be used both for table transformations and for global transformations. For a description of the parameters available for these functions, see [Input Parameters](#).

Table 15–12 SQLite Data Enrichment Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
<code>source_lookup(TTL, 'SCHM', 'TBL', 'EXP', 'COND', COND_PARAMS)</code>	Use to retrieve additional data from the source database.
<code>target_lookup(TTL, 'SCHM', 'TBL', 'EXP', 'COND', COND_PARAMS)</code>	Use to retrieve additional data from the target database.

Input Parameters

The possible input parameters for the lookup functions are described in the table below. For a usage example, see [Data Enrichment Example](#).

Table 15–13 Lookup Input Parameters for Data Enrichment Functions

Parameter	Description
TTL	TTL (Time to Live) is the amount of time the 'COND' return value will be cached. Caching the 'COND' return value improves performance by reducing the frequency that Amazon RDS Migration Tool needs to access the source/target database. As there is no default, you must specify a TTL value, which can be one of the following: <SECONDS> - The time to cache the 'COND' return value in seconds. Specify a short caching time (e.g. 3) for data that is frequently updated or a long caching time for data that rarely changes. 'NO_CACHING' - Specify 'NO_CACHING' if you do not want to cache the 'COND' return value. This is recommended for data that is constantly updated (e.g. share prices). 'NO_EXPIRATION' - For data that is never updated (e.g. a street name), specify 'NO_EXPIRATION' to store the 'COND' return value permanently in the cache.
'SCHM'	The schema name.
'TBL'	The table on which to perform the lookup.

Table 15–13 (Cont.) Lookup Input Parameters for Data Enrichment Functions

Parameter	Description
'EXP'	<p>The expression to retrieve data from the lookup table.</p> <p>Note: The expression syntax must be native to the database it accesses. The result should be a single column. Possible expressions include: <code>col1</code>, <code>col1+5</code>, <code>max(col1)</code>.</p> <p>Note: Full LOB columns are not supported. For information on including Limited-size LOB columns in the replication, see the description of the Metadata tab.</p> <p>Input Columns (Transforms and Filters only), Header Columns, and Metadata (Global Transformations only) can also be used in the expression and are evaluated before the lookup statement is performed against the database.</p>
'COND'	<p>The condition for the lookup statement.</p> <p>Note: The condition syntax must be native to the database it accesses. The COND is a single field referencing all required fields.</p> <p>Example if the lookup table is located in Oracle:</p> <pre>'Fieldname1=:1 and Fieldname2=:2 and Fieldname3 =:3'</pre> <p>Example if the lookup table is located in Microsoft SQL Server:</p> <pre>'Fieldname1=? and Fieldname2=? and Fieldname3=?'</pre> <p>Input Columns (Transforms and Filters only), Header Columns, and Metadata (Global Transformations only) can also be used in the condition and are evaluated before the lookup statement is performed against the database.</p>
COND_PARAMS	<p>Any parameters required by the COND parameter.</p> <p>The COND_PARAMS (condition parameters) is not a single field, but a list of fields.</p> <p>Syntax:</p> <pre>\$FIELDNAME1 , \$FIELDNAME2 , \$FIELDNAME3</pre> <p>Full example:</p> <pre>source_lookup(10000 , 'HR' , 'DEPARTMENTS' , 'DEPARTMENT_NAME' , 'COMPANY_ID=? and DIVISION_ID=? and DEPT_ID=?' , \$COMP_ID , \$DIV_ID , \$DEPT_ID)</pre>

Note: To improve efficiency, the source/target lookup tables should be indexed for the specified lookup fields.

Data Enrichment Example

In the following example, Mike needs to add the `DEPARTMENT_NAME` column to the `HR.JOB_HISTORY` table. The `DEPARTMENT_NAME` column is located in the `HR.DEPARTMENTS` table in the source database.

This is how the `HR.JOB_HISTORY` table appears before the column is added:

EMPLOYEE_ID	START_DATE	END_DATE	JOB_ID	DEPARTMENT_ID
102	2001-01-13 00:00:00	2006-07-24 00:00:00	IT_PROG	60
101	1997-09-21 00:00:00	2001-10-27 00:00:00	AC_ACCOUNT	110
101	2001-10-28 00:00:00	2005-03-15 00:00:00	AC_MGR	110
201	2004-02-17 00:00:00	2007-12-19 00:00:00	MK_REP	20
114	2006-03-24 00:00:00	2007-12-31 00:00:00	ST_CLERK	50
122	2007-01-01 00:00:00	2007-12-31 00:00:00	ST_CLERK	50
200	1995-09-17 00:00:00	2001-06-17 00:00:00	AD_ASST	90
176	0001-09-09 09:30:25	2010-09-09 09:30:00	SA_REP	60
176	2007-01-01 00:00:00	2010-09-09 09:30:00	SA_MAN	60
200	2002-07-01 00:00:00	2006-12-31 00:00:00	AC_ACCOUNT	90

To add the `DEPARTMENT_NAME` column, Mike needs to:

1. Create a new task and select the `HR.JOB_HISTORY` table for replication.
2. Apply a “New Column” transformation to the `HR.JOB_HISTORY` table. For more information on defining transforms, see [Define Transformations on a Single Table](#).
3. Open the **Expression Builder** and choose **Data Enrichment** from the **Functions** tab. For more information on the Expression Builder, see [Using the Expression Builder \(for Filters, Transforms, and Global Transformations\)](#).
4. Select the `source_lookup` function and configure it as follows (using the native syntax of the source database):

If the lookup table is located in Oracle:

```
source_lookup(10000, 'HR', 'DEPARTMENTS', 'DEPARTMENT_NAME',
'DEPARTMENT=:1', $DEPARTMENT_ID)
```

If the lookup table is located in Microsoft SQL Server:

```
source_lookup
(10000, 'HR', 'DEPARTMENTS', '[DEPARTMENT_NAME]',
'[DEPARTMENT]=?' , $DEPARTMENT_ID)
```

Where:

10000 is the TTL parameter.

HR is the schema name.

DEPARTMENTS is the table name.

DEPARTMENT_NAME is the expression.

DEPARTMENT_ID=:1 (or ? on Microsoft SQL Server) is the condition.

\$DEPARTMENT_ID is the condition parameter.

5. Run the task.

This is how the HR.JOB_HISTORY table appears after the Full Load completes.

EMPLOYEE_ID	START_DATE	END_DATE	JOB_ID	DEPARTMENT_ID	DEPARTMENT_NAME
102	2001-01-13 00:00:00	2006-07-24 00:00:00	IT_PROG	60	Warehouse
101	1997-09-21 00:00:00	2001-10-27 00:00:00	AC_ACCOUNT	110	ProductManagment
101	2001-10-28 00:00:00	2005-03-15 00:00:00	AC_MGR	110	ProductManagment
201	2004-02-17 00:00:00	2007-12-19 00:00:00	MK_REP	20	HRoductManagment
114	2006-03-24 00:00:00	2007-12-31 00:00:00	ST_CLERK	50	Bookkeepingggment
122	2007-01-01 00:00:00	2007-12-31 00:00:00	ST_CLERK	50	Bookkeepingggment
200	1995-09-17 00:00:00	2001-06-17 00:00:00	AD_ASST	90	Managmentnggment
176	0001-09-09 09:30:25	2010-09-09 09:30:00	SA_REP	60	Warehousenggment
176	2007-01-01 00:00:00	2010-09-09 09:30:00	SA_MAN	60	Warehousenggment
200	2002-07-01 00:00:00	2006-12-31 00:00:00	AC_ACCOUNT	90	Managmentnggment

Operation

The following table describes the Operation functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–14 SQLite Operation Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
<code>operation_indicator(value_on_delete, value_on_update, value_on_insert)</code>	<p>When the <code>operation_indicator</code> function is invoked on its own or as part of an expression, records deleted from the source database will <i>not</i> be deleted from the target database. Instead, the corresponding target record will be flagged (with a user-provided value) to indicate that it was deleted from the source. The <code>operation_indicator</code> function also requires you to provide values to indicate records that were inserted or updated in the source database.</p> <p>Note: The <code>operation_indicator</code> function is not supported on tables that do not have a Primary Key.</p> <p>Note: It is recommended to add a dedicated column for the flag values, for example, <code>OPERATION</code>. For an explanation of how to add a column, see Using the Transform Tab.</p> <p>To specify the function values:</p> <p>Replace <code>value_on_delete</code>, <code>value_on_insert</code> and <code>value_on_update</code> with the values that you want to appear in the target database.</p> <p>Values should be formatted according to the corresponding column type.</p> <p>Example when the column type is <code>INT4</code>:</p> <pre>operation_indicator('1', '0', '0')</pre> <p>Example when the column type is <code>STRING</code>:</p> <pre>operation_indicator('Deleted', 'Updated', 'Inserted')</pre>

Other Functions

The following table describes additional functions used by the Amazon RDS Migration Tool Expression Builder:

Table 15–15 SQLite Functions used by the Amazon RDS Migration Tool Expression Builder

Function	Description
<code>length(x)</code>	<p>For a string value <code>x</code>, the <code>length(x)</code> function returns the number of characters (not bytes) in <code>x</code> before to the first NULL character.</p> <p>If <code>x</code> is NULL then <code>length(x)</code> is NULL. If <code>x</code> is numeric then <code>length(X)</code> returns the length of a string representation of <code>X</code>.</p>
<code>like(x,y,z)</code>	<p>The <code>like()</code> function is used to implement the "<code>Y LIKE X [ESCAPE Z]</code>" expression. The <code>ESCAPE (z)</code> clause is optional. If there is a <code>z</code> clause, then the <code>like()</code> function is invoked with three arguments. Otherwise, it is invoked with two arguments.</p>
<code>typeof(x)</code>	<p>The <code>typeof(x)</code> function returns a string that indicates the datatype of the expression <code>x</code>: <code>null</code>, <code>integer</code>, <code>real</code>, <code>text</code>, or <code>BLOB</code>.</p>

Header Columns

By default, header columns for source tables are not replicated to the target. You can determine which, if any, header columns to replicate when you define a transformation by creating an expression that includes the header field.

You can create a filter using header field values. Header column filters are applied during change processing. See [Using Filters](#) for additional information.

Note: The Header Column tab in the Expression builder is available for Filters and Transformations. It is available for Global Transformations only when you select **Add Columns**. See [Selecting the Transformation Type](#).

The following table describes the header field columns.

Table 15–16 Header Columns

Header Column Name	Value in Change Process	Value in Full Load	Data Type
AR_H_STREAM_POSITION	The stream position value from the source (For example, the SCN or LSN depending on the source database).	Empty string	STRING
AR_H_TIMESTAMP	Change timestamp	Current timestamp	DATETIME
AR_H_COMMIT_TIMESTAMP	Commit timestamp	Current timestamp	DATETIME
AR_H_OPERATION	INSERT/UPDATE	INSERT	STRING
AR_H_USER	The user name, ID or any other information that the source provides about the change initiator. This header column is supported on the Microsoft SQL Server, IBM DB2 on iSeries (AIS), and Oracle (version 11.2.0.3 and higher) source endpoints only.	Empty	STRING

Task Settings

Each task has settings that you can configure according to your needs for replication. You configure the settings in the Task Settings dialog box.

To open the Select Tables dialog box

1. Open the task you are working with if it is not displayed in the Amazon RDS Migration Console. For information on opening a task, see [Editing a Replication Task](#).
2. Click **Task Settings**.
3. In the Task Settings dialog box, select the tab on the left with the task setting you want to configure. The Task Settings dialog box has the following tabs:
 - [Metadata](#)
 - [Full Load](#)

- [Change Processing](#)
- [Error Handling](#)
- [Logging](#)

Metadata

When you click **Metadata** in the Task Settings dialog box, you can configure the **Target Metadata Settings** for a replication task.

Target Metadata

Target table schema: (if empty, use the schema from the source table): This will automatically add the owner prefix for the target database to all tables if no source schema is defined.

Include LOB columns in replication (BLOB, CLOB and similar large object datatypes): Select this for if you are using tables that include LOBs. For information on LOB support in Amazon RDS Migration Tool endpoints, see [Support for Large Object Data Types \(LOBs\)](#).

Note: LOB data types are supported only in tables that include a primary key.

If you select **Include LOB columns**, you must select one of the following:

- **Full LOB mode (slower unlimited LOB size):** If you select this option then enter a value for the following parameter:
 - Chunk size (KB):** Use the counter or type in the size of the LOB chunks to use when replicating the data to the target.
- **Limited-size LOB mode:** If you select this option then enter a value for the following parameter:
 - Max LOB size (KB):** Use the counter or type the maximum size (in kilobytes) for an individual LOB. Enter **0** for unlimited size.

Control Tables

Control Tables provide information about the replication task as well as useful statistics that can be used to plan and manage both the current replication task and future replication tasks. Aside from the **Apply Exceptions** table which is always created, users can choose which of the following Control Tables to create on the target:

- **Replication Status:** Provides details about the current task including task status, amount of memory consumed by the task, number of changes not yet applied to the target and the position in the source database from which Amazon RDS Migration Tool is currently reading.
- **Suspended Tables:** Provides a list of suspended tables as well as the reason they were suspended.
- **Replication History:** Provides information about the replication history including the number and volume of records processed during a replication task, latency at the end of a CDC task, among others.

For a detailed description of these tables, see [Control Tables](#).

Create Amazon RDS Migration Tool control table in target using schema: Enter the database schema for the RDS Migration Tool target Control Tables. If you do not enter any information in this field, then the tables are copied to the default location in the database.

Note: When this field is left empty, the target endpoint is MySQL, and the [Multiple databases](#) option is enabled, a default database named `amazon_control` will be created on the MySQL server. The selected control tables will be created in this database.

Replication history time slot (minutes): The length of each time slot in the Replication History table. The default is 5 minutes.

Table Selection

In addition to the **Apply Exceptions** table (required), select which of the following Control Tables you want Amazon RDS Migration Tool to create on the target database: **Replication Status**, **Suspended Tables** and **Replication History**.

Full Load

When you click **Full Load** in the Task Settings dialog box, you can configure the following:

- [Full Load Settings](#)
- [Full Load Tuning](#)

Full Load Settings

Click the **Full Load Settings** sub-tab to configure the following:

Full is ON/OFF.

Click this button to toggle full load on or off. The initial setting is determined when [Setting up Tasks](#).

When full load is ON, Amazon RDS Migration Tool loads the initial source data to the target database.

Note: Full load can be turned on or off at any stage even if change processing is on. Once the task begins to process changes, the full load on/off switch is used only as additional protection against accidental or unauthorized reload.

Target table preparation:

If target table already exists: Select one of the following from the list to determine how you want to handle loading the target at *full-load start up*:

- **DROP and Create table:** The table is dropped and a new table is created in its place.
- **TRUNCATE before loading:** Data is truncated without affecting the table metadata.
- **Do nothing:** Data and metadata of the target table are not affected.

Resuming an incomplete full load: This lets you set the policy for determining whether to allow tables to resume processing from the point of interruption during a full load. Allowing the restarting of tables from the point of interruption can cause a slowdown in the full-load processing. However, starting extremely large tables from the beginning may cause long delays in completing the full load. In addition, for tables to be restarted from the point of interruption, they must have a unique index. It is recommended that they have a clustered primary key.

For tables interrupted while in full load: Select how you want to handle loading tables that were interrupted during a full-load operation.

Use the following options to determine the policy for reloading tables after an interruption of a full-load process:

- **Allow resuming table full load:** Select this to allow tables loading when a full-load process is interrupted to begin loading from the point where the loading was interrupted. When you select this option, you must define which tables can be reloaded in the **Resume interrupted full load only for tables** option.
- **Always restart table full load:** Select this option to restart the loading for all tables even those that were partially loaded when the full-load process was interrupted.

Resume interrupted full load only for tables: This option is available when you select **Allow resuming table full load**. Use this option to define which tables should resume loading when the full-load operation begins again. You can define the following options:

- **Larger than (rows):** Use the arrows (or type) to select the minimum number of rows that a table must have to resume loading. Tables with fewer rows than the number entered in this field will restart the full load. The default value is 1000000.
- **That have:** You must select one of the following. Tables must have a clustered primary key or a unique index to be restarted automatically.
 - **A clustered primary key:** Only tables with a clustered primary key are reloaded. This is the default selection.
 - **Any unique index:** Tables with any type of unique index are reloaded. Note that this may slow down the full-load process.

Create primary key or unique index after full load completes: Select this option if you want to delay primary key or unique index creation until after full load completes.

Note: When this option is selected, resuming of incomplete Full Load tasks is not supported.

Stop the task after Full Load completes and: You can set the task to stop automatically after Full Load completes. This is useful if you need to perform DBA operations on the target tables before the task's CDC (Change Data Capture) phase begins.

During Full Load, any changes to the source tables are cached. When Full Load completes, the cached changes are automatically applied to the target tables.

Select either **Cached changes have not yet been applied** to stop the task before the cached changes are applied or **Cached changes have been applied** to stop the task after the cached changes are applied.

Selecting the **Cached changes have been applied** option will stop the task as soon as data is consistent across all tables in the task. Selecting the **Cached changes have not yet been applied** option will stop the task immediately after Full Load completes.

Note: Choosing to stop the task *before* cached changes have been applied may adversely affect performance, since the cached changes will only be applied to tables (even those that have already completed Full Load) *after* the last table completes Full Load.

Note: When working with the Amazon RDS Migration Tool File Channel endpoint, these options should be set in the remote File Channel task and not in the local File Channel task.

For more information on the Amazon RDS Migration Tool File Channel endpoint, see [Chapter 12, "Using the Amazon RDS Migration Tool File Channel"](#).

Full Load Tuning

Click the **Full Load Tuning** sub-tab to configure the following:

Tuning settings

- **Maximum number of tables to load in parallel:** Enter the maximum number of tables to load into the target at one time. The default value is **5**.
- **Transaction consistency timeout (seconds):** Enter the number of seconds that Amazon RDS Migration Tool waits for transactions close, if they are open when the task starts, before beginning the full-load operation. The default value is **600** (10 minutes). Amazon RDS Migration Tool will begin the full load after the timeout value is reached even if there are open transactions.
- **Commit rate during full load:** The maximum number of events that can be transferred together. The default value is **1000**.

Change Processing

When you click **Change Processing** in the Task Settings dialog box, you can configure the following:

- [Apply Changes Settings](#)
- [Changes Processing Tuning](#)

Apply Changes Settings

Click the **Apply Changes Settings** sub-tab to configure the following:

Apply Changes is ON/OFF.

Click this button to toggle Apply Changes (Change Processing) on or off. The initial setting is determined when [Setting up Tasks](#).

When Apply Changes is ON, Amazon RDS Migration Tool processes the changes. You can view the change processing in the Monitor. For more information, see [Monitoring Change Processing Operations](#).

Note: When you turn on apply changes you must reload the task or position back to the point of the previous reload.

DDL handling policy: Determine how to handle the target table for the change capture:

- **When source table is dropped:** Select one of the following:
 - DROP target table
 - Ignore Drop
- **When source table is truncated**
 - TRUNCATE target table
 - Ignore TRUNCATE
- **When source table is altered**
 - ALTER target table
 - Ignore ALTER

Changes Processing Tuning

Click the **Change Processing Tuning** sub-tab to fine-tune the Apply Changes settings.

Change processing mode

Determine which method will be used to apply changes.

- **Transactional apply:** Select this to apply each transaction individually, in the order it is committed. In this case, strict referential integrity is ensured for all tables.
- **Batch optimized apply (TurboStream CDC):** Select this to commit the changes in batches. In this case, a pre-processing action occurs to group the transactions into batches in the most efficient way. This may affect transactional integrity. Therefore, you must select one of the following to determine how the system will handle referential integrity issues:
 - Preserve transactional integrity
 - Allow temporary lapses in transactional integrity to improve performance

Note: The SAP Sybase ASE target database does not support applying binary data types in Batch Optimized Apply mode.

Note: When LOB columns are included in the replication, **Batch optimized apply (TurboStream CDC)** can only be used in **Limited-size LOB mode**. For more information about including LOB columns in the replication, see Task Settings | [Metadata](#).

Batch tuning

The following options are available when **Batch optimized apply (TurboStream CDC)** **Change processing mode** is selected:

Apply batched changes in intervals:

- **Longer than:** The minimum amount of time to wait between each application of batch changes. The default value is 1.
- **But less than:** The maximum amount of time to wait between each application of batch changes (before declaring a timeout). The default value is 30.
- **Force apply a batch when processing memory exceeds (MB):** The maximum amount of memory to use for pre-processing in **Batch optimized apply mode**. The default value is 500.
- **Limit the number of changes applied per change processing statement to:** To limit the number of changes applied in a single change processing statement, select this check box and then optionally change the default value. The default value is 10,000.

The following options are available when the **Transactional apply Change processing mode** is selected.

- **Minimum number of changes per transaction:** The minimum number of changes to include in each transaction. The default value is 1000.
- **Maximum time to batch transactions before applying (seconds):** The maximum time to collect transactions in batches before declaring a timeout. The default value is 60.

Transaction offload tuning

The following tuning options are available, regardless of which **Change processing mode** is selected:

Offload transaction in progress to disk if:

Amazon RDS Migration Tool usually keeps transaction data in memory until it is fully committed to the source and/or target. However, transactions that are larger than the allocated memory or that are not committed within the specified time limit will be offloaded to disk.

- **Transaction memory size exceeds (MB):** The maximum size that all transactions can occupy in memory before being offloaded to disk. The default value is 1000.
- **Transaction duration exceeds (seconds):** The maximum time that each transaction can stay in memory before being offloaded to disk. The duration is calculated from the time that Amazon RDS Migration Tool started capturing the transaction. The default value is 60.

Miscellaneous tuning

Statements cache size (number of statements): The maximum number of prepared statements to store on the server for later execution (when applying changes to the target). The default is 50. The maximum is 200.

Error Handling

When you click **Error Handling** in the Task Settings dialog box, you can configure the following:

- [Error Handling Settings](#)
- [Environmental Errors](#)
- [Data Error](#)
- [Table Error](#)

- [Apply Conflicts](#)

For more information on error handling in Amazon RDS Migration Tool, see [Error Handling](#).

Error Handling Settings

In each of the Error Handling sub-tabs, you can decide whether to use the Server Error Handling settings (the default) or override them.

To override the Server Error Handling settings:

1. Click the **Change** button in any of the **Error Handling** sub-tabs.
A window opens allowing you to override the server error handling policy.
2. Click **Create task specific error handling policy**.
3. Click **Change Policy**.

To revert to the Server Error Handling settings:

1. Click the **Change** button in any of the **Error Handling** sub-tabs.
A window opens allowing you to disable the task-specific error handling policy.
2. Click **Global server error handling policy**.
3. Click **Change Policy**.

For information on setting the error policy for the server, see [Error Handling](#) in the [Amazon RDS Migration Server Settings](#) chapter.

Environmental Errors

Click the **Environmental Errors** sub-tab to configure the following:

- **Maximum retry count:** Select this option and then specify the maximum number of attempts to restart a task when an environmental error occurs.
Specify "0" to never restart a task.
When the check box is not selected, Amazon RDS Migration Tool will attempt to restart the task an infinite number of times.
When the system attempts to restart the task the designated number of times, the task is stopped and manual intervention is required.
- **Interval between retry attempts:** Use the counter to select or type the number of seconds that the system waits between attempts between attempts to restart a task.
- **Increase retry interval for long outages:** Select this check box to increase the retry interval for long outages. When this option is enabled, the number of seconds between retry attempts increases each time.
- **Maximum retry interval:** Use the counter to select or type the number of seconds to wait between attempts to restart a task when the **Increase retry interval for long outages** option is enabled.

For information about environmental errors and the configuration properties, see [Environmental Errors](#) and [Error Handling Properties](#) in the [Error Handling](#) appendix.

Data Error

Click the **Data Error** sub-tab to configure the following:

- **Policy:** Click the triangle to open the list and select what happens when an error occurs in one or more specific records. You can select one of the following from the list:
 - **Ignore record:** The task continues and the error is ignored.
 - **Log error** (default): The task continues and the error is written to the task log.
 - **Suspend table:** The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.
- **Escalation:** Select the Escalation check box to implement the escalation policy for data errors. If this is not selected, escalation is not implemented.
- **Escalation count:** Use the counter to select or type the number of errors that can occur to the data for a specific record before carrying out the escalation policy.
- **Escalation policy:** Click the triangle to open the list and select what happens when the Escalation Count you entered is reached. You can select one of the following from the list:
 - **Log error:** The task continues and the error is written to the task log.
 - **Suspend table** (default): The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.
- **Truncation policy:** Click the triangle to open the list and select what happens when a truncation occurs in one or more specific records. You can select one of the following from the list:
 - **Ignore record:** The task continues and the error is ignored.
 - **Log error** (default): The task continues and the error is written to the task log.
 - **Suspend table:** The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.

For information about environmental errors and the configuration properties, see [Data Errors](#) and [Error Handling Properties](#) in the [Error Handling](#) appendix.

Table Error

Click the **Table Error** sub-tab to configure the following:

- **Policy:** Click the triangle to open the list and select what happens when an error occurs in the general table data for one or more tables being replicated. You can select one of the following from the list:
 - **Suspend table** (default): The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.
- **Escalation:** Select the Escalation check box to implement the escalation policy for table errors. If this is not selected, escalation is not implemented.
- **Escalation count:** Use the counter to select or type the number of errors that can occur to the data for general table data for a specific table before the task is stopped.

- **Escalation policy:** You cannot take any action in this option. The escalation policy for table errors is automatically set to **Stop task**.

For information about environmental errors and the configuration properties, see [Table Errors](#) and [Error Handling Properties](#) in the [Error Handling](#) appendix.

Apply Conflicts

Click the **Apply Conflicts** sub-tab to configure the following:

- **No record found for applying a DELETE:** Click the triangle to open the list and select what happens when there is a conflict with a DELETE operation. You can select one of the following from the list:
 - **Ignore record** (default): The task continues and the error is ignored.
 - **Log record to the exceptions table:** The task continues and the record is written to the exceptions table.
 - **Suspend table:** The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.
- **Duplicate key when applying an INSERT:** Click the triangle to open the list and select what happens when there is a conflict with an INSERT operation. You can select one of the following from the list:
 - **Ignore record:** The task continues and the error is ignored.
 - **Log record to the exceptions table** (default): The task continues and the record is written to the exceptions table.
 - **Suspend table:** The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.
 - **Update the existing target record:** The target record with the same primary key as the INSERTED source record is updated.
- **No record found for applying an UPDATE:** Click the triangle to open the list and select what happens when there is a conflict with a UPDATE operation. You can select one of the following from the list:
 - **Ignore record:** The task continues and the error is ignored.
 - **Log record to the exceptions table** (default): The task continues and the record is written to the exceptions table.
 - **Suspend table:** The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.
 - **Insert the missing target record:** The missing target record will be inserted into the target table. When the source database is Oracle, selecting this option requires supplemental logging to be enabled for all the source table columns.
- **Escalate on repeating table apply conflicts:** Select this check box to implement the escalation policy for table errors. If this is not selected, escalation is not implemented.
- **Escalation count:** Use the counter to select or type the number of apply errors that can occur before carrying out the escalation policy.

- **Escalation policy:** Click the triangle to open the list and select what happens when the Escalation Count you entered is reached. You can select one of the following from the list:
 - **Log error** (default): The task continues and the error is written to the task log.
 - **Suspend table:** The task continues but data from the table with the error record is moved into an error state and its data is not replicated
 - **Stop task:** The task is stopped and manual intervention is required.

For information about environmental errors and the configuration properties, see [Apply Errors](#) and [Error Handling Properties](#) in the [Error Handling](#) appendix.

Note: When you select **Fix record** you must be sure that you are using full supplemental logging to ensure that an UPDATE is not turned into an INSERT. In other cases, FIX_RECORD can cause an async full load of a record similar to the LOB channel.

Logging

When you click **Logging** in the Task Settings dialog box, you can carry out the actions in the following sub-tabs:

- [Logging Level](#)
- [Log File](#)

Logging Level

Click the **Logging Level** sub-tab to set the logging level for task logs. The level you set determines what information is written to the log

Note: You can also set the task logging level from the **Tools** menu in the Monitor view. For more information, see [Monitor Mode](#). and [Working with Task Logs at Runtime](#).

The following are the available logging levels. The list is in order from the lowest level to the highest level.

1. Error
2. Warnings
3. Info
4. Debug
5. Detailed Debug

The higher levels always include the messages from the lower levels. Therefore, if you select **Error**, only error messages are written to the log. However, if you select **Info**, informational messages, warnings, and error messages are included. Selecting **Detailed Debug** writes all possible messages to the log.

For information on how to set the logging level, see [Set the Task Logging Level](#).

Log File

You can manually roll the log file for the task you are working with in the **Log File** sub-tab. This lets you stop logging to the current log file and begin to log to a new log

file. The current log file is called `reptask_<name of task>` and saved (older) log files have the file name `reptask_<name of task>_xxxxxxxxxxxx` where `xxxxxxxxxxxx` represents a 12-digit timestamp.

Click **Roll Log File** to immediately roll the current server log file for the task you are working with.

Working with Tasks at Runtime

This chapter describes how to work with tasks that you design. For information on how to design a task, see [Designing Tasks](#). This chapter contains information on running tasks, viewing the task status, and viewing messages about the task. Information on monitoring and working with tasks during runtime is in the chapter [Monitoring and Controlling Replication Tasks](#).

This chapter contains the following topics:

- [Running a Task](#)
- [Viewing the Task Status](#)
- [Reading Messages about a Task](#)

Running a Task

After you design a task (see [Designing Tasks](#)), you can run and monitor its progress with one click in Amazon RDS Migration Tool. This simple Click-2-Replicate function is described in this topic. In addition, the various types of run options available are also described. This topic has the following sub-topics.

- [How to run a Task](#)
- [Using the Run Button Options](#)

Note: The task run buttons area available in the button bar at the top of the console in the following views:

- [The Tasks View](#) (In both the [Design Mode](#) and [Monitor Mode](#))
 - When [Viewing Specific Tasks](#)
-

How to run a Task

Click the **Run** button to execute a replication task. The task process continues to run until you click the **Stop** button to stop the task.

Note:


When you click **Run** the following occurs:

- If this is the first time that a task is run, the [Start Replication](#) operation is run.
 - If the task has been started and stopped, the **Resume Processing** operation described in [Use Advanced Run Options](#) is run.
 - If changes were made to the database, change processing takes place after the full load operation. If you do not want change processing to occur or if you want to start change processing from a pre-determined point, you must make the appropriate [Use Advanced Run Options](#) selection.
-
-

In some cases, task replication may stop due to an error although the task process is still running. See [Task Icons](#) for information on the task status and how Amazon RDS Migration Tool displays information on the current task status.

The following figure shows the **Run** and **Stop** buttons in the button bar.

Figure 16–1 Run and Stop Buttons



The **Run** button is available in the following views:

- The Tasks view when you select a task from the Task List.
- In the individual task, both the Design mode and Monitor mode have the **Run** and **Stop** buttons available.

Note: You must be in the Monitor mode to view the task progress.

Using the Run Button Options

Clicking the **Run** button runs a full-load replication task from the source to the target. This is a first time task that creates the target databases and loads the source data to the target according to your task definitions.

Subsequent runs allow you to resume processing from a specific point and process changes. In addition, you can also specify from what point you want the replication to start.

The following options are available:

- [Start Replication](#)
- [Reload Target](#) (Full Load Only)
- [Use Advanced Run Options](#)

Start Replication

This is available the first time you run the task only. This will execute the initial full load operation. If CDC is also enabled for the task, change processing will start as soon as any changes are made to the source database.

Reload Target

Starts the full load and change processing (if enabled) from the beginning. Already processed tables are handled according to the Task Setting, **Initial Table Creation; If target table already exists**. See [Full Load Tuning](#) for more information about this setting.

File Channel Note: To replicate tables that were added to the local file channel task after the initial full load, you need to reload both the local and the remote file channel tasks.

Use Advanced Run Options

Advanced Run Options provide you with additional options for resuming and restarting tasks.

To use Advanced Run Options, click the triangle next to the **Run** button and select **Advanced Run Options**. One of the dialog boxes shown below will open, depending on whether the replication task is Full Load + CDC or CDC only.

Figure 16–2 Advanced Run Options - Full Load

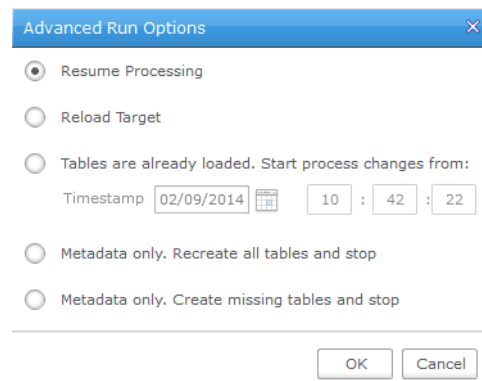
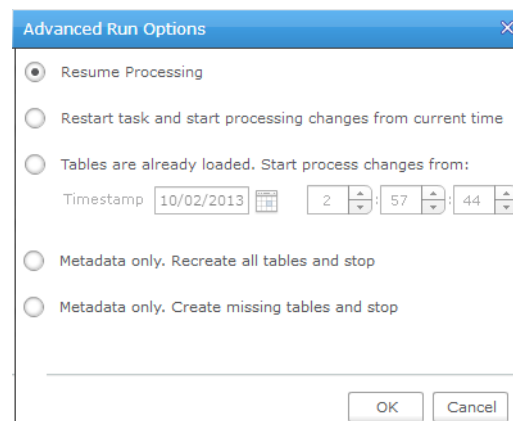


Figure 16–3 Advanced Run Options - CDC Only



The **Advanced Run Options** dialog box lets you do the following:

- Resume Processing:** Resumes task execution from the point that it was stopped. You can also resume processing by clicking the **Run** button if the task has been stopped.

Note: If the schema or a filter was changed after the task stopped, the task should be reloaded as opposed to resumed (see below).

- ***Reload Target:** This will re-start the full-load replication process if the task was previously run. See [Reload Target](#) for more information.
*Only available for Full Load tasks.
- ****Restart task and start processing changes from current time:** This starts the CDC replication task from the beginning (as if the task has not run before).
**Only available for CDC tasks.
- **Tables are already loaded. Start process changes from:** Select the date and time to create a timestamp to define the point from where you want to process changes.
Note: The timestamp uses the local time of the browser machine.
Click the calendar to select the date.

The "Metadata only" options described below allow you to:

- Create empty tables on the target and then manually edit them.
- Create tables during a task.

Enabling the options will also ensure that supplemental logging is set up correctly on the source tables before starting the actual replication task.

- **Metadata only. Recreate all tables and stop:** Select this option to recreate the target tables as defined in the **Full Load Settings** tab. To use this option, stop the existing task, run the task with this option enabled (the task will stop automatically) and finally, resume the task.
- **Metadata only. Create missing tables and stop:** Select this option to create missing target tables. To use this option, stop the existing task, run the task with this option enabled (the task will stop automatically) and finally, resume the task.

Viewing the Task Status

In the [The Tasks View](#), you can see the task status by viewing the icon for the task. After a task is run, the [Task Icons](#) in the Task view display the current status of the task. For additional information on the possible statuses, see [Task Icons](#).

The following icon represents a task that is in an error status.

There are two types of errors:



Recoverable error: A recoverable error indicates that there is a temporary problem, such as a missing connection. The task icon is blue indicating that the task is still active. In this case, Amazon RDS Migration Tool attempts to restart the task automatically. As soon as the error state is resolved, the task is restarted.

The task remains active but paused throughout the error state. You can stop the task at any time and resolve the error manually, if necessary.

Note: Amazon RDS Migration Tool will continue to check the task for 30 minutes to determine whether it is no longer in an error status. If the error is not resolved in 30 minutes the error becomes a fatal error and you must resolve the error manually.



Fatal Error: When a fatal error occurs, the task stops and you must resolve the error manually. You cannot start the task again until the error is resolved. Use the logs or the messages in the Alerts pane to see the error type.

See also:

[View Log Messages for a Task](#)

[Viewing Notifications](#)

Reading Messages about a Task

Task messages are displayed in the Messages section of the Amazon RDS Migration Console. The Messages section is located at the bottom right of the console in the [Monitor Mode](#) and when [Viewing Specific Tasks](#).

The Message section has two types of messages that provide information about events that occur in a task. Each type of message is displayed in the following tabs:

- [Notifications](#)
- [Log Messages](#)

Viewing Notifications

The **Notifications** tab displays notifications about the task. These messages alert you to specific events encountered by a task, such as the task starting or stopping, a specific error type, or information about latency and disk space.

The **Notifications** tab displays the time of a notification and a description of the notification. You define the notifications that are sent for each task and a description for each notification in the Settings area. For more information, see [Define the Notification Message](#).

Using the Notifications List

When a notification is sent, it is displayed in the **Notifications** tab. You can:

- [Open a Notification](#)

- [Clear a Notification](#)

Open a Notification When you open a notification, you can see the full message presented in a dialog box. The dialog box contains a button to copy the text so that you can use it somewhere else for trouble shooting and the timestamp for the notification.

To open a notification

1. In the Messages section of the console, click the **Notifications** tab. The **Notifications** tab opens.
2. Select the notification you want to open from the list.
3. Double-click the notification or click **Open** from the button bar at the top of the list.

Clear a Notification You can clear notifications from the list to make sure that you are seeing only those that are relevant to you.

To clear a notification

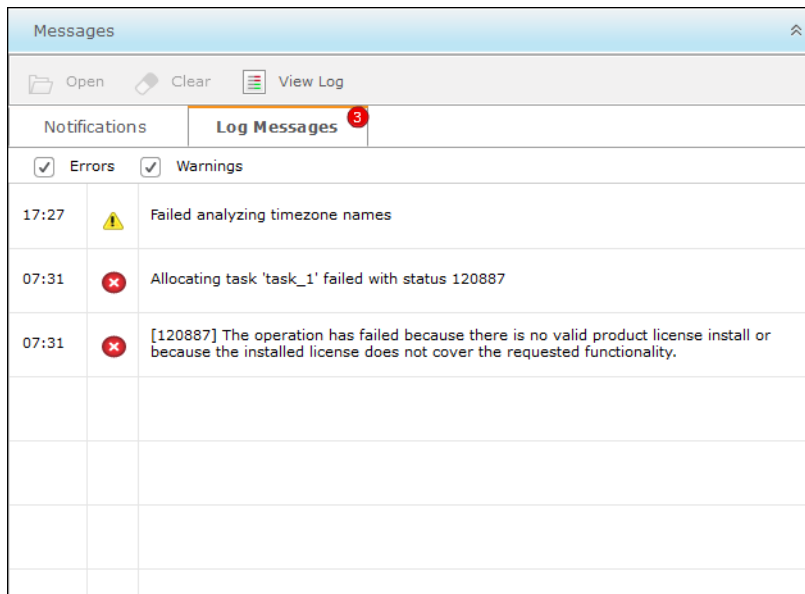
1. In the Messages section of the console, click the **Notifications** tab.
2. Select the notification you want to clear from the list.
3. Click **Clear** from the button bar at the top of the list.

View Log Messages for a Task

The **Log Messages** tab displays log messages for errors or warnings from a task. The errors are listed in this tab with the time of the error or warning and the log entry for the event. You can choose to view both errors and warnings or only one of them.

If errors or warnings exist in the task, a red circle with the total number of errors and warnings is displayed. The number displayed may be the number of errors, the number of warnings, or the total of number of errors and warnings depending on what you select to view in this tab. The **Log Messages** tab is shown in the figure below.

Figure 16–4 Log Messages



Using the Log Messages List

When a log error or warning is sent, it is displayed in the **Log Messages** tab. You can:

- [Select the Log Message Type](#)
- [Open a Log Message](#)
- [Clear a Log Message](#)

Select the Log Message Type Two types of log messages are displayed in the Log Messages List. You can view Errors, Warnings, or both.

To select the log message type

Select the check box or boxes for the type messages you want to view. The check boxes are located at the top of the Log Messages List.

Open a Log Message When you open a log message, you can see the full log text presented in a dialog box. The dialog box contains a button to copy the text so that you can use it somewhere else for trouble shooting and the timestamp for the log message.

To open a log message

1. In the Messages section of the console, click the **Log Messages** tab.
2. Select the log message you want to open from the list.
3. Double-click the log message or click **Open** from the button bar at the top of the list.

Clear a Log Message You can clear log messages from the list to make sure that you are seeing only those that are relevant to you.

To clear a log message

1. In the Messages section of the console, click the **Log Messages** tab.
2. Select the log message you want to clear from the list.
3. Click **Clear** from the button bar at the top of the list.

Viewing the Log file in the Log Viewer

In addition to viewing the log messages, you can view the entire log file in the log viewer.

To view the log in the log viewer

1. From the **Messages** section, click **View Log**.
2. The log for the task you ran is displayed.

You can select a different log from the **Select Log** drop-down list in the top right of the Live Log Viewer window.

Monitoring and Controlling Replication Tasks

Monitoring and running a task let you use the Click-2-Replicate function to carry out the replication task and view its functions in near real time. This chapter contains the following topics about running and monitoring a replication task.

- [Viewing Information in the Monitor](#)
- [Monitoring the Full-Load Operation](#)
- [Monitoring Change Processing Operations](#)
- [Viewing Messages](#)
- [Using the Monitor Tools](#)

Viewing Information in the Monitor

You access the Monitor view when you open a specific task. The monitor provides near real-time information for the task you select.

To access the Monitor

1. When [Viewing Specific Tasks](#), select the task you want to monitor.
2. From the button bar at the top of the console, click **Open**.
3. From the button bar at the top right, click **Monitor**.

The Monitor opens. To view the information in real time, you need to run the task (if the task has not already started). For information on running a task, see [Running a Task](#).

Monitoring the Full-Load Operation

You can view the progress of the full-load operation in the left section of the Monitor.

To make sure you are viewing the information for a full-load operation, select the **Full Load** tab.

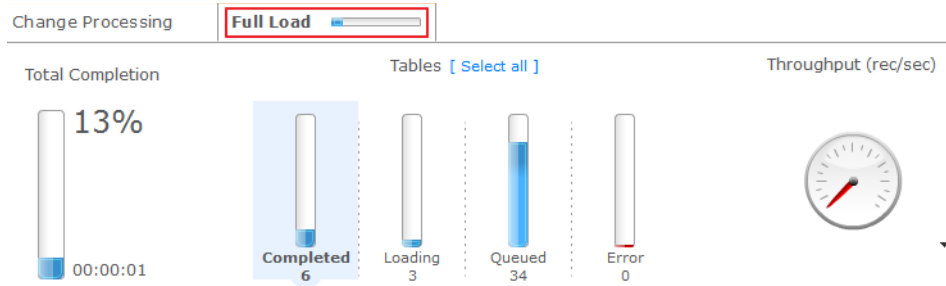
You can view the following:

- [General Information for a Full Load](#)
- [Detailed Information for the Full Load](#)
- [Monitoring Throughput in a Full Load Operation](#)

General Information for a Full Load

General information about the full load is presented in a graphical format. The following figure shows the graphical information displayed when a task is running.

Figure 17-1 Full Load Status Bars



This section has the following information:

- Status bars: Indicates the status of the tables being loaded.
 - **Completed:** The number of tables that finished loading into the target database.
 - **Loading:** The number of tables that are in the process of loading into the target database.
 - **Queued:** The number of tables that are waiting to load into the target database.
 - **Error:** The number of tables that could not be loaded due to an error. See [Reading Messages about a Task](#) for information about error messages.
- Full-load total completion bar: Displays the progress of all records being loaded to the target database. The bar is located in the **Full Load** tab at the top of the graph section.
- Throughput gauge: Displays the current throughput. Throughput displays the number of events read in the task for a specified amount of time.

You can also view [Detailed Information for the Full Load](#).

Detailed Information for the Full Load

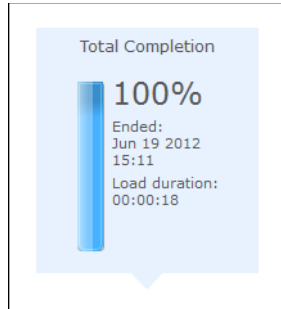
For each of the status bars displayed in the [General Information](#) graphs, a table is displayed in the section below with specific information about the current loading status. The following information is available:

- [General Information for Completed Task](#)
- [Information for Each Table in the Task](#)
- [Information for Tables that have Completed Loading](#)
- [Information for Tables that are Currently Loading](#)
- [Information for Tables that are in the Loading Queue](#)
- [Information for Tables with Errors](#)

General Information for Completed Task

This section displays a table with information for all of the completed tables in a task. To view this table, click the **Total Completion** bar, shown in the figure below.

Figure 17–2 Total Completion Status



This table displays the following Progress Details:

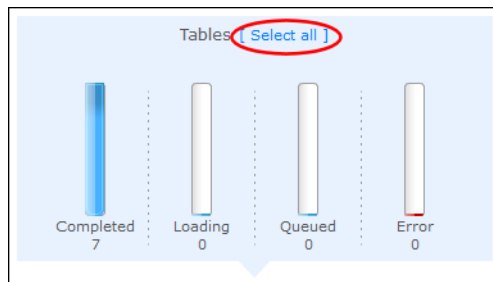
Table 17–1 Progress Details for all Tables in the Task

	Total	Completed	Remaining	Notes
Tables	The total number of tables that are included in the task.	The total number of tables that completed loading at the current time.	The total number of tables waiting to be loaded.	Additional information.
Records	The total records that completed loading at the current time.	The total number of records that completed loading at the current time.	The total number of records waiting to be loaded.	Additional information.
Time	The estimated time to load all of the selected tables in the task.	The total elapsed time.	The estimated amount of time to load the remaining tables.	Additional information.

Information for Each Table in the Task

This section describes the progress of each of the tables being processed for the task. To display this information, click **[Select all]** at the top of the page as shown in the figure below.

Figure 17–3 Select All Tables



The progress is displayed in a table that has the following columns:


- **Table:** The names of the source tables that are included in the task.

- **Status:** This is a statement that describes the status for the table. The following are the statuses that can be displayed:
 - **Queued:** The table is in the queue waiting to be loaded to the target database.
 - **Loading:** The table is being processed but is not finished loading.
 - **Completed:** All of the table records are loaded to the target.
 - **Error:** The table stopped loading due to an error. See [Reading Messages about a Task](#) for more information about task errors.
- **Count:** The number of records that are loaded to the target.
- **Elapsed Time:** The total elapsed time since the table records began processing.
- **Progress:** The table status and the time the table entered that status.
- **Reload:** Click the **Reload** icon to reload the data for selected tables and run the full-load operation again.

Information for Tables that have Completed Loading

This section displays a table with information for each of the completed tables. To view this table, click the **Completed** bar, shown in the figure below.

Figure 17–4 Completed Tables Status




The information is displayed in a table that has the following columns:

- **Table name:** The names of the source tables that have completed loading.
- **Loaded On:** The time that the table completed loading all of its records to the target.
- **Transferred Count:** The number of records loaded to the target.
- **Transferred Volume:** The volume of the records (in KB) loaded to the target.
- **Load Duration:** The amount of time that it took for all records to load to the target.
- **Throughput Records:** The average throughput rate for the table. Throughput describes the number of records read per second. For more information on throughput, see [Monitoring Throughput in a Full Load Operation](#).
- **Throughput Volume:** The average throughput rate for the table. Throughput describes the volume of records (in KB) read per second. For more information on throughput, see [Monitoring Throughput in a Full Load Operation](#).
- **Reload:** Click the **Reload** icon to reload the data for selected tables and run the full-load operation again.

Information for Tables that are Currently Loading

This section displays a table with information for each of the tables that are currently loading. To view this table, click the **Loading** bar, shown in the figure below.

Figure 17–5 Loading Tables Status




The information is displayed in a table that has the following columns:

- **Table Name:** The names of the source tables that are currently loading.
- **Load Duration:** The amount of time that it took for all records to load to the current point in time.
- **Estimated Count:** The estimated number of rows that are to be loaded in the full load operation.
- **Transferred Count:** The number of records that are loaded to the target database.
- **Current Throughput:** The current throughput rate for the table. Throughput describes the number of records read per second. For more information on throughput, see [Monitoring Throughput in a Full Load Operation](#).
- **Estimated Finish Time:** The approximate time the task finished loading the tables. The timestamp displayed indicates the date and time.
- **Progress:** The table status and the time the table entered that status.
- **Reload:** Click the **Reload** icon to reload the data for selected tables and run the full-load operation again.

Information for Tables that are in the Loading Queue

This section displays a table with information for each of the tables that are waiting to be loaded. To view this table, click the **Queued** bar, shown in the figure below.

Figure 17–6 Queued Tables Status




The information is displayed in a table that has the following columns:

- **Table Name:** The names of the source tables that are currently in the queue waiting to be loaded.
- **Estimated Count:** The estimated number of rows that are waiting to be loaded in the full load operation.

Information for Tables with Errors

This section displays a table with information for each of the tables that stopped loading or suspended CDC due to an error. To view this table, click the Error bar, shown in the figure below.

Figure 17–7 Error Tables Status




The information is displayed in a table that has the following columns:

- **Table Name:** The names of the source tables that stopped due to an error.
- **Failed On:** The time that the error occurred.
- **Loaded Count:** The number of records loaded when the error occurred.

Monitoring Throughput in a Full Load Operation

Throughput values for a full-load operation provide information on how fast the table records are being replicated to the target database. The information is displayed in a gauge on the right side of the full-load graph section. The following figure shows the throughput gauge.

Figure 17–8 Throughput Gauge




You can set the throughput measurement values either to the number of records replicated per second, or to the number of kilobytes replicated per second. The display is always based on the current load operation.

To set the unit of throughput measurement:

- Select either **rec/sec** or **kbyte/sec** from the drop-down menu below the Throughput gauge.

Click the Throughput gauge to display a graph with the throughput details as shown in the figure below. To view the graph only and hide the status bars and throughput gauge, click the double arrow above the graph.

Figure 17–9 Throughput Details



Monitoring Change Processing Operations

You can view the progress of the change-processing operation in the left section of the Monitor.

To make sure you are viewing the information for a change-processing operation, select the **Change Processing** tab.

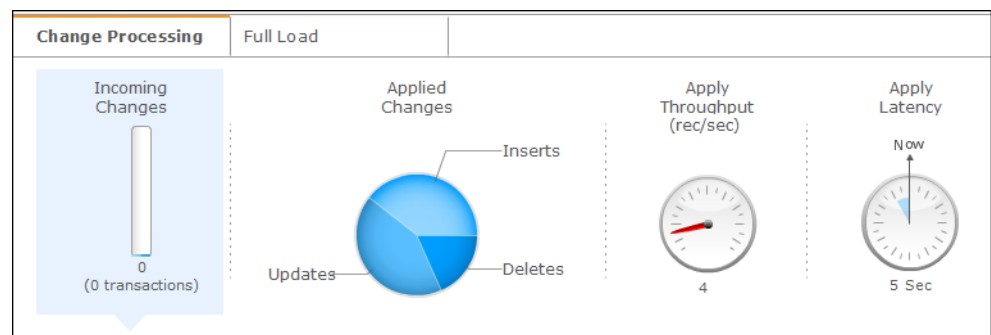
You can view the following:

- [General Change Processing Information](#)
- [Detailed Change Processing Information](#)

General Change Processing Information

General information about the full load is presented in a graphical format. The following figure shows the graphical information displayed:

Figure 17–10 Change Processing Status



This section has the following information:

- **Incoming Changes:** The total number of records that were processed for the task.
- **Applied Changes:** A circle graph that shows information about the processed changes. It displays the following:
 - The number of INSERT operations processed. Roll over the Insert section with your mouse to see the number and per cent of the accumulated inserts.

- The number of UPDATE operations processed. Roll over the Update section with your mouse to see the number and percentage of the accumulated updates.
- The number of DELETE operations processed. Roll over the Delete section with your mouse to see the number and percentage of the accumulated deletes.
- The number of metadata changes (DDL) processed. DDL changes include information about events like changes to table names or to column names.
- **Apply Throughput** gauge: A gauge that describes the number of change events read per second. For additional details, you can also view a graph with [Information about Change Processing Throughput](#).
- **Apply Latency** gauge: A gauge that displays the latency information.

The latency values displayed in the Amazon RDS Migration Console measure the time delay (latency) between the time when a change is visible to the source (and committed), and the time when this same change is visible to the target. The display is always based on the current change being applied.

You should take the following into consideration:

- Latency when applying large transactions:

For example, when the most recent latency value was 10 seconds and now a transaction of one million rows gets committed at the source database, Amazon RDS Migration Tool starts to apply that transaction to the selected target and it will take some time to write all the changes to the target (for example 60 seconds). During the next 60 seconds, the latency value gradually grows to 70 seconds for the last change in the transaction. Once the transaction is committed, the latency drops back to the 'regular' latency (10 seconds in this case).
- Latency when no transactions are being applied:

When a time period passes with no changes applied to the target, the latency calculation is based on the time difference between the current time and the timestamp of the last change event read from the transaction log. This could happen if, for example, there is high activity on tables which are not selected for replication in the current task.

For additional details, you can also view a graph with [Information about Apply Latency](#).

Detailed Change Processing Information


For each of the status indicators displayed in the [General Change Processing Information](#) section, a table or graph is displayed in the section below with detailed information about the change processing status. The following information is available.

- [Information about Incoming Changes](#)
- [Information about Applied Changes](#)
- [Information about Change Processing Throughput](#)
- [Information about Apply Latency](#)

Information about Incoming Changes

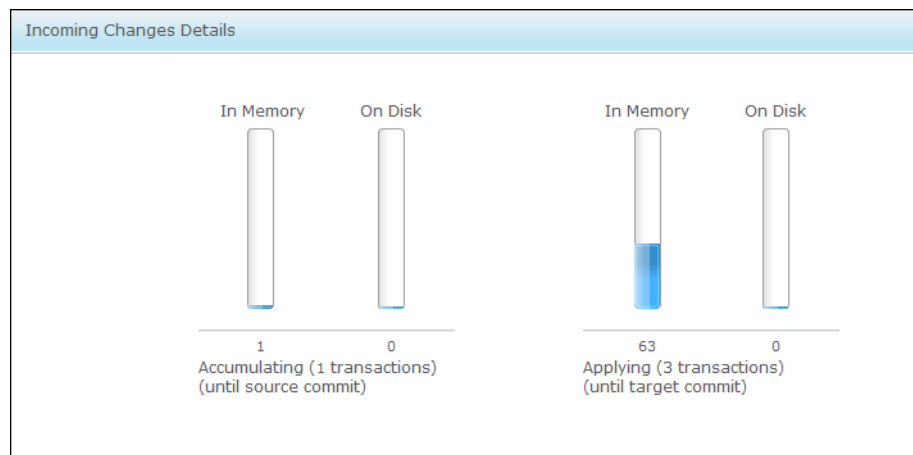
This section displays two bar graphs with information about incoming changes. Incoming changes displays the number of change records currently being read from the source database and written to the target database. To view these graphs, click the Incoming Changes bar, shown in the figure below.

Figure 17–11 Incoming Changes



The following graphs are displayed.

Figure 17–12 Incoming Change Graphs



The graphs have the following information:

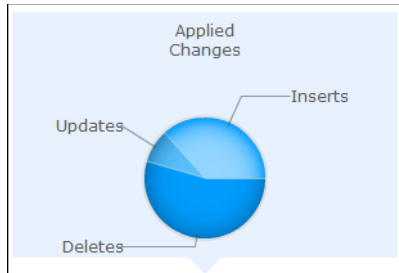
- **Accumulating:** These bars display the number of records currently being read from the source database. These records are accumulated in a queue until they are applied to the target. The following is displayed:
 - **In Memory:** The number of accumulating records that are currently in the computer memory.
 - **On Disk:** The number of accumulating records that are currently stored on disk.
- **Applying:** The number of records currently being written to the target. These are the applied changes. The following is displayed:
 - **In Memory:** The number of records being applied that are currently in the computer memory.
 - **On Disk:** The number of records being applied that are currently stored on disk.

Information about Applied Changes

This section displays two tables with information about the applied changes.

To view these tables, click the Applied Changes pie graph, shown in the figure below.

Figure 17–13 Applied Changes



The following tables are available when you select Applied Changes.

- [Recent Activity](#)
- [Aggregates](#)

Recent Activity

Click **Recent Activity** at the top of the Applied Changes Details section to view information about which changes occurred in each table. It has the following information:

- **Table Name:** The names of the source tables that are included in the task.
- **Insert:** The number of INSERT operations processed for the specific table.
- **Delete:** The number of DELETE operations processed for the specific table.
- **Update:** The number of UPDATE operations processed for the specific table.
- **DDL:** The number of metadata changes (DDL) processed. DDL changes include information about events like changes to table names or to column names.
- **Total Applied:** The total number of changes applied to the target.
- **Last Modified:** The time the last change occurred for the specific table.
- **Reload:** Click the **Reload** icon to reload the data for selected tables and run the full-load operation again.

Aggregates

Click **Aggregates** at the top of the Applied Changes Details section to view information about total changes for each change type and transaction type.

The Aggregate table displays the total changes (for all tables) applied for each of the following types of operations:

- INSERT
- UPDATE
- DELETE
- DDL

The Aggregate table also displays the information about transactions. It displays the total number and volume of:

- COMMITS
- ROLLBACKS

Information about Change Processing Throughput

Throughput values for apply throughput in a change-processing operation provide information on how fast the change records are loaded to the target database. The information is displayed in a gauge in the Change-Processing graph section. The following figure shows the Apply Throughput gauge:

Figure 17–14 Apply Throughput Gauge




You can set the Apply Throughput measurement values either to the number of change records replicated per second, or to the number of kilobytes replicated per second. The display is always based on the current load operation.

To set the unit of throughput measurement:

- Select either **rec/sec** or **kbyte/sec** from the drop-down menu below the **Apply Throughput** gauge.

Click the **Apply Throughput** gauge to display a graph with the throughput details as shown in the figure below. To view the graph only and hide the status bars and Change-Processing gauges (including the **Apply Throughput** gauge), click the double arrow above the graph.


Figure 17–15 Apply Throughput Details Graph



Information about Apply Latency

Latency values for apply latency in a change-processing operation provide information about the time delay (latency) between the time when a change is visible to the source (and committed), and the time when this same change is visible to the target. The information is displayed in a gauge in the Change-Processing graph section. The following figure shows the Apply Latency gauge.


Figure 17–16 Apply Latency



The latency values displayed in the Amazon RDS Migration Console measure the time delay (latency) between the time when a change is visible to the source (and committed), and the time when this same change is visible to the target. The display is always based on the current change being applied. For more information about latency, see [Apply Latency](#).

Select the Apply Latency gauge to display a graph with the latency details. To view the graph only and hide the Change-Processing gauges including the latency gauge, click the double arrow as shown in the figure below.

Figure 17–17 Apply Latency Details Graph



Viewing Messages

You can see messages sent for the task while in the monitor view. For information on viewing messages, see [Reading Messages about a Task](#).

Using the Monitor Tools

The monitor tools let you view additional information about the task. The following topics describe the information available through these tools:

- [Viewing History Information](#)
- [Working with Task Logs at Runtime](#)

Viewing History Information

The **History** window displays information for each event carried out in the task. To access the **History** information, from the **Monitor Mode**, click **Tools** and then select **History**.

You can view the following information in the **History** window:

- **Event type:** The type of event that occurred, for example **Task started** or **Task table load finished**.
- **Timestamp:** A timestamp that indicates when the event took place. The timestamp is in the format, `YYYY-MM-DD hh:mm:ss.milliseconds` (to six places).
- **Table Name:** The name of the table where the event takes place if the event is related to a table.
- **Description:** A description of the event. This is not displayed for all events.

You can double-click the description cell to view a window with the full message if the entire description is not available.

Working with Task Logs at Runtime

Use the **Logging** window to define how you want to work with logs for the task you are currently monitoring.

Note: For information about configuring and viewing logs for a specific task in the **Logging Level** sub-tab in the Task Settings dialog box, see [Logging Level](#).

To open the logging window from the Monitor

1. Open the task you are working with if it is not displayed in the Amazon RDS Migration Console. For information on opening a task, see [Editing a Replication Task](#).
2. On the right side of the button bar, click the **Monitor**.
3. Click **Tools** menu in the button bar then select **Logging**. The Logging window is displayed.

You can do the following in the Logging window.

- [Set the Task Logging Level](#)
- [Carry Out Task Log File Actions](#)
- [View the Task Log File](#)

Note: If you want to view the log file, select **Launch Log Viewer** to view the log file in the viewer window. For more information, see [View Log Messages for a Task](#).

To set the log level for the server, see [The Server View](#).

Set the Task Logging Level

You change the logging level for a task that you are monitoring from the Logging dialog box.

To change the task logging level

1. Open the Logging window. See [Working with Task Logs at Runtime](#).
2. At the top of the Logging window, move the **Task Logging Level** slider to the log level you want. This sets the logging level for all server log modules. Note that all of the sliders for the individual modules move to the same position that you set in the main slider.
3. Make any changes to the sliders for the individual modules. This is optional. Note that if you change the main slider, all of the individual sliders are reset to the new position. If you want to maintain a different logging level for a specific module, you need to reset it.

See [Figure 0–2, "The Logging Window \(Monitor View\)"](#) for a list of the logging modules you can set.

4. Click **Apply** at the top of the Logging window. Changes to the logging level take place immediately. There is no need to restart the Amazon RDS Migration Server.

Carry Out Task Log File Actions

You can carry out the following actions for the log file for the task you are monitoring.

Roll Log File

You can manually roll the log file for the task you are monitoring in the Logging window. This lets you stop logging to the current log file and begin to log to a new log file. The current log file is called `reptask_<name of task>` and saved (older) log files have the file name `reptask_<name of task>_xxxxxxxxxxxx` where `xxxxxxxxxxxx` represents a 12-digit timestamp.

Click **Roll Log File** to immediately roll the current server log file.

Click **Delete Logs** to delete the current logs.

View the Task Log File

You can view any of the task log files in the Log Viewer.

To view the task log files

1. Open the Logging window. See [Working with Task Logs at Runtime](#).
2. At the top of the Logging window, click **Log Viewer**. The Live Log Viewer is displayed.
3. To view a different log file, select the file you want to view from the **Select Log** list at the top right of the window.

Amazon RDS Migration Server Settings

This chapter describes how to configure Amazon RDS Migration Tool using the Server page. For information on opening and viewing the Server page, see [The Server View](#).

Note: Configurations made in the Server page affect all Tasks that are created in the Amazon RDS Migration Tool instance you are working with.

You can do the following in the Server page:

- Create notifications, define mail settings and default recipients for sending notifications.
For more information, see [Notifications Settings](#).
- Register or request a License to work with Amazon RDS Migration Tool.
For more information, see [License Settings](#).
- Set the error-handling policy for the Amazon RDS Migration Server.
For more information, see [Error Handling](#).
- Set the logging level for the Amazon RDS Migration Server.
For more information, see [Logging Settings \(Server\)](#).
- Add File Transfer Service hosts
For more information, see [File Transfer Service](#).
- Schedule jobs
For more information, see [Scheduling Jobs](#).

Notifications Settings

The following can be defined in the Notifications settings:

- [Defining Notifications](#)
- [Setting up Mail Parameters](#)
- [Creating a Default Recipient List](#)

To view and edit the Notification settings

In **Server** view, click the **Notifications** tab on the left. Then click the **Notifications** sub-tabs to enter your settings.

Defining Notifications

To configure and create notifications, click the **Notifications** sub-tab.

You use notifications to send messages about events that occur when running tasks in Amazon RDS Migration Tool. Notifications are sent to inform users of any change in the system state, including:

- A task is started or stopped
- Latency is too high
- Memory utilization is too high
- Disk utilization is too high
- An error or a specific type of error occurred

You can manage notifications that you create from the Notifications list. This list provides you with information about each notification defined and lets you activate/deactivate a notification. In addition, you can make changes to the definitions of existing notifications or delete them.

The following topics describe how to define notifications in Amazon RDS Migration Tool:

- [Creating a New Notification](#)
- [Using the Notification List](#)
- [Editing a Notification](#)
- [Deleting a Notification](#)

To open the Notifications page

From the Server view, click **Notifications** from the menu list at the left. The **Notifications** sub-tab is displayed.

Notifications are sent by:

- An email message to the default list of users and/or to a custom list of users.
- Writing an entry in the Windows Event Log.
- Displaying a message in the Amazon RDS Migration Console.

Creating a New Notification

Use the New Notification Rule wizard to determine the notifications that are sent and who receives them.

To start the New Notification Wizard

1. From the Settings page, click **Notifications**.
2. At the top of the Notifications page, click **New Notification**.
The New Notification Rule wizard, **Notify When?** page, opens.
3. In the **Notification Name** field, type a name for this notification.
4. Continue to work with this wizard to define the notification. You must do the following to create the notification:
 - [Define the Action that Triggers the Notification](#)
 - [Define Which Changes of Status Trigger the Notification](#)

- [Define Errors That Trigger the Notification](#)
- [Define the Notification Distribution Properties](#)
- [Determine the Email-Message Recipients for the Notification](#)
- [Define the Notification Message](#)
- [Associate Tasks with the Notification](#)
- [Review the Notification Rule](#)

Define the Action that Triggers the Notification In the **Operator** section of the **Notify When?** page, you can determine the action that triggers the notification. If the **Operator** section is not displayed, click on the header with the word **Operator** to display the options for this section. Select one of the following:

- **Task was started:** To send the notification when the task starts.
- **Task was stopped manually or scheduled:** To send the notification when the task is stopped either manually or by the Scheduler.
- **Task was stopped after Full Load: Cached changes were not applied:** To send the notification when the task is stopped after Full Load completes but before cached changes (changes to the source tables that occurred during Full Load) are applied to the target.
- **Task was stopped after Full Load: Cached changes were applied:** To send the notification when the task is stopped after Full Load completes and cached changes (changes to the source tables that occurred during Full Load) have been applied to the target.
- **Full load started:** To send the notification when the Full Load process starts.
- **Full load completed:** To send the notification when the Full Load process completes.

Once you determine when to send the notification, you can decide whether specific changes in status trigger the notification.

If you want to send a message about problems in latency, memory utilization, or disk utilization, click **Performance/Resources**. See [Define Which Changes of Status Trigger the Notification](#) for an explanation.

If you want to send the notification when certain errors occur, click **Errors**. See [Define Errors That Trigger the Notification](#) for an explanation.

Or you can click **Next** to [Define the Notification Distribution Properties](#).

Define Which Changes of Status Trigger the Notification In the **Performance/Resources** section of the **Notify When?** page, you can define specific parameters for latency, disk utilization, or memory utilization that trigger a notification.

To set up notifications about latency, disk utilization, or memory utilization

1. In the [New Notification Wizard](#), **Notify When?** page, click **Performance/Resources**.
2. Select one of the following:
 - **Latency is higher than Value seconds.**
 - **Memory utilization exceeded Value MB**
 - **Disk utilization exceeded Value MB**

- Define the value for the option you select. See the table below for an explanation on each of these options and how to set the value.

Note: If you select one of these options, the notification is sent only when the selected parameter is true. However, you must also [Define the Action that Triggers the Notification](#).

Table 18–1 Set Values for Latency, Disk Utilization, Memory Utilizations

Notification	Set Value	Notes
Latency is higher than Value seconds	Click [N] and enter a value in the field that is displayed. Latency is the time interval in seconds between the time a change was committed in the source system and the time it is applied and committed in the target system.	
Clear notification when latency drops below <n> seconds.	Use this to set the value that determines when latency returns to "normal limits." Click [N] and enter a value.	When latency is below the value entered in this field, it is considered to be in the "normal" range and the notification status ends. If selected, a notification is sent to indicate that latency returned to "normal" status. For more information, see Define the Notification Message .
Memory utilization exceeded Value MB	Click [N] and enter a value in the field that is displayed. Memory utilization is the amount of memory used by the task.	
Clear notification when memory utilization is below <n> MB	Use this to set the value that determines when memory utilization returns to "normal limits." Click [N] and enter a value.	When memory utilization is below the value entered in this field, it is considered to be in the "normal" range and the notification status ends. For more information, see Define the Notification Message .
Disk utilization exceeded Value MB	Click [N] and enter a value in the field that is displayed. Disk utilization is the amount of disk space used. Set a value that indicates that the current amount of disk space used is problematic to running a replication task.	
Clear notification when disk utilization is below <n> MB	Use this to set the value that determines when disk utilization returns to "normal limits." Click [N] and enter a value.	When disk utilization is below the value entered in this field, it is considered to be in the "normal" range and the notification status ends. For more information, see Define the Notification Message .

Once you determine the status changes that trigger a notification, you can decide whether specific errors trigger a notification.

If you want to send the notification when certain errors occur, click **Errors**. See [Define Errors That Trigger the Notification](#) for an explanation.

Or you can click **Next** to [Define the Notification Distribution Properties](#).

Define Errors That Trigger the Notification In the **Errors** section of the **Notify When?** page, you can determine whether notifications are sent when an error occurs. You can determine whether to send the notification for all errors or only for specific error types.

To set up notifications for errors

1. In the [New Notification Wizard](#), **Notify When?** page, click **Errors**.
2. Select one of the following:
 - **Task encountered a non-retriable error and was stopped:** Select this to receive a notification when an error that cannot be retried is returned and a task or tasks are stopped due to this error.
 - **Table encountered more than [N] apply errors:** Select this to receive a notification when a specified number of errors that occur in a CDC operation are applied to a table. In this case, the table is not loaded but the task continues to run.

Click **Value** and type the number of errors to trigger the notification. For example, type 50 to send a notification after fifty records fail to be applied to the target table. All apply errors are logged in the `amazon_apply_exceptions` table. The count of apply errors is reset each time the task starts.

- **Table processing suspended due to errors:** Select this to receive a notification when an error causes a table to stop processing during a full-load operation or suspend CDC. In this case, the table process stops, but the task continues.
- **Any Error:** Select this to receive a notification when an error occurs in the system.
- **Any Warning:** Select this to receive a notification when a warning is issued in the system.

Once you determine the error types that trigger a notification, you can:

- [Define the Action that Triggers the Notification](#), if you have not done this already.
- [Define Which Changes of Status Trigger the Notification](#) if you have not done this.
- Or you can click **Next** to [Define the Notification Distribution Properties](#).

Define the Notification Distribution Properties In the **Notify How/Who?** page of the New Notification Rule wizard, you determine which users receive the notification and how the notification is sent.

To determine the delivery notification properties

Select any of the following to determine where (how) the notification is sent:

- **Replication Console:** This selection is selected by default. You cannot change this selection. All notifications are sent to the Replication Console. Notifications are displayed in the **Messages** section for a specific task. This section is displayed in:
 - The Monitor for a specific task. For more information, see [Reading Messages about a Task](#).

- The right pane of the Tasks page. For more information, see [The Tasks View](#).
- **Event Log:** Select this if you want the notification message to be written to the Windows/Linux Event log. For information on how to view the Windows/Linux Event log, see the online help for the version of Windows or Linux you are using.
- **Default Notification Email List:** Select this option if you want to send an email message to the all the recipients on the Default Notification Email List. For more information, see [Creating a Default Recipient List](#).

If you choose to add additional recipients to email notifications or send the email message only to a custom list of recipients, then stay on the **Notify How/Who?** page to [Determine the Email-Message Recipients for the Notification](#).

If you do not need to create a custom email recipient list for this notification, then click **Next** to [Define the Notification Message](#).

Determine the Email-Message Recipients for the Notification In addition to sending an email message for a specific notification to the default notification list, you can also create a custom notification list of users who receive this notification only or you can also send the email message to a custom list of users only.

This section describes how:

- [To create a custom list of users](#)
- [To send notification email messages to a custom list of recipients only](#)

To create a custom list of users

1. In the [New Notification Wizard](#), **Notify How/Who?** page, click **Add**. The add button and custom list are in the middle of the **Notify How/Who?** page.

The **Name** field in the first available row in the list is activated.

2. Type the name of the user that you want to receive the message.
3. In the **Name** column of the activated cell, type the name of the user you want to add to the list of default recipients.

Note: If you click another part of the Amazon RDS Migration Console, the cell will become inactive. You can double-click the cell to enter additional information.

4. Press the [tab] key or double click in the in the **Email** cell, then type the email address for the user you entered in the **Name** cell.
5. Click **Next** to [Define the Notification Message](#).

To send notification email messages to a custom list of recipients only

1. In the [New Notification Wizard](#), make sure that the **Default Notification Email List** option is *not* selected.
2. Create a Custom Notification list as described in [To create a custom list of users](#).

Define the Notification Message You can create a message for your notification. By default, a standard message is created based on the definitions you entered when [Define Which Changes of Status Trigger the Notification](#) and [Define Errors That Trigger the Notification](#).

To create a notification message

1. In the [New Notification Wizard](#), **Notification Message** page, double-click in any of the table cells to open the Edit Notification Message dialog box. See the [Creating](#)

a [Notification Message](#) table for an explanation of the information to enter in each field.

2. Click in the right pane of the dialog box and begin to type your message. In some cases a default message is displayed based on the information you entered in the previous pages of the New Notification Rule wizard. You can edit or delete the message, or create a new message to be sent with the notification in this dialog box.
3. Add variables in messages and email headers you define for notifications, if necessary. You can enter variables in one of two ways:
 - Type a variable into the message pane on the right, Use the following format:

```

{{<VARIABLE_NAME      >}}

```

 For example: {{TASK_NAME}}.
 - Use the variables from the left pane of the Edit Notification Dialog box. To add a variable to the notification message you can:
 - Double-click the variable. The variable is inserted where your cursor is located in the notification message in the right pane.
 - Select the variable you want to use and click the arrow key in the middle of the Edit Notification Message dialog box. The variable is inserted where your cursor is located in the notification message in the right pane.
 - Drag the variable from the left pane to the location you want to place it in the notification message in the right pane.
 For more information, see the [List of variables that can be used to create notifications](#).
4. Click **OK** to enter the message.
5. After you define the message sent with the notification, click **Next** to [Associate Tasks with the Notification](#).

The following table describes how to enter the information in the **Notification Message** page.

Table 18–2 Creating a Notification Message

To where:	Notification On Message	Notification Off Message
This column describes where the message is sent. For more information, see Define the Notification Distribution Properties .	The Notification On Message is sent when the replication task meets the conditions for the notification to be sent. For more information, see Define the Action that Triggers the Notification, Define Which Changes of Status Trigger the Notification, and Define Errors That Trigger the Notification .	The Notification Off Message is sent when the replications task returns to its normal state. This type of message is sent for notifications about latency, disk utilization, and memory utilization. For more information, see Define Which Changes of Status Trigger the Notification .

Table 18–2 (Cont.) Creating a Notification Message

To where:	Notification On Message	Notification Off Message
<p>Console:</p> <p>The messages in this row are sent to the Amazon RDS Migration Console. They are displayed in the Messages section for a specific task. This section is displayed in:</p> <ul style="list-style-type: none"> ■ The Monitor for a specific task. For more information, see Reading Messages about a Task. ■ The right pane of the Tasks page. For more information, see The Tasks View. <p>Note: This message is also sent to the Windows Event log if you select this option. For more information, see Define the Notification Distribution Properties.</p>	<p>In this field, you can edit, change or delete the message that is sent to the Amazon RDS Migration Console when the replication task meets the conditions for the notification to be sent.</p> <p>Example:</p> <pre>[{{SERVER_NAME}}\{{NOTIFICATION_NAME}}]{{TASK_NAME}} replication task latency exceeds defined limits. Current latency is {{LATENCY}} seconds.</pre> <p>This message is sent to the console when latency reaches a value higher than the value you defined.</p>	<p>In this field, you can edit, change, or delete the message that is sent to the Amazon RDS Migration Console when the replication task returns to the normal range as you defined when you Define Which Changes of Status Trigger the Notification.</p> <p>This field is relevant only for notifications about latency, disk utilization, and memory utilization.</p> <p>Example:</p> <pre>Latency is back to normal, latency is {{LATENCY}} seconds</pre> <p>This message is sent when latency returns to within its normal limits.</p>
<p>Email Subject:</p> <p>This is the subject of the email messages sent for the notification.</p> <p>See Define the Notification Distribution Properties for information about sending a notification as an email.</p>	<p>In this field, you can edit, change or delete the subject line for an email that is sent when the replication task meets the conditions for the notification to be sent.</p> <p>Example:</p> <pre>[{{SERVER_NAME}}\{{NOTIFICATION_NAME}}]{{TASK_NAME}} high latency notification</pre> <p>This is the subject for an email message sent when latency reaches a value higher than the value you defined.</p>	<p>In this field, you can edit, change or delete the subject line for an email that is sent when the replication task returns to the normal range as you defined when you Define Which Changes of Status Trigger the Notification.</p> <p>This field is relevant only for notifications about latency, disk utilization, and memory utilization.</p> <p>Example:</p> <pre>RDS Migration Tool notification '{{NOTIFICATION_NAME}}' for task '{{TASK_NAME}}'</pre> <p>This is the subject for an email message sent when latency returns to within its normal limits.</p>

Table 18–2 (Cont.) Creating a Notification Message

To where:	Notification On Message	Notification Off Message
<p>Email Message:</p> <p>This is the body of the email message sent for the notification.</p> <p>See Define the Notification Distribution Properties for information about sending a notification as an email.</p>	<p>In this field, you can edit, change or delete the message that is sent by email when the replication task meets the conditions for the notification to be sent.</p> <p>Example:</p> <p>The latency for replication task <code>{{TASK_NAME}}</code> exceeds defined limits.</p> <p>The current latency is <code>{{LATENCY}}</code> seconds.</p> <p>----- ----- -----</p> <p>This is an automated message generated by Amazon RDS Migration server <code>{{SERVER_NAME}}</code> for notification <code>{{NOTIFICATION_NAME}}</code>.</p> <p>This is an email message sent when latency reaches a value higher than the value you defined.</p>	<p>In this field, you can edit, change, or delete the message that is sent by email when the replication task returns to the normal range as you defined when you Define Which Changes of Status Trigger the Notification.</p> <p>This field is relevant only for notifications about latency, disk utilization, and memory utilization.</p> <p>Example</p> <p>Latency is back to normal, latency is <code>{{LATENCY}}</code> seconds</p> <p>This is an email message sent when latency returns to within its normal limits.</p>
<p>Event viewer</p>	<p>In this field, you can edit, change or delete the message that is sent to the Windows/Linux event viewer when the replication task meets the conditions for the notification to be sent.</p> <p>Note: This field is available only when you select Event log when you Define the Notification Distribution Properties.</p> <p>Example:</p> <pre>[{{SERVER_NAME}}\{{NOTIFICATION_NAME}}]{{TASK_NAME}} high latency notification</pre> <p>The latency for replication task <code>{{TASK_NAME}}</code> exceeds defined limits.</p> <p>The current latency is <code>{{LATENCY}}</code> seconds.</p> <p>This message is sent to the event viewer when latency reaches a value higher than the value you defined.</p>	

After you define the message sent with the notification, click **Next** to [Associate Tasks with the Notification](#).

List of variables that can be used to create notifications

You can use the following variables when creating messages that are sent as notifications:

- TASK_STATUS
- LATENCY
- MEMORY_USAGE
- DISK_USAGE
- COUNT_ACTIVE_TABLES
- ACTIVE_TABLES
- COUNT_ERROR_TABLES
- ERROR_TABLES
- COUNT_ACTIVE_TRANSACTION
- COUNT_DATA_ERRORS
- TLOADED_RECORDS
- CHANGES_RECORDS
- COMMAND
- COMMAND_CODE
- COMMAND_PARAMETERS
- FULLLOAD_COUNT_REQUESTED_TABLES
- FULLLOAD_COUNT_COMPLETED_TABLES
- FULLLOAD_COUNT_ERROR_TABLES
- FULLLOAD_REQUESTED_TABLES_LIST
- FULLLOAD_COMPLETED_TABLES_LIST
- FULLLOAD_ERROR_TABLES
- TABLE_NAME
- TABLE_OWNER
- RECORD_COUNTER
- ERROR_TEXT
- ERROR_CODE
- SQL_STMT
- TASK_NAME
- NOTIFICATION_NAME
- TABLE_COUNT_APPLY_ERRORS

Associate Tasks with the Notification By default, notifications are sent for all tasks that are defined in the Amazon RDS Migration Tool instance you are using. You can determine whether to send the notification to specific tasks defined in the Amazon RDS

Migration Tool instance you are using. For example, you can define a different latency rate for a specific task that is replicating from a slow system.

To associate the notification with tasks

1. In the [New Notification Wizard, Apply To Tasks](#) page, select one of the following:
 - **All Tasks:** To associate this notification with all tasks that are defined in the Amazon RDS Migration Tool instance you are working with. In this case all tasks that were previously defined and any future task will be associated with this notification.

If you choose to associate this notification with All Tasks, then click Next to [Review the Notification Rule](#).
 - **Selected Tasks:** To associate this notification with one or more specific tasks only. Continue with the next step.
2. Select the check box next to any of the tasks you want to associate with this notification. You can select one or more tasks.

Note: The Task check box at the top of the check-box column lets you select all of the tasks that are displayed. When you select this check box it is as if you select each of the tasks individually. Therefore, if you add tasks in the future they will not be included.
3. Click **Next** to [Review the Notification Rule](#).

Review the Notification Rule The **Summary** page lets you review the notification rule that you defined so that you can determine whether the selections you made in the wizard are correct. If you want to make changes, click **Back** and go to the page or pages you want to change.

When you are sure that the notification rule is defined in the way that you want, click **Finish** to close the wizard and add the rule to the [Notification List](#).

After you close the wizard, make sure to click **Save** at the top of the **Settings** page. This will save the information for *all* settings, not only for the notification rule that you created. If you made changes that you do not want to keep, click **Discard** to discard all changes before you make changes to any of the other settings.

Using the Notification List

The Notification List lists all of the notification rules that are defined for the Amazon RDS Migration Tool instance you are working with. It has the following information:

- **Name:** Displays the name of the notification rule.
- **Condition:** Displays the condition that triggers the notification to be sent. For more information, see [Define Which Changes of Status Trigger the Notification](#).
- **Send Message To:** Displays custom users that receive the message. For more information, see [Determine the Email-Message Recipients for the Notification](#).
- **Tasks:** Displays the tasks that are associated with this notification rule. For more information, see [Associate Tasks with the Notification](#).
- **Active:** Select the check box in this column to activate the notification. If this check box is cleared, notifications defined by this rule are not sent. This check box is selected by default.

Editing a Notification

You can make changes to any notification rule.

To edit a notification rule

1. From the [Notification List](#) select the notification you want to edit.
2. Click **Open** (at the top of the list).
or
Double-click the notification you want to edit.
The Edit Notification Rule wizard opens.
3. Make any changes you need in the wizard. For information on how to work with each of the pages in the New Notification Rule wizard, see [Creating a New Notification](#).

Notes:

You can only make changes to those sections that you defined when [Creating a New Notification](#).

- You cannot change name of the notification.
 - If you defined a notification to let you know when the task or full load started or stopped, this cannot be edited. For example, if you created a notification rule for starting a task and you now also want to get notified when the task stops, you must create a new notification rule.
 - In the Notify When? page, you can make changes to the data you defined in the original notification rule. For example, if you defined a **Memory utilization** message in the **Notify when?** page, **Performance/Resources** section, you can only change this parameter. If you want to add information about something that was not defined in the original notification rule, for example you want to add errors to your notification or you want to get information about latency, you must create a new notification rule.
-
-

Deleting a Notification

You can delete notification rules that you no longer want to use.

To delete a notification

1. From the [Notification List](#) select the notification you want to delete.
2. Click **Delete** (at the top of the list).

Note: When you delete a notification, it is deleted permanently.

Setting up Mail Parameters

The Mail parameters define the mail server used to send notifications.

To set the Mail parameters, click the **Mail Settings** sub-tab and enter the following information.

Mail server: Type the outgoing mail server you are using to send the notifications that you define in Amazon RDS Migration Tool, for example, `smtp.example.com`.

Port: Type the port number where the mail server is located. The default value is 25.

Use SSL: Select this check box if you want to use SSL security to connect to the mail server for the notifications that are sent.

Anonymous login: Check this to allow an Amazon RDS Migration Tool user to access the mail server to receive messages without having to provide any user credentials.

User name: Type an email user name for the user account that is sending the notifications. For SMTP authentication be sure to supply a valid user name.

Password: Type the password for the email user account that is sending the notifications. For SMTP authentication be sure that password provided is valid.

Sender email address: Enter the email address that sends the email notifications. This is the address that appears in the **From** field of the email notification.

Send test email: Click **Send test email** to open the Send Test Email dialog box.

In the **Email address for test email**, type an email address to receive a test email message from the server you configured. Use this to determine that the **Mail Parameters** you defined are valid.

When you finish typing the information, click **Save** at the top of the screen to save all of the changes you made.

Note: Click **Save** to save the information for *all* settings, not only for the recipient list. If you made changes that you do not want to keep, click **Discard** to discard all changes before you make changes to any of the other settings.

Creating a Default Recipient List

Click the Default Recipient List sub-tab to create a default recipient list.

A default recipient list is a list of recipients that receive all of the notifications that you define for any task created in the Amazon RDS Migration Tool instance you are working with. This allows you to use one list for all email notifications without having to define the list each time you create a notification.

Note: You can choose to send notifications to a different list or to additional users for any specific notification. You define these exceptions when you create the specific notification. For more information, see [Define the Notification Distribution Properties](#).

To create a Default Recipient List

1. At the top of the **Default Recipient List** settings page, click **Add Recipient**.

The next row in the Recipient List table becomes available.

2. Type the name of the user you want to add to the list of default recipients. Continue to enter a name and email address for each recipient you want to include in the default list.
3. Press the [tab] key or double click in the in the **Email** cell, then type the email address for the user you entered in the **Name** cell.

4. Click **Save** at the top of the screen to save all of the changes you made.

Note: Click **Save** to save the information for *all* settings, not only for the recipient list. If you made changes that you do not want to keep, click **Discard** to discard all changes before you make changes to any of the other settings.

License Settings

You need to register the software before you can use Amazon RDS Migration Tool. Your Amazon vendor should provide you with a text file called `license.txt`. This file contains details such as the product expiration date (if any).

Use the License settings page for:

- [Requesting a License](#)
- [Registering a License](#)
- [Viewing a License](#)

To open the License settings page

From the **Server** view, click **License** from the menu list at the left. The **License** sub-tab is displayed.

Requesting a License

You must have a valid license to work with Amazon RDS Migration Tool. You can request a license from the License settings page in the Amazon RDS Migration Console. In the License Request dialog box, fill out the required information and submit the request by email. Once your request is approved, the license file is sent to you by email. To use Amazon RDS Migration Tool, register the license by using the procedure described in [Registering a License](#).

To request a license

1. From the **Server** page, click **License**.
2. At the top of the **License** tab, click **Request License**.
The Replication License Request dialog box opens.
3. Enter the requested information:
 - **Request type:** Select one of the following:
 - **New License:** Select this if this is your initial license request.
 - **Extend License:** Select this if you have a license and want to extend its period of validity.
 - **Alter License:** Select this if you want to make changes to an existing license. For example, if you want to add additional sources or targets or change the host computer.
 - **License to:** Type the name of the company or group that is requesting a license to use Amazon RDS Migration Tool.
 - **License type:** Select one of the following:
 - **Permanent:** Select this if the license will always be valid. Permanent licenses do not require an expiration date.

- **Evaluation:** Select this if you are requesting a temporary license to use Amazon RDS Migration Tool for a trial period.
 - **Term:** Select this if you are requesting a license that is valid for a specific period of time. In this case you must be sure to include an expiration date in your request.
 - **Expiration date:** Click in this field to select the expiration date using the pop-up calendar. This is required only if you selected **Evaluation** or **Term** in as the License type.
 - **Hosts:** Type the name of the local computer where Amazon RDS Migration Tool is installed. By default the name of the local computer is displayed in this field. You can change this or add additional computers if you are installing Amazon RDS Migration Tool in a different or an additional location.
 - **Source Types:** Click **Edit** to open the Edit Source Types dialog box. Check the database types you are working with as your replication sources. You can select one or more database endpoints as necessary. If you need to work with all available databases, click **All**.
 - **Target Types:** Click **Edit** to open the Edit Target Types dialog box. Check the database types you are working with as your replication targets. You can select one or more database endpoints as necessary. If you need to work with all available sources, click **All**.
4. Click **Send by Mail** to open an email request for the license. Send the email to the address entered in the recipient field of your default email client.

Click **Copy to Clipboard** to copy the information to the computer's clipboard. You can paste this information into the Advanced license request and edit it as necessary. For more information, see [Using the Advanced License Request Option](#).

Using the Advanced License Request Option

The advanced license request option lets you request a license by manually typing the information. Make sure to include all of the information required as described in [Requesting a License](#). The following is a suggested format for the advanced option:

```
Request type:New License
License to: <company name>
License type: Permanent
Expiration date:
Hosts: bee01-xp.company.local
Source Types: Oracle
Target Types: SQLServer
```

Registering a License

You must have a valid license to work with Amazon RDS Migration Tool. If you did not receive a license.txt file, you can request a license using the procedure described in [Requesting a License](#). Once you receive the license, you must register it to work with Amazon RDS Migration Tool.

To register a license

1. Copy the `license.txt` file to your computer or any computer in your network you have access to.
2. From the Server page, click **License**.
3. At the top of the **License** tab, click **Register License**.

The Register License dialog box opens.

4. Click **Load** and browse to find and select the license file.

The license text is displayed in the dialog box as shown above. Check to be sure that the details are correct.

5. Click **Register License** to register the license. A message indicating the license was registered successfully is displayed.

Note: A message is displayed at the top of the Amazon RDS Migration Console that indicates that you have a valid license and when it expires. If the license is expired or invalid, the message indicates this.

You can also click on this message link to request, register, or view license information.

Viewing a License

You can view the license information in the Amazon RDS Migration Console at any time.

To view the license information

From the Server page, click **License**.

The **License** tab is displayed. All of the license information is displayed in the **License** tab.

Error Handling

You can configure how Amazon RDS Migration Tool responds to specific types of errors. You can define error handling on the task level or the server level. The configurations you make in the Server Settings affect all tasks created for this instance of Amazon RDS Migration Tool unless you define a task to use the definitions you create for that task. For information on how to configure error handling for a specific task, see [Error Handling](#) in the [Customizing Tasks](#) chapter.

To open the Error Handling page

From the Server view, click **Global Error Handling** from the menu list at the left.

The following tabs are available:

- **Environmental Errors:** An error that is caused by an environmental problem in the source or target database or on the network. Environmental errors can be restarted.

The information you enter in this tab are the same as the information you enter in the **Environmental Errors** tab for tasks. For information about the options available in this tab, see [Environmental Errors](#) in the Customizing Tasks chapter.

- **Data Error:** An error related to data processing at the record level.

The information you enter in this tab are the same as the information you enter in the **Data Error** tab for tasks. For information about the options available in this tab, see [Data Error](#) in the Customizing Tasks chapter.

- **Table Error:** An error in processing data or metadata for a specific table. This only includes general table data and not an error that relates to a specific record.

The information you enter in this tab are the same as the information you enter in the **Table Error** tab for tasks. For information about the options available in this tab, see [Table Error](#) in the Customizing Tasks chapter.

- **Apply Error:** Errors that occur when the target database is not synchronized with the source database when processing changes. This can cause duplicate key errors on INSERT operations or zero rows affected on UPDATE/DELETE operations.

The information you enter in this tab are the same as the information you enter in the **Apply Error** tab for tasks. For information about the options available in this tab, see [Apply Conflicts](#) in the Customizing Tasks chapter.

For additional information about errors handling in Amazon RDS Migration Tool, see the [Error Handling](#) appendix.

Logging Settings (Server)

The following topics describe the server logging settings:

- [Set the Server Logging Level](#)
- [Configure Log Files](#)

To view and configure the server Logging Settings

From the Server view, click **Logging** from the menu list at the left. When you click the Logging tab, the following sub-tabs are displayed:

- **Logging Level:** Lets you [Set the Server Logging Level](#).
- **Log File:** Lets you [Configure Log Files](#).

Set the Server Logging Level

You set the logging level for server logs in the Server view. The level you set determines what information is written to the server log. This log provides information about the Amazon RDS Migration Server instance you are working with and not for individual tasks. For information on configuring the task logs, see [Working with Task Logs at Runtime](#).

The following are the available logging levels. The list is in order from the lowest level to the highest level.

1. Error
2. Warning
3. Info
4. Debug

The higher levels always include the messages from the lower levels. Therefore, if you select **Error**, only error messages are written to the log. However, if you select **Info**, informational messages, warnings, and error messages are included. Selecting **Debug** writes all possible messages to the log.

You can set the logging level generally for all log components. In addition, you can create a separate logging level for each component.

To set the server logging level

1. On the left side of the Server view, click **Logging** then click the **Server Log** sub-tab.
The **Logger settings** sliders are displayed.

2. Move the **Logger settings** slider to the log level you want. This sets the logging level for all server log modules. Note that all of the sliders for the individual modules move to the same position that you set in the main slider.
3. Make any changes to the sliders for the individual modules. This is optional. Note that if you change the main slider, all of the individual sliders are reset to the new position. If you want to maintain a different logging level for a specific module, you need to reset it.
4. Click **Save** at the top of the screen. Changes to the logging level take place immediately. There is not need to restart the Amazon RDS Migration Server.

Note: Click **Save** to save the information for *all* settings, not only for the logging settings. If you made changes that you do not want to keep, click **Discard** to discard all changes before you make changes to any of the other settings.

Note: To view the log files, open the **Log File** sub-tab.

For more information, see [View a Server Log](#).

Configure Log Files

The **Log File Management** sub-tab lets you define how to handle the physical log files in the system and to view the physical log file. You can define the following.

- [Schedule Maintenance Operations for Log Files](#)
- [Carry out Log File Actions](#)
- [View a Server Log](#)

Schedule Maintenance Operations for Log Files

You can schedule when a log file is rolled over or when to clean old log files from the system. The following are the scheduling options you can define in the **Log File** sub-tab.

Enable automatic roll over: Select this check box if you want to determine when to stop logging to a specific file and begin to log to a new log file. The current log file is called `repserv` and saved (older) log files have the file name `repserv_XXXXXXXXXXXX` where `XXXXXXXXXXXX` represents a 12-digit timestamp.

- **Roll over the log if the log file is older than (days):** Use the counter or type in the number of days that to use a specific log file. After the specified number of days the old log is saved with a timestamp appended to its name and a new log file is started.

The default value is 0 days.

- **Roll over the log if the log file is larger than (MB):** Use the counter or type in the maximum amount of megabytes for a specific log file. When the log file reaches the specified size, the old log is saved with a timestamp appended to its name and a new log file is started.

The default value is 0 megabytes.

Enable automatic cleanup: Select this check box if you want to define the maximum number of days or what size a log file can be before it is deleted.

- **Delete log files that are more than (days):** Use the counter or type in the maximum number of days to keep a log file. Log files that are older than the

number of days entered are automatically deleted from the system. For example, if you enter **4**, then all log files are deleted on the fifth day.

The default value is 3 days.

- **Delete log files that are larger than (MB):** Use the counter or type in the maximum amount of megabytes for a specific log file. When the log file reaches the specified size, it is deleted from the system.
- The default value is 10 megabytes.

Carry out Log File Actions

You can roll a log file or delete one or more old log files at any time.

To carry out log file actions, from the top of the **Server Log** sub-tab, click one of the following:

Roll log file: Click this button to immediately roll the current server log file. The old log file is saved with a timestamp appended to its name.

Delete Logs: Click this button delete the log files.

View a Server Log

You can view the information written to the log in the log viewer.

To view the log in the log viewer

1. From the top of the **Logging** tab (**Server Log** or **Log File Management** sub-tabs), click **View Log**. The server log is displayed.
2. To view a different log file, select the file you want to view from the **Select Log** list at the top right of the window.

File Transfer Service

The Amazon File Transfer Service (FTS) is a robust and reliable file transfer engine designed to efficiently transfer files over the WAN. This can dramatically improve transfer speeds when the source database and the target database are located on different LANs.

How it Works

A solution using FTS consists of two Amazon RDS Migration Servers: A local Amazon RDS Migration Server installed on the source database LAN and a remote Amazon RDS Migration Server installed on the target database LAN.

A local task on the local server is defined from the source database to a File Channel target. A remote task on the remote Amazon RDS Migration Server is defined from a File Channel source to the target database.

The FTS runs on the remote Amazon RDS Migration Server only and transfers the File Channel files from the storage location defined in the local task to the storage location defined in the remote task.

Upon file transfer, and before [Compression](#) and [Encryption](#), large files are split into smaller blocks which form recoverable transport units, and small files are merged into bigger blocks to be sent at the same time. The blocks are then transferred and reconstructed into File Channel files when received by the FTS server.

For information on setting up a File Channel source or target to use FTS, see [Using Advanced Properties for a File-Channel Source](#) and [Using Advanced Properties for a File-Channel Target](#) respectively.

Compression

File Channel files are compressed upon sending using GZIP. You can disable the compression and control the compression level.

Encryption

After compression, File Channel files are encrypted using a randomly generated AES-256 session key. The session key is exchanged between the client and server using the Diffie-Hellman key exchange protocol which is authenticated using a secret key that is shared between the client and the server.

Note: The File Transfer Service should be configured on the remote Amazon RDS Migration Server *only*.

Defining a File Transfer Service

Define a file transfer service as described below.

To add a File Transfer Service:

1. Switch to **Server** view as described in [The Server View](#).
2. In the left side of the **Server** view, click **File Transfer Service**.

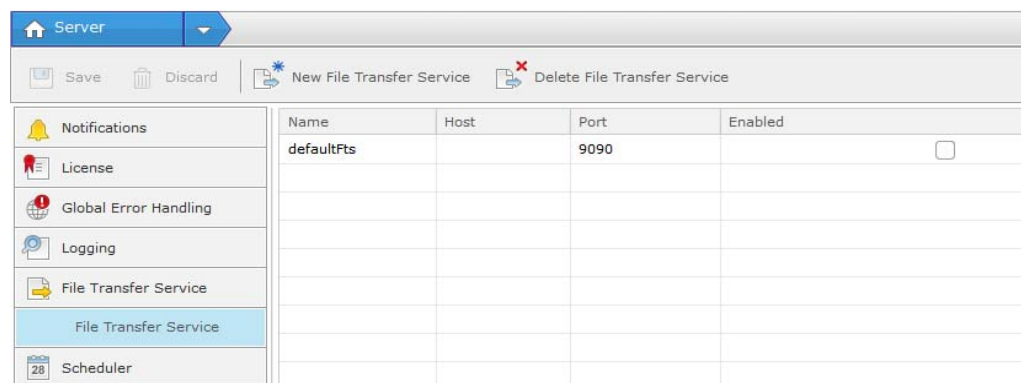
The File Transfer Service list is displayed.

3. In the Actions toolbar, click **New File Transfer Service**.

A new row is added to the File Transfer Service list.

4. Edit the values in the **Name**, **Host** and **Port** columns as follows (click the cells to make them editable):
 - **Name:** The name of the File Transfer Service.
 - **Host:** The host name or IP address of machine on which the remote Amazon RDS Migration Server is installed. The default is 0.0.0.0 (all interfaces). If the server has multiple NICs (Network Interface Cards), you can define a different File Transfer Service for each card.
 - **Port:** The port through which the File Channel files are received.
5. To enable the File Transfer Service, select the check box in the **Enabled** column.
6. Click **Save** to save your settings.

Figure 18–1 The File Transfer Service List



Name	Host	Port	Enabled
defaultFts		9090	<input type="checkbox"/>

Editing a File Transfer Service

You can edit a File Transfer Service as described below.

To edit a File Transfer Service

1. Select the File Transfer Service you want to edit.
2. Edit the values in the **Name**, **Host** and **Port** columns as follows:
 - a. Click the cell to make it editable.
 - b. Change the value as required and then click **Save**.

Note: When you edit a File Transfer Service, make sure that any File Channel targets configured to use the File Transfer Service are also updated accordingly. For more information on File Channel Targets, see [Setting up a File Channel Target using Amazon RDS Migration Tool](#).

Deleting a File Transfer Service

You can delete File Transfer Services that you no longer want to use.

To delete a File Transfer Service

1. In the File Transfer Services List, select the item you want to delete.

2. In the Actions toolbar, click **Delete File Transfer Service**.

Scheduling Jobs

Use the Amazon RDS Migration Tool Scheduler to schedule a one-time job or a recurrent job for specific task operations. A job is essentially an operation that can be scheduled to occur once, daily, weekly or monthly.

The following operations can be scheduled:

- Run a task
- Stop a task
- Reload a task

To schedule a new job:

1. Switch to **Server** view as described in [The Server View](#).
2. In the left side of the **Server** view, click the **Scheduler** tab.

The Scheduler tab has two sub-tabs: **Active Jobs** and **Expired Jobs**. The **Active Jobs** tab contains a list of jobs that are scheduled to run periodically or once only while the **Expired Jobs** tab contains a list of jobs that have already run.

3. Click the **New Scheduled Job** toolbar button.

The **New Scheduled Job** wizard opens.

4. In the **Schedule what** screen, specify a **Job Name** and then choose one of the following as appropriate:

- **Run task** to run the task(s) at the scheduled time.

Note: For Full Load only tasks, it is preferable to select **Reload task** rather than **Run task** when the scheduling is set to Daily, Weekly or Monthly. This will update the table's load whereas **Run task** will replace the existing table.

- **Stop task** to stop the task(s) at the scheduled time.

- **Reload task** to reload the task(s) at the scheduled time.

5. Click **Next**. In the **When?** screen, define when you want the job to run. The following options are available:
 - Once - The first option
 - Daily - The second option
 - Weekly - The third option
 - Monthly - The fourth option

Optionally, select the **Run job as soon as possible after scheduled time is missed** check box. If the Amazon RDS Migration Server is offline for whatever reason (e.g. for maintenance), any jobs that were scheduled to run during that time, will be submitted after the machine is brought back online (at the earliest opportunity).

6. Click **Next**. In the **Apply to tasks** screen, select which tasks to schedule. Select **All tasks** to apply the job to all current and future tasks or **Selected tasks** to apply the job to specific tasks. If you chose **Selected tasks**, select which tasks to apply the job to.
7. Click **Next**. Review your settings and then click **Finish**.

To enable or disable a scheduled job:

- In the **Active Jobs** tab, select or clear the check box in the **Enabled** column as required.

To edit a scheduled job:

1. Select the job in the **Active Jobs** or **Expired Jobs** list.
2. Double-click the job or click the **Open** toolbar button and edit the job as required.

To delete a scheduled job:

1. Select the job in the **Active Jobs** or **Expired Jobs** list.
2. Click the **Delete** toolbar button.

A



Error Handling

Amazon RDS Migration Tool must handle different types of errors during its operation. The way the system should respond to these errors depends on several aspects, including the component where the error occurred, the type of error, and the scope of the error. Because different sites may have different requirements for error behavior, Amazon RDS Migration Tool lets you configure the Error Handling.

Error Types

This topic provides information on the different types of errors that you may encounter using Amazon RDS Migration Tool. The following is a list of the error types:

- **Environmental Errors:** An error that is caused by an environmental problem in the source or target database or on the network. Some examples of environmental errors are loss of communication with the source or target database, restarting a database, or network problems.
- **Data Errors:** An error related to data processing at the record level. Some examples of data errors are conversion errors, errors in transformations, or bad data.
- **Table Errors:** An error in processing data or metadata for a specific table. This only includes general table data and not an error that relates to a specific record.
- **Apply Errors (conflicts):** Errors that occur when the target database is not synchronized with the source database when processing changes. This can cause duplicate key errors on INSERT operations or zero rows affected on UPDATE/DELETE operations.
- **Fatal Errors:** An error that is not related to a specific table but is not environmental. Some examples of fatal errors are incorrect database configuration or incorrect password information.
- **Abnormal Termination:** An error that is returned when the task terminates for an abnormal or unknown reason.

Error Handling Properties

This topic describes the properties that you can set to determine how errors are handled. For information on how to configure error handling behavior in Amazon RDS Migration Tool, see [Error Handling in Customizing Tasks](#) or [Error Handling in Amazon RDS Migration Server Settings](#). See the following topics for information about configuring the error behavior for each type of error.

- [Environmental Errors](#)
- [Data Errors](#)

- [Table Errors](#)
- [Apply Errors](#)
- [Fatal Errors](#)
- [Abnormal Termination](#)

Environmental Errors

When an environmental error occurs, the system will attempt to restart the task if it stops. You can define exactly how and when the system will try to restart a task.

The following table describes the properties used to determine how to handle environmental errors.

Table A-1 Properties to Define Environmental Error Handling

Name	Description	Default
RETRY_COUNT	The maximum number of attempts made to restart a task when an environmental error occurs. After the system attempts to restart the task the designated number of times, the task is stopped and manual intervention is required. Enter -1 to continuously attempt to restart a task. Enter 0 to never attempt to restart a task.	0
RETRY_INTERVAL	The number of seconds that the system waits between attempts to restart a task.	5
RETRY_THROTTLING	When active, the interval between attempts to restart a task gets longer each time.	Y
RETRY_THROTTLING_MAX	The maximum number of seconds to wait between attempts to restart a task if RETRY_THROTTLING is on.	1800

Data Errors

When a data error occurs (to a specific record), you can define what you want the task to do and how to handle the error.

The following table describes the properties used to determine how to handle data errors.

Table A–2 Properties to Define Data Error Handling

Name	Description	Default
DATA_TRUNCATION_ERROR_POLICY	<p>Determine what to do when data is truncated in one or more records. You can:</p> <p>IGNORE_RECORD: The task continues and the data for that record is ignored. The error counter for the DATA_ERROR_ESCALATION_COUNT property is increased so that if you set a limit on errors for a table, this error will count toward this limit.</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its data is not replicated.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	LOG_ERROR
DATA_ERROR_POLICY	<p>Determine what to do when an error occurs to the data from a specific record being replicated. You can:</p> <p>IGNORE_RECORD: The task continues and the data for that record is ignored. The error counter for the DATA_ERROR_ESCALATION_COUNT property is increased so that if you set a limit on errors for a table, this error will count toward this limit.</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its data is not replicated.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	LOG_ERROR
DATA_ERROR_ESCALATION_COUNT	<p>Set the maximum number of errors that can occur to the data for a specific record. When this number is reached, the data for the table that contains the error record is handled according to the policy set in the DATA_ERROR_ESCALATION_ACTION.</p>	0, no escalation
DATA_ERROR_ESCALATION_ACTION	<p>Determines the action to carry out when the maximum number of errors (set in the DATA_ERROR_ESCALATION_COUNT) is reached. You can select one of the following options:</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its data is not replicated.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	SUSPEND_TABLE

Table Errors

The following table describes the properties used to determine how to handle errors that occur to the general data or metadata for a specific table. It does not affect the error policy for specific records in the table.

Table A-3 Properties to Define Environmental Error Handling

Name	Description	Default
TABLE_ERROR_POLICY	Determine what to do when an error occurs to the general table data being replicated. You can: SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its data is not replicated. STOP_TASK: The task is stopped and manual intervention is required.	SUSPEND_TABLE
TABLE_ERROR_ESCALATION_COUNT	Set the maximum number of errors that can occur to the general data or metadata for a specific table. When this number is reached, the data for the table is handled according to the policy set in the TABLE_ERROR_ESCALATION_ACTION.	0, no escalation
TABLE_ERROR_ESCALATION_ACTION	Determines the action to carry out when the maximum number of errors (set in the TABLE_ERROR_ESCALATION_COUNT) is reached. You can select one of the following options: STOP_TASK: The task is stopped and manual intervention is required.	STOP_TASK

Apply Errors

The following table describes the properties used to determine how to handle Apply errors. Apply errors occur when the source and target databases are not properly synchronized during a change processing operation. This causes the following record-level issues:

- DELETE or UPDATE operations occur with zero rows affected.
- INSERT operation occurs with a duplicate key.

Table A-4

Name	Description	Default
APPLY_DELETE_CONFLICT_POLICY	<p>Determine what to do when there is a conflict with DELETE operation. You can:</p> <p>IGNORE_RECORD: The task continues and the data for that record is ignored. The error counter for the APPLY_CONFLICT_ESCALATION_COUNT property is increased so that if you set a limit on errors for a table, this error will count toward this limit.</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>FIX_RECORD: The record is corrected so that it shows the correct DELETE operation.¹</p> <p>SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its changes are not processed.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	IGNORE_ERROR
APPLY_INSERT_CONFLICT_POLICY	<p>Determine what to do when there is a conflict with an INSERT operation. You can:</p> <p>IGNORE_RECORD: The task continues and the data for that record is ignored. The error counter for the APPLY_CONFLICT_ESCALATION_COUNT property is increased so that if you set a limit on errors for a table, this error will count toward this limit.</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>FIX_RECORD: The record is corrected so that it shows the correct INSERT operation.¹</p> <p>SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its changes are not processed.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	LOG_ERROR

Table A-4 (Cont.)

Name	Description	Default
APPLY_UPDATE_CONFLICT_POLICY	<p>Determine what to do when there is a conflict with an UPDATE operation. You can:</p> <p>IGNORE_RECORD: The task continues and the data for that record is ignored. The error counter for the APPLY_CONFLICT_ESCALATION_COUNT property is increased so that if you set a limit on errors for a table, this error will count toward this limit.</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>FIX_RECORD: The record is corrected so that it shows the correct UPDATE operation.¹</p> <p>SUSPEND_TABLE: The task continues but data from the table with the error record is moved into an error state and its changes are not processed.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	LOG_ERROR
APPLY_CONFLICT_ESCALATION_COUNT	<p>Set the maximum number of APPLY conflicts that can occur for a specific table during a change process operation. When this number is reached, the data for the table is handled according to the policy set in the APPLY_CONFLICT_ESCALATION_ACTION</p>	0, no escalation
APPLY_CONFLICT_ESCALATION_ACTION	<p>Determines the action to carry out when the maximum number of errors (set in the APPLY_CONFLICT_ESCALATION_COUNT) is reached. You can select one of the following options:</p> <p>LOG_ERROR: The task continues and the error is written to the task log.</p> <p>SUSPEND_TABLE: The task continues but data from the table with the conflict is moved into an error state and its data is not replicated.</p> <p>STOP_TASK: The task is stopped and manual intervention is required.</p>	LOG_ERROR
APPLY_CONFLICT_DURING_FULL_LOAD	<p>Determine whether to load the conflicting data when carrying out a full-load operation after the change processing is complete.</p>	Conflicts are ignored

¹ FIX_RECORD requires full supplemental logging in order to ensure that an UPDATE is not turned into an INSERT. In other cases, FIX_RECORD can cause an async full load of a record similar to the LOB channel.

Fatal Errors

There is no manual configuration for handling fatal errors. In case of a fatal error, the task is stopped and manual intervention is required.

Abnormal Termination

There is no manual configuration for handling fatal errors. In case of abnormal termination, Amazon RDS Migration Tool attempts to restart the task five times. If the task does not successfully restart for at least 15 minutes, it is stopped and manual intervention is required.

Using Encryption

This appendix describes how to set up a secure connection between the Amazon RDS Migration Console and the Amazon RDS Migration Server and covers the following topics:

- [Setting up HTTPS with IIS](#)
- [Setting up HTTPS without IIS](#)
- [Changing the Server Password](#)
- [Protecting RDS Migration Tool Passwords](#)
- [Changing and Protecting the Master Key](#)

Setting up HTTPS with IIS

The Amazon RDS Migration Console supports working over HTTPS to prevent third-party eavesdropping on sensitive information that can be exchanged, such as setting up user names and passwords.

To use HTTPS, the server must be set up with a valid certificate. If your organization has its own certificate, you can set up the browsers you are using with the Amazon RDS Migration Console to trust the certificate.

You must do the following for your browser to trust the certificate.

- **Create a new certificate:**
 - Create a new site binding and assign it as HTTPS on port 443. See information on creating certificates with IIS on the Microsoft Web site ([How to: Create and Install Temporary Certificates in WCF for Transport Security During Development](#)) or see the information provided by the developer of the Web server you are using.

Note: If you are not using IIS, see [Setting up HTTPS without IIS](#) for information on how to create a certificate, then [Connect using a Web browser and add the certificate to the browser](#):

- You need to make sure that port 443 is open on your firewall.
- Make sure that the certificate is installed on the server.
- **Connect using a Web browser and add the certificate to the browser:**

- Connect to the server using a supported Web browser.

For a list of Web browsers supported by Amazon RDS Migration Tool, see [Additional Required Software](#).

- When you connect using HTTPS, you may receive a message stating the certificate was not signed by a known authority.
- For Internet Explorer and Chrome:
 - Select the option to use the certificate anyway. This will be slightly different depending on the Web browser you are using.
 - You may receive a certificate error. Click to open the properties for this error. This error is caused by a mismatched address. Select the option to install the certificate.
 - Place the certificate in the following store:
Trusted Root Certification Authorities
 - For Internet Explorer only, you must also install the certificate in the personal store.
- For Firefox:
 - You will receive a message that the connection is untrusted. Expand **I Understand the Risks**, and then click **Add Exception**.
 - In the Add Security Exception dialog box, click **Confirm Security Exception**.

Setting up HTTPS without IIS

To facilitate initial evaluation and development, Amazon RDS Migration Tool is installed with both HTTP and HTTPS protocols enabled. However, when using Amazon RDS Migration Tool in a production environment, it is strongly recommended to use HTTPS protocol exclusively.

The following section explains how to set up Amazon RDS Migration Tool to work with the HTTPS protocol without using IIS. Working with HTTPS protocol does not require IIS to be installed on the Amazon RDS Migration Console machine and can be implemented using a self-signed certificate, thereby eliminating the need for a certificate issued by a trusted CA.

Note: Before starting the procedure, it is advisable to first check whether your organization's corporate security policy approves the use of self-signed certificates along with the various certificate parameters (encryption type, key lengths, and so on).

Creating and Installing a Self-Signed Server Authentication Certificate

You can use a self-signed certificate to quickly set up a secure Amazon RDS Migration Tool environment for evaluation purposes. Such an environment will be secure as long as the self-signed certificate is trusted only within the context of its usage with a specific server.

To create and install a self-signed server authentication certificate:

1. Create a text file named `self-signed-web-cert.inf` with the following contents:

```
[NewRequest]
Subject="CN={FQDN of your web server host machine, e.g.
amazonrdsmigrationtool.acme.com}"
MachineKeySet=true
RequestType=cert
HashAlgorithm =SHA256
KeyAlgorithm=RSA
KeyLength=2048
KeyUsage="CERT_KEY_ENCIPHERMENT_KEY_USAGE|CERT_DATA_ENCIPHERMENT_KEY_USAGE"
[EnhancedKeyUsageExtension]
OID=1.3.6.1.5.5.7.3.1
OID=1.3.6.1.5.5.7.3.2
```

2. To create and install the certificate, run the following command:

```
$ certreq -new self-signed-web-cert.inf amazonrdsmigrationtool-https-cert.cer
```

The certificate will also be available in the file:

`amazonrdsmigrationtool-https-cert.cer`. This file does not contain any sensitive information, so it is safe to keep it.

Setting up the Self-Signed Certificate on the Server

To set up the certificate:

1. Retrieve the certificate's thumbprint by running the following command:

```
$ certutil -dump amazonrdsmigrationtool-https-cert.cer
```

The command will output detailed information about the certificate.

- a. Locate the line with the certificate **SHA1** hash. It should look something like this:

```
Cert Hash(sha1): fc fc c8 be 55 4d ca c4 ac e7 39 01 8e 04 0d 60 f2 3b a1 26
```

- b. To create the thumbprint, remove the spaces from the hexadecimal string. For example, removing the spaces from the string above would produce the following thumbprint:

```
fcfcc8be554dcac4ace739018e040d60f23ba126
```

2. To associate the certificate with the HTTPS endpoint, run the following command:

```
$ netsh http add sslcert ipport=0.0.0.0:443  
appid="{4dc3e181-e14b-4a21-b022-59fc669b0914}"  
certhash=fcfcc8be554dcac4ace739018e040d60f23ba126
```

Where:

- The `ipport` value should match the `host:port` part of the `url=` attribute of the **ServiceConfiguration.xml** file. For example:

```
<?xml version="1.0" encoding="utf-8"?>  
<ServiceConfigurationxmlns:xsd="http://www.w3.org/2001/XMLSchema"  
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
  name="AmazonRDSMigrationToolConsole"  
  displayName="Amazon RDS Migration Console"  
  description="Web server for Amazon RDS Migration Console Web  
  application."  
  url="https://rds migration  
  tool.acme.com:443/AmazonRDSMigrationToolConsole">  
  ...  
</ServiceConfiguration>
```

`0.0.0.0:443` means that the certificate can be used with all local IP addresses.

- The `appid` value is the IIS Application ID. The value can also be the Amazon RDS Migration Console assembly GUID: `{3477c544-694e-46ee-93d8-ba4b9a0d24e4}`, or any other GUID.
Note: When setting up a port other than 443 for HTTPS, the IIS Application ID should not be used.
- The `certhash` value is the thumbprint value created earlier.

Removing the Self-Signed Certificate from the Server

To remove the self-signed certificate from the server, run the following command:

```
$ netsh http del sslcert ipport=0.0.0.0:443
```

Where the `ipport` value is the same as the one specified when the certificate was added.

Setting Up the Self-Signed Certificate on a Client

Web browsers connecting to a server on which a self-signed certificate is installed have no way of verifying the authenticity of the server certificate. As a result, the browser will warn you that the site poses a potential security threat. To prevent such warnings from appearing, you can instruct the browser to trust the certificate as described below.

For Chrome and Internet Explorer:

1. Either save the certificate from the browser or copy the `amazonrdsmigrationtool-https-cert.cer` file that you created earlier to the client machine.
2. Right-click the certificate file and select **Install Certificate**.
3. Complete the wizard, making sure to selected **Trusted Root Certification Authorities** as the certificate store.

For Firefox:

- Simply add an exception for the certificate when the **Add Security Exception** prompt is displayed.

Changing the Server Password

By default, the Amazon RDS Migration Console establishes a secure connection to the Amazon RDS Migration Server using the default password. The default password should only be used for evaluation purposes. Before using Amazon RDS Migration Tool in a production environment, you should change the default password as described below.

To change the server password (on Windows or Linux):

1. On Windows, run the following command:

```
repctl SETSERVERPASSWORD <new_password>
```
2. On Linux:
 - a. Run the following command:

```
repctl
```
 - b. Press [Enter] and then type the following:

```
SETSERVERPASSWORD <new_password>
```
 - c. Press [Enter] again to set the password.
3. Restart the Amazon RDS Migration Tool services (Windows) or the Amazon RDS Migration daemon (Linux).

Protecting RDS Migration Tool Passwords

RDS Migration Tool stores secrets (i.e passwords and keys) in its internal repository, enabling it to perform secure operations during runtime, such as connecting to a database, connecting to a web server, connecting to an email server, connecting to a remote file transfer server, and so on.

Because this information is so sensitive, RDS Migration Tool applies stringent security measures to protect it. These are described below.

All passwords are stored as strings in the following format:

```
"{Zxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx}"
```

Where:

- 'Z' is a fixed character
- 'x' is the protection method (see below for available methods)

- 'hhhhh...' is a hexadecimal representation of the encrypted secret

Each of the stored secrets is encrypted using the AES-256 algorithm and protected with a 256-bit master key.

The following techniques are used to strengthen the protection:

- The secret is prefixed with a nonce in order to prevent known cipher attacks.
- A salt based on the nature of the secret (e.g. "DB" for a database password) and related data - for example, the database username associated with the password - is used in the encryption. This makes it difficult to use an encrypted password for a purpose other than for what it was intended.
- The 256-bit master key which protects the encrypted secrets is automatically generated on the first use of the product. The master key is stored in a special file located in <product-dir>/bin/**mk.dat** and is encrypted with a fixed key. The procedure for changing the master key as well as measures that can be taken to protect the file containing the master key are described below in [Changing and Protecting the Master Key](#).

Changing and Protecting the Master Key

The following section describes how to change the master key as well as how to protect the master key file from misuse.

Master Key Replacement

You can change the master key at any time. Note that changing the master key will also require you to re-enter the passwords in all endpoints (as they were previously encrypted using the old key).

To change the Master Key:

1. Stop any running tasks.
2. Stop the RDS Migration Tool services.
3. Open a command prompt as an administrator.
4. Change the working directory to ~Amazon\RDS Migration Tool\bin (Windows) or ~Amazon/RDS Migration Tool/bin (Linux) and then issue the following command:

```
repctl setmasterkey <your_new_master_key>
```

Example:

```
repctl setmasterkey 78543043vuiyfyrf64454555jy65
```

5. Start the **Amazon RDS Migration Server** service.
6. Re-enter the access passwords in all endpoints.
7. Start the Tasks.

Protecting the Mk.dat File from Misuse

The mk.dat file should be granted the proper permissions to protect it against unauthorized access. The following protection methods are available for protecting mk.dat from misuse:

- **Machine Profile** - In this case, the master key is further encrypted with the machine profile (available only on Windows). The mk.dat file can only be used

on the same machine on which it was created. This method is useful when multiple local accounts need to use RDS Migration Tool (e.g., the system account for the service and user X for command-line work).

- **User Profile** - In this case, the master key is further encrypted with the user profile which was used to set the master key (typically the user who installed the product). The mk.dat file can only be used by the same user account. If the user account has a roaming profile then the master key can be used on another machine as long as RDS Migration Tool is run from the same account. This method is useful for clustering scenarios.
- **Constant** - In this case, no additional encryption is applied to the mk.dat file so it is important to key it safe. This is the default mode.

Master Key Protection

As well as being able to change the master key, you can also protect it from misuse. Note that changing the master key will require you to re-enter the passwords in all endpoints (as they were previously encrypted using the old key).

To protect the Master Key:

1. Stop any running tasks.
2. Stop the RDS Migration Tool services.
3. Open a command prompt as an administrator.
4. Change the working directory to `~\Amazon\RDS Migration Tool\bin` (Windows) or `~\Amazon/RDS Migration Tool/bin` (Linux) and then issue the following command:

```
repctl setmasterkey <your_new_master_key> master_key_scope=<enum>
```

where `<enum>` can be any of the following:

- 0 - DEFAULT
- 1 - CONSTANT
- 2 - USER
- 3 - MACHINE

For a detailed description of each of the enumerators, see [Protecting the Mk.dat File from Misuse](#) above.

Example:

```
repctl setmasterkey f2544390946gffjhf463434r master_key_scope=2
```

5. Start the **Amazon RDS Migration Server** service.
6. Re-enter the access passwords in all endpoints.
7. Start the Tasks.

Exporting and Importing Tasks

To be able to export tasks from one machine and then import them to another, the same master key must exist on both machines. Meaning that if you change the master key on one machine, you must also change it on the other machine (using the procedure described above).

Note: After you change the master key, the mk.dat file will have administrator permissions. Consequently, in order to export a task, you will also need to open the command prompt as an administrator.

For more information on importing and exporting RDS Migration Tool tasks, see [How to: Export and Import a Task](#).

C



Setting up Amazon RDS Migration Tool in a Cluster Environment

This appendix describes how to set up Amazon RDS Migration Tool in Windows and Linux clustering environments and contains the following topics:

- [Setting up Amazon RDS Migration Tool in a Windows Server Cluster \(HA\)](#)
- [Setting up Amazon RDS Migration Tool in a Linux Cluster](#)

Setting up Amazon RDS Migration Tool in a Windows Server Cluster (HA)

This section describes how to set up Amazon RDS Migration Tool in a Windows Server 2008 or Windows Server 2012 Cluster environment. The following are the required steps for setting up the cluster environment. Please note that these steps provide a general explanation of how to set up the environment to work with Amazon RDS Migration Tool. If you need directions on how to work with specific programs in Microsoft Windows, see the Microsoft documentation.

- [Preparing to Set Up the Cluster Environment](#)
- [Create a Failover Cluster and a Service to Use in Your Environment](#)
- [Add Storage](#)
- [Define Client Access](#)
- [Install Amazon RDS Migration Tool](#)
- [Add the Amazon RDS Migration Tool Services](#)
- [Define the Dependencies for Each Service](#)
- [Enabling Different Console Configurations in a High Availability Environment](#)

Preparing to Set Up the Cluster Environment

You must make sure that your system is ready to use a cluster environment. Carry out the following procedures before you set up your cluster and services.

1. Make sure that the storage, hardware, and hardware settings on each of your servers, are compatible with failover clustering. See the Microsoft documentation for an explanation on how to validate your system.
2. Open the Server Manager console on one of the servers Amazon RDS Migration Tool is installed. You will need to install Amazon RDS Migration Tool on this server and all other servers in your cluster environment. See [Install Amazon RDS Migration Tool](#).

Note: You must make sure that the Failover Clustering snap-in for the Microsoft Management Console is installed on the server where you are configuring the cluster. For more information, see the Microsoft documentation.

3. In the left column of the Server Manager, expand **Features**.
4. Select the **Failover Cluster Manager**.

Create a Failover Cluster and a Service to Use in Your Environment

Carry out this procedure for as many services that you are using.

The first thing you should do is to create a failover cluster that defines the servers or nodes that are included in your cluster environment. You must also create a service to manage the Amazon RDS Migration Tool application and services.

To create a failover cluster and a service

1. Right-click the highest node called **Failover Cluster Manager** and select **Create a Cluster**. Follow the directions in the Create Cluster wizard to create a cluster.
2. Expand the cluster you are working with.
3. Right-Click **Services and applications** and point to **More Actions**, then select **Create Empty Service or Application**. Follow the directions in the Wizard to create a new service. If you need additional information on how to create a new service, see the Microsoft documentation.
4. Right-click the service that you created and select **Properties**. In the **Name** field, change the default name to a name that you want to use for this service. The name should have a relevant meaning so you can remember it in the future.

Note: If necessary, you can carry out any of the options in the menu displayed when you right-click the service.

Add Storage

You must define the disk that you use for storage. Use the storage to hold all of the information in the Amazon RDS Migration Tool data folder.

To add storage

1. Right-click the Storage node and select **Add a Disk**. This disk will contain the data in the data folder. The data folder is located in the root folder where Amazon RDS Migration Tool is installed. This disk is bound to the RDS Migration Tool Services. This means that in case of failover, this disk will be moved to the another cluster running the Amazon RDS Migration Tool Services.
2. On the Storage node, you can add the new disk to the failover cluster.
3. Right-Click one of the services in the cluster and select **Add Storage**. Then select the new disk to add this disk to the cluster you are working with.

Note: Make sure that there is a disk available on the server where you are adding the disk. Also, to ensure that the system recognizes the disk, make sure that the storage configuration for the disk is correct. The disk you want to use must be configured as a basic disk, not a dynamic disk. For more information, see the Microsoft online help.

Define Client Access

A network name must be added so you can access this service from other machines using the Network Name.

To define client access

1. Right-click one of the services that you created and point to **Add a resource** then select **Client Access Point**. The services are sub-nodes of Services and applications.
2. Follow the directions in the New Resource Wizard to set the Network Name. The network name must be the same as the service.

You set the name on the first page (Client Access Point) of the wizard.

Install Amazon RDS Migration Tool

Carry out the following two procedures to install Amazon RDS Migration Tool:

1. Install Amazon RDS Migration Tool on all the computers (nodes) that are part of the cluster you configured (see [Create a Failover Cluster and a Service to Use in Your Environment](#)).

Note: The Amazon RDS Migration Tool destination folder should be local to all of the computers that are part of the cluster. Use the default location `C:\Program Files\Amazon\RDS Migration Tool`.

In the install wizard, set the **data** folder location to a shared storage location, for example, `F:\data`.

For more information, see [Installing Amazon RDS Migration Tool on Windows](#).

2. To install Amazon RDS Migration Tool on the rest of the clusters when the storage is not a shared drive and is managed by the Failover Cluster manager, do one of the following (according to your Windows Server version):
 - **Windows Server 2008 Cluster:** Right-click the service and select **Move this service or application to another node**. Then select the node on which to install Amazon RDS Migration Tool. This ensures that the storage is available on that node.
 - **Windows Server 2012 Cluster:** Right-click the cluster role and select **Move > Select Node**. Then select the node on which to install Amazon RDS Migration Tool. This ensures that the storage is available on that node.

Add the Amazon RDS Migration Tool Services

The two Amazon RDS Migration Tool services must be added as resources to the service (Windows Server 2008 Cluster) or role (Windows Server 2012 Cluster) you added in the [Create a Failover Cluster and a Service to Use in Your Environment](#) procedure.

The Amazon RDS Migration Tool services are called Amazon RDS Migration Console and Amazon RDS Migration Server.

To add the Amazon RDS Migration Tool services

1. Do one of the following (according to your Windows Server version):
 - **Windows Server 2008 Cluster:** Right-click the service you are working with and point to **Add a resource**. Then select **Generic Service**.
 - **Windows Server 2012 Cluster:** In the left pane of the Failover Cluster Manager, select **Roles**. The available roles will be listed in the right pane of the console. Right-click the role you are working with and point to **Add a resource**. Then select **Generic Service**.

2. In the **Select Service** screen of the New Resource wizard, select **Amazon RDS Migration Console** from the List.
3. Click **Next** and follow the directions in the wizard to create the resource. For information on how to use this wizard, see the Microsoft online help.
Note: Amazon RDS Migration Tool must be installed on the computers where you defined the service for the Amazon RDS Migration Tool services to be available in the list.
4. Repeat the same steps for the **Amazon RDS Migration Server**.

Define the Dependencies for Each Service

You should define dependencies for the Amazon RDS Migration Tool services. This allows the Storage and the Network names to start before the Amazon RDS Migration Tool services. If these resources do not start up before the services, Amazon RDS Migration Tool will not start because it will continue to search for the data location.

To define the dependencies

1. Do one of the following (according to your Windows Server version):
 - **Windows Server 2008 Cluster:** In the left pane of the Failover Cluster Manager, select the **Amazon RDS Migration Console** service.
The properties for this service are displayed in the center pane.
 - **Windows Server 2012 Cluster:** In the left pane of the Failover Cluster Manager console, select **Roles**. The available roles will be listed in the right pane of the console. Select the role you are working with and then, in the bottom right pane, select the **Resource** tab. From the list of the available roles, select **Amazon RDS Migration Console**.
2. Do one of the following (according to your Windows Server version):
 - **Windows Server 2008 Cluster:** In the **Other Resources** section, double-click the **Amazon RDS Migration Console** service.
 - **Windows Server 2012 Cluster:** Right-click the **Amazon RDS Migration Console** role and select **Properties**.
The **Amazon RDS Migration Console Properties** dialog box opens.
3. In the **Amazon RDS Migration Console Properties** dialog box, select the **Dependencies** tab.
4. Click **Insert**. A new line is added to the Resource list.
5. In the Resource column, click the arrow and select the RDS Migration Tool Data storage resource from the list.
6. Click **Insert** and add the Network Name resource (its name should be the same as the cluster name).
7. Repeat the steps for the Amazon RDS Migration Server service.
8. Start the Services using the Failover Cluster Manager and access the console using the Network name. See [Define Client Access](#) for information on how to configure the network name.
9. Register the license. The license should contain all host names of the cluster.

Note: To open Amazon RDS Migration Tool Console, it is recommended to use an address that includes the name or IP address of the cluster machine (as opposed to the specific node name).

Example:

`http://cluster_name_ip/AmazonRDSMigrationToolConsole/4.0.5.118/#`

Enabling Different Console Configurations in a High Availability Environment

In a High Availability active-passive scenario, it is recommended to install the Amazon RDS Migration Tool data folder on a shared storage device. The data folder contains various configuration files in XML format (e.g. `LoggingConfiguration.xml`, `ServiceConfiguration.xml`, and so on). In a standard installation, you do not need to change the names of these files. However, in a High Availability environment in which Amazon RDS Migration Console is installed on two different machines, you may want each machine to have its own unique settings. In such a situation, changing the name of these files to include the hostname of the Amazon RDS Migration Console machine (as it appears in the output of the Windows hostname command) will allow you to store a different configuration for each machine.

The file name should be formatted as follows:

```
[Name]Configuration-[hostname].xml
```

For example, let's assume that one Amazon RDS Migration Console is installed on a machine called `rds migration tool-main` and the other Amazon RDS Migration Console is installed on a machine called `rds migration tool-failover`. To set up a different configuration for each machine, simply create two copies of the `ServiceConfiguration.xml` file and then rename them to `ServiceConfiguration-rds migration tool-main.xml` and `ServiceConfiguration-rds migration tool-failover.xml` accordingly. After renaming the files, edit each of them with your desired settings.

Setting up Amazon RDS Migration Tool in a Linux Cluster

This section describes how to set up Amazon RDS Migration Tool in a Linux Cluster Environment. There are several commercially available clustering solutions for Linux including Veritas Cluster Server, Red Hat Cluster Suite and IBM HACMP for Linux.

When one of the available clustering solutions is already in place, Amazon RDS Migration Tool can be set up like any other cluster application while adhering to the following guidelines:

- RDS Migration Tool should be installed on the participating cluster nodes.
For more information, see [Installing Amazon RDS Migration Tool on Linux](#).
- RDS Migration Tool only supports the failover cluster configuration (active-passive).
- Amazon RDS Migration Tool data (the data folder, tasks folder, file channel, etc.) should be stored in a SAN for shared access between the cluster nodes. To change the default location of the Amazon RDS Migration Tool data folder, run the following command on the primary cluster node when the installation completes:

```
./repctl-d <shared_storage_path> service start
```

- Only one instance of RDS Migration Tool can be active at a given data location. The cluster software should be set so that during failover, one RDS Migration Tool instance is stopped and the other is started.

D



Control Tables

This appendix describes the Amazon RDS Migration Tool Control Tables which are created on the target database when the corresponding table is selected in the [Control Tables](#) tab.

The following tables are described:

- [Apply Exceptions](#)
- [Replication Status](#)
- [Suspended Tables](#)
- [Replication History](#)

Note: All Control Table timestamps are in UTC format.

Apply Exceptions

All task-related errors are recorded in the `amazon_apply_exceptions` table, which is described below.

Table D-1 *amazon_apply_exceptions Table*

Column	Type	Description
TASK_NAME	nvchar	The name of the Amazon RDS Migration Tool task.
TABLE_OWNER	nvchar	The table owner.
TABLE_NAME	nvchar	The table name.
ERROR_TIME	timestamp	The time the exception (error) occurred.
STATEMENT	nvchar	The statement that was being executed when the error occurred.
ERROR	nvchar	The actual error.

Replication Status

The `amazon_status` table contains the current status of the replication task and the target data.

Table D-2 *amazon_status Table*

Column	Type	Description
SERVER_NAME	nvchar	The name of the machine on which Amazon RDS Migration Tool is installed.
TASK_NAME	nvchar	The name of the Amazon RDS Migration Tool task.
TASK_STATUS	varchar	One of the following: <ul style="list-style-type: none"> ■ FULL LOAD ■ CHANGE PROCESSING Task status is FULL LOAD as long as there is at least one table in full load. After <i>all</i> tables have been loaded, the task status changes to CHANGE PROCESSING.
STATUS_TIME	timestamp	When the status was last updated.
PENDING_CHANGES	int	The number of change records that were not applied to the target.
DISK_SWAP_SIZE	int	The amount of disk space is occupied by old or offloaded transactions.
TASK_MEMORY	int	Current memory consumption in MB.
SOURCE_CURRENT_POSITION	varchar	The POSITION in the source database that Amazon RDS Migration Tool is currently reading from.
SOURCE_CURRENT_TIMESTAMP	timestamp	The TIMESTAMP in the source from which Amazon RDS Migration Tool is currently reading.
SOURCE_TAIL_POSITION	varchar	The POSITION of the oldest start transaction that is still not committed. This represents the newest position that you can revert to, without losing any changes. There may, of course, be duplicates.
SOURCE_TAIL_TIMESTAMP	timestamp	The TIMESTAMP of the oldest start transaction that is still not committed. This represents the newest TIMESTAMP that you can revert to, without losing any changes. There may, of course, be duplicates.
SOURCE_TIMESTAMP_APPLIED	timestamp	This is the timestamp of the last transaction commit. In a bulk apply this will be the timestamp for the commit of the last transaction in that batch. It will only be changed as part of the last transaction in the batch.

Suspended Tables

The following section describes the `amazon_suspended_tables` table.

Table D-3 *amazon_suspended_tables Table*

Column	Type	Description
SERVER_NAME	nvchar	The name of the machine on which Amazon RDS Migration Tool is installed.
TASK_NAME	nvchar	The name of the Amazon RDS Migration Tool task.

Table D-3 (Cont.) amazon_suspended_tables Table

Column	Type	Description
TABLE_OWNER	nvchar	The owner of the suspended table.
TABLE_NAME	nvchar	The name of the suspended table.
SUSPEND_REASON	vvarchar	The reason why the table was suspended.
SUSPEND_TIMESTAMP	timestamp	The date and time the table was suspended.

Replication History

The following section describes the amazon_history table.

Table D-4 amazon_history Table

Column	Type	Description
SERVER_NAME	nvchar	The name of the machine on which Amazon RDS Migration Tool is installed.
TASK_NAME	nvchar	The name of the Amazon RDS Migration Tool task.
TIMESLOT_TYPE	vvarchar	One of the following: <ul style="list-style-type: none"> ■ FULL LOAD ■ CHANGE PROCESSING (CDC) When FULL LOAD and CHANGE PROCESSING are running in parallel (some tables in full load, some in CDC), two history records will occupy the same time slot.
TIMESLOT	timestamp	The end timestamp of the time slot.
TIMESLOT_DURATION	int	The duration of each history record in minutes.
TIMESLOT_LATENCY	int	The latency at the end of the time slot. This is only applicable to CDC time slots. Note that this value contains the value of the <i>target</i> latency only.
RECORDS	int	The number of records processed during the time slot.
TIMESLOT_VOLUME	int	The volume of data processed in MB.

Impact of DST Change on Amazon RDS Migration Tool

This appendix describes how Amazon RDS Migration Tool is affected by Daylight Saving Time (DST) and provides guidelines for handling changes brought about by DST.

There are two types of DST changes:

- **DST On** - Occurs approximately when Summer starts (actual date is country specific). Its impact on local time is that local time is moved one hour forward (so, for example, 01:00AM becomes 02:00AM). This DST change does not impact Amazon RDS Migration Tool as it does not result in time overlap.
- **DST Off** - Occurs approximately when Winter starts (actual date is country specific). Its impact on local time is that local time is moved back one hour (so, for example, 02:00AM becomes 01:00AM). This DST change results in time overlap where local time travels over the same hour twice in a row.

The comments below assume that the customer has not changed the time but rather the timezone or the DST setting. Changing the actual time (not for minor time adjustments) is a sensitive operation and is best done when Amazon RDS Migration Tool is stopped.

Running Amazon RDS Migration Tool tasks do not depend on the timezone or DST (daylight saving time) for correctly scanning and processing the transaction logs. Internally, Amazon RDS Migration Tool timers use UTC.

Still, there are several places where DST may have an effect:

1. Timestamps in logs and audit messages are in local time. As a result, when Winter time starts, the logs will show the time going back an hour; conversely, when Summer time starts, the logs may appear to be missing one hour.
2. The global and table manipulation variables `timestamp` and `commit_timestamp` are expressed in local time so their use would also be affected. The impact of this depends on the manipulation done and on the intended use of the timestamp based data. If the backward/forward change prescribed by the timestamp may be disruptive, then it is recommended to stop the Amazon RDS Migration Server for the duration of the first overlap period.

Going in to Winter time, for example, if at 02:00AM the clock is to be set back to 01:00AM then when the time is 00:55AM the Amazon RDS Migration Server should be stopped and, after an hour and ten minutes (at 01:05AM), should be started again.

3. Statistics shown on the console are also sensitive to local time and thus may also show confusing/inaccurate data in the overlap period (going in to Winter time) or for the skipped period (going into Summer time).
4. If the clock on Amazon RDS Migration Server machine is one hour behind the clock on the Amazon RDS Migration Console (UI) machine, the following issues are known to occur:
 - The [Applied Changes](#) circle graph will be updated as the changes are applied, but the information in the [Recent Activity](#) tab will not be updated.
 - Scheduled jobs will start according to the Amazon RDS Migration Server time (as expected), but will remain in the **Active Jobs** list after execution instead of moving to the **Expired Jobs** tab.

For more information on scheduling jobs, see [Scheduling Jobs](#).

In general, it is recommended to avoid non-critical task design changes during the first overlap period (going in to Winter time) so as to prevent confusion about when the changes took place.

In addition to Amazon RDS Migration Tool, other components are also involved including:

- The source database system
- The target database system
- The local operating system
- The task design (specifically using timestamp based variables)

Given the complexity of the topic and the involvement of many independent components and settings, Amazon generally recommends that customers first verify the impact of DST changes in their test environment.

Glossary

Change Data Capture (CDC)

Captures changes in the source data or metadata as they occur and applies them to the target database as soon as possible, in near-real-time. The changes are captured and applied as units of single committed transactions and several different target tables may be updated as the result of a single source Commit. This guarantees transactional integrity in the target database. The CDC process for any file or table starts as soon as the data Load operation for the file or table begins.

Full Load

Creates all defined files or tables at the target database, automatically defines the metadata that is required at the target, and populates the tables with data from the source.

Latency

The time interval between the completion of a transaction on the source database and the appearance of changes in the target. The latency at a target may vary due to various reasons (for example, the target system may be down or inaccessible), but the Amazon RDS Migration Tool attempts to keep the latency at a target in the several-seconds-to-several-minutes range.

Source Database

A collection of files or tables managed by a database management system (such as, Oracle, SQL Server) that is part of the main computing service of the IT organization of an enterprise. This source continuously updated, may need to provide a high throughput rate, may have strict 24/7 up-time requirements, and may reference or update a number of tables in the course of a single logical transaction while providing transactional consistency and integrity for the data.

Target Database

A collection of files or tables managed by a Database Management System (DBMS), which may be different from the DBMS managing the source database. It contains data that is derived from the source. It may contain only a subset of the tables, columns, or rows that appear in the source. Its tables may contain columns that do not appear in the source but are transformations or computations based on the source data.

A

- abnormal termination, 19-7
- Accessing the console, 1-13
- add database, 14-3
 - Amazon Redshift target, 10-5
 - Microsoft SQL Server (source), 6-12
 - Microsoft SQL Server (target), 6-21
 - Microsoft SQL Server source, 6-12
 - Microsoft SQL Server target, 6-21
 - MySQL source, 8-9
 - MySQL target, 8-15
 - ODBC source, 9-6
 - Oracle source, 5-14
 - Oracle target, 5-25
 - PostgreSQL target, 11-12
 - SAP Sybase ASE source, 7-4
 - SAP Sybase ASE target, 7-9
- add databases
 - file channel source, 12-8
 - file channel target, 12-13
 - source, 14-4
 - target, 14-4
- add tables, 14-5
- advance license request options, 18-15
- advanced properties
 - Amazon Redshift target, 10-6
 - file channel, 12-15
 - file channel source, 12-10
 - Microsoft SQL Server source, 6-14
 - Microsoft SQL Server target, 6-23
 - MySQL source, 8-11
 - MySQL target, 8-18
 - ODBC target, 9-8
 - Oracle source, 5-16
 - Oracle target, 5-27
 - PostgreSQL source, 11-8
 - PostgreSQL target, 11-14
 - SAP Sybase ASE source, 7-6
 - SAP Sybase ASE target, 7-12
- alerts, 16-5
- alternate backup folder, 6-15, 6-16
- Amazon RDS
 - Oracle, 5-12
- Amazon Redshift, 10-1
 - advanced properties (target), 10-6

- Amazon Redshift target, adding a, 10-5
- Cloud prerequisites, 10-2
- data types, 10-3
- prerequisites, 10-1
- Amazon Redshift registration, 10-2
- Amazon Redshift requirements, 10-1
- apply changes task settings, 15-50
- apply changes tuning settings, 15-51
- apply errors, 19-4
- apply errors settings, 15-55
- apply global transformation rule, 15-15
- ARCHIVELOG (Oracle source), 5-10

B

- backup
 - Microsoft SQL Server, 6-10
- browsers, 1-2

C

- CDC
 - monitoring, 17-7
 - monitoring tables, 17-8
 - processes, 3-4
- certificates, 20-1
- change processing, 3-4
 - monitoring, 17-7
 - monitoring tables, 17-8
- Cluster
 - Linux, 21-5
- concepts, 3-1
- concepts and design, 3-1
- configuration
 - Microsoft SQL Server source, 6-9
 - Oracle target, 5-24
- configuration properties
 - file channel source, 12-7
 - file channel target, 12-11
 - Microsoft SQL Server source, 6-5
 - Microsoft SQL Server target, 6-17
 - MySQL source, 8-5
 - MySQL target, 8-13
 - ODBC target, 9-2
 - Oracle source, 5-4
 - Oracle target, 5-20
 - SAP Sybase ASE source, 7-2
 - SAP Sybase ASE target, 7-8
- configure log files, 18-18
- configuring error handling behaviors, 19-1
 - abnormal termination, 19-7
 - apply errors, 19-4
 - data errors, 19-2
 - environmental errors, 19-2
 - fatal errors, 19-6
 - table errors, 19-4
- configuring tasks, 15-46
- control tables
 - creating on target, 15-48
- create a notification message, 18-7

- create notifications, 18-2
- creating expressions for transforms, 15-6
- customizing tasks, 15-1

D

- data enrichment, 15-41
- data errors, 19-2
- data errors settings, 15-53
- Data Types, 4-2
- data types
 - Amazon Redshift, 10-3
 - Microsoft SQL Server
 - not supported, 6-9
 - Microsoft SQL Server source data types, 6-7
 - Microsoft SQL Server target data types, 6-19
 - MySQL source data types, 8-6
 - MySQL target data types, 8-13
 - ODBC source, 9-3
 - Oracle
 - not supported, 5-10
 - Oracle source data types, 5-8
 - PostgreSQL, 11-10
 - PostgreSQL source, 11-4
 - SAP Sybase ASE
 - not supported, 7-4, 7-9
 - SAP Sybase ASE source data types, 7-3
 - SAP Sybase ASE target data types, 7-8
- database configuration
 - Microsoft SQL Server source, 6-9
 - Oracle target, 5-24
- databases, 4-1
 - adding, 14-3
 - Amazon Redshift, 10-1
 - editing, 14-3
 - file channel, 12-1
 - files, 12-1
 - managing, 14-2
 - Microsoft SQL Server, 6-1
 - MySQL, 8-1
 - ODBC, 9-1
 - Oracle, 5-1
 - PostgreSQL, 11-1
 - SAP Sybase ASE, 7-1
 - view information, 14-3
 - working with, 14-2
- default recipient list, 18-13
- define changes that trigger the notification, 18-3
- define errors that trigger a notification, 18-5
- define the notification distribution properties, 18-5
- define the notification message, 18-6
- define the transformation rule (global transformation), 15-18
- define trigger to send notification, 18-3
- define what to transform (global transformation), 15-15
- delete global transformation rule, 15-29
- deleting a notification, 18-12
- design mode, 13-4
- designing tasks, 14-1

determine email message recipients for the notification, 18-6

E

edit global transformation rule, 15-29
editing a notification, 18-12
editing database, 14-3
encryption, 20-1
endpoints, 4-1
 Amazon Redshift, 10-1
 file channel, 12-1
 files, 12-1
 Microsoft SQL Server, 6-1
 MySQL, 8-1
 ODBC, 9-1
 Oracle, 5-1
 PostgreSQL, 11-1
 SAP Sybase ASE, 7-1
environmental errors, 19-2
environmental errors settings, 15-53
error handling, 18-16
 error types, 19-1
 properties, 19-1
 abnormal termination, 19-7
 apply errors, 19-4
 data errors, 19-2
 environmental errors, 19-2
 fatal errors, 19-6
 table errors, 19-4
error handling behaviors
 properties, 19-1
 abnormal termination, 19-7
 apply errors, 19-4
 data errors, 19-2
 environmental errors, 19-2
 fatal errors, 19-6
 table errors, 19-4
error handling settings, 15-52, 15-53
error types, 19-1
export tasks, 14-9
expression builder, 15-30
 build expression, 15-32
 data enrichment, 15-41
 date and time functions, 15-39
 evaluate expression, 15-33
 functions, 15-34, 15-37
 global transformations, 15-35
 header columns, 15-46
 input columns, 15-34
 inputs, 15-34
 LOBs, 15-38
 logical operators, 15-35
 mathematical operators, 15-37
 metadata
 expression builder
 using variables, 15-35
 null check functions, 15-39
 numeric functions, 15-38
 operations, 15-45

 operators, 15-34, 15-35
 other functions, 15-45
 overview, 15-30
 stings, 15-35
 strings, 15-38
 test expression, 15-33
 using elements, 15-34
expression builder for transforms and filters, 15-30

F

fatal errors, 19-6
features, 3-1
file channel, 12-1
 advanced properties, 12-15
 directory structure, 12-4
 distribution, 12-2
 file-channel files, 12-4
 directory structure, 12-4
 file-channel files, working with, 12-4
 installation requirements, 12-5
 limitations, 12-6
 local tasks, 12-1
 remote tasks, 12-1, 12-2
 source, 12-6
 target, 12-11
 tasks, 12-1, 12-2
file channel source
 advanced properties, 12-10
 configuration properties, 12-7
file channel source, adding a, 12-8
file channel target
 configuration properties, 12-11
file channel target, adding a, 12-13
file transfer service
 settings, 18-20
File_Channel_settingUpTarget, 12-13
file-channel directory structure, 12-4
file-channel files, working with, 12-4
files, 12-1
filters, 15-7
 expression builder, 15-30
 range builder, 15-11
 SQLite syntax, 15-12
full load, 3-4
 monitoring, 17-1
 monitoring tables, 17-2
full load task settings, 15-48
full load tuning, 15-50
function
 data enrichment, 15-41

G

general settings
 default recipient list, 18-13
 mail parameters, 18-12
general table settings, 15-2
getting started, 2-1
global transformations, 15-13

- apply rule, 15-15
- define the transformation rule, 15-18
- delete transformation rule, 15-29
- edit transformation rule, 15-29
- expression builder, 15-30
- new rule wizard, 15-14
- transformation rule
 - add column, 15-27
 - convert data type, 15-28
 - drop column, 15-28
 - rename column, 15-25
 - rename schema, 15-19
 - rename table, 15-22
- transformation rules list, 15-28
- transformation type, 15-14

H

- header columns, 15-46
- high availability
 - different configurations, 21-5
- history, 17-13
- HTTPS, 20-1

I

- import tasks, 14-9
- input columns (for expressions), 15-34
- installation, 1-4
 - file channel, 12-5
 - Linux, 1-10
 - pre-installation, 1-10
 - procedures, 1-10
 - uninstall, 1-12
 - Oracle, 5-1
 - Windows, 1-4

L

- license settings, 18-14
- licensing, 18-14
 - advanced license request options, 18-15
 - registering, 18-15
 - request a license, 18-14
 - viewing a license, 18-16
- limitations
 - file channel, 12-6
 - Microsoft SQL Server, 6-2
 - MYSQL, 8-4
 - ODBC, 9-1
 - Oracle, 5-2
 - SAP Sybase ASE, 7-2
- Linux
 - required software, 1-2
- Linux installation procedures, 1-10
- Linux pre-installation tasks, 1-10
- Linux uninstall, 1-12
- list of notifications, 18-11
- log file, 16-7
- log file actions (task), 17-14
- log file maintenance operations, 18-18

- log messages (task), 16-6
- log viewer, 16-7
- logging
 - task log settings, 15-56
- logging file settings, 15-56
- logging level settings, 15-56
- logging settings (server), 18-17

M

- mail parameters, 18-12
- manage databases, 14-2, 14-3
- messages, 16-5, 17-12
 - log messages (task), 16-6
 - notifications, 16-5
- Microsoft SQL Server, 6-1
 - advanced properties (source), 6-14
 - advanced properties (target), 6-23
 - backup and recovery, 6-10
 - configuration properties, 6-5, 6-17
 - data types
 - not supported, 6-9
 - limitations, 6-2
 - set up MS Replication, 6-10
 - source data types, 6-7
 - supported versions, 6-1
 - tables without a primary key, replicating, 6-10
 - target data types, 6-19
- Microsoft SQL server
 - alternate backup folder, 6-15, 6-16
- Microsoft SQL Server file, 12-1
- Microsoft SQL Server source
 - add Microsoft SQL Server, 6-12
 - database configuration, 6-9
- Microsoft SQL Server source data types, 6-7
- Microsoft SQL Server target
 - add Microsoft Microsoft SQL Server, 6-21
 - security, 6-19
- Microsoft SQL Server target data types, 6-19
- monitor mode, 13-6
- monitor tools, 17-12
- monitoring, 17-1
 - alerts, 16-5
 - CDC, 17-7
 - applied changes, 17-10
 - incoming changes, 17-9
 - latency, 17-11
 - tables, 17-8
 - throughput, 17-11
- change processing, 17-7
 - applied changes, 17-10
 - incoming changes, 17-9
 - latency, 17-11
 - tables, 17-8
 - throughput, 17-11
- full load
 - all tables, 17-3
 - completed loading, 17-4
 - errors, 17-6
 - loading tables, 17-4

- queued, 17-5
 - tables, 17-2
 - throughput, 17-6
 - total completion, 17-3
- messages, 17-12
- run options, 16-2
- running a task, 16-1
- status tab
 - full load, 17-1
- task status, 16-5
- tools, 17-12
 - history, 17-13
 - task logging, 17-13
- view, 17-1
- monitoring CDC
 - tables, 17-8
- monitoring change processing, 17-7
 - tables, 17-8
- monitoring full load, 17-1
 - tables, 17-2
- multitenant
 - Oracle, 5-15
- MYSQL
 - configuration properties, 8-5
 - limitations, 8-4
- MySQL, 8-1
 - add a MySQL source, 8-9
 - add a MySQL target, 8-15
 - advanced properties (source), 8-11, 8-18
 - configuration properties, 8-13
 - source data types, 8-6
 - target data types, 8-13
- MySQL source data types, 8-6
- MySQL target
 - security, 8-13
- MySQL target data types, 8-13

N

- new task, 14-1
- notification list, 18-11
- notification rule wizard, 18-2
 - notification message page, 18-6
 - notify how/when? page, 18-6
 - notify how/who? page, 18-5
 - Notify When? page (changes of status), 18-3
 - Notify When? page (define errors), 18-5
 - Notify When? page (define when), 18-3
 - summary page, 18-11
- notifications, 16-5, 18-2, 18-7
 - create, 18-2
 - create a notification message, 18-7
 - define, 18-2
 - define changes of status that trigger the notification, 18-3
 - define errors that trigger the notification, 18-5
 - define the notification distribution properties, 18-5
 - define the notification message, 18-6
 - define trigger, 18-3
 - deleting, 18-12

- determine email recipients for notification, 18-6
 - editing, 18-12
 - notification list, 18-11
 - review the notification rule, 18-11
 - set values for latency, disk utilization, memory utilization, 18-4
 - using variables, 18-7
- notifications settings, 18-1

O

- ODBC
 - advanced properties (target), 9-8
 - limitations, 9-1
 - source, 9-1
- ODBC data source, 9-1
- ODBC prerequisites, 9-1
- ODBC source
 - configuration properties, 9-2
 - data types, 9-3
- ODBC source, adding a, 9-6
- Oracle, 5-1
 - add Oracle source, 5-14
 - add Oracle target, 5-25
 - advanced properties (source), 5-16
 - advanced properties (target), 5-27
 - Amazon RDS, 5-12
 - ARCHIVELOG (source), 5-10
 - configuration properties, 5-4
 - data types
 - not supported, 5-10
 - installation, 5-1
 - limitations, 5-2
 - multitenant support, 5-15
 - Oracle source data types, 5-8
 - Oracle target data types, 5-23
 - supplemental logging, 5-11
 - supported versions, 5-1
- Oracle file, 12-1
- Oracle on Amazon RDS
 - working with, 5-12
- Oracle source
 - security requirements, 5-5
- Oracle source data types, 5-8
- Oracle target
 - configuration properties, 5-20
 - database configuration, 5-24
 - security requirements, 5-21
- Oracle target data types, 5-23

P

- PostgreSQL, 11-1
 - add PostgreSQL target, 11-12
 - advanced properties (source), 11-8
 - advanced properties (target), 11-14
 - data types, 11-10
 - source, 11-1
- PostgreSQL source
 - data types, 11-4

- prerequisites, 1-2
- properties
 - file channel source, 12-7
 - file channel target, 12-11
 - Microsoft SQL Server source, 6-5
 - Microsoft SQL Server target, 6-17
 - MYSQL source, 8-5
 - MySQL target, 8-13
 - ODBC target, 9-2
 - Oracle source, 5-4
 - Oracle target, 5-20
 - SAP Sybase ASE source, 7-2
 - SAP Sybase ASE target, 7-8

R

- recovery
 - Microsoft SQL Server, 6-10
- registering a license, 18-15
- remote file channel task
 - adding tables to, 12-4
- replication, 3-2
- replication task, 3-3
- replication topologies, 3-4
- requesting a license, 18-14
 - advance license request options, 18-15
- required software, 1-1, 1-2
- review the notification rule, 18-11
- running a task, 16-1
 - run options, 16-2

S

- SAP Sybase ASE, 7-1
 - add SAP Sybase ASE source, adding a, 7-4
 - advanced properties (source), 7-6
 - advanced properties (target), 7-12
 - configuration properties, 7-2
 - data types
 - not supported, 7-4, 7-9
 - limitations, 7-2
 - SAP Sybase ASE target, adding a, 7-9
 - source data types, 7-3
 - target data types, 7-8
- SAP Sybase ASE source
 - security, 7-3
- SAP Sybase ASE source data types, 7-3
- SAP Sybase ASE target
 - configuration properties, 7-8
 - security, 7-8
- SAP Sybase ASE target data types, 7-8
- scheduler settings (server), 18-23
- security
 - encryption, 20-1
 - https, 20-1
 - Microsoft SQL Server source, 6-6
 - Microsoft SQL Server target, 6-19
 - MySQL target, 8-13
 - SAP Sybase ASE source, 7-3
 - SAP Sybase ASE target, 7-8

- security requirements
 - Oracle source, 5-5
 - Oracle target, 5-21
- selecting tables, 14-6
- server log level, 18-17
- server log settings, 18-17
- server settings
 - error handling, 18-16
- server view, 13-8
- setting up a new task, 14-1
- setting up Microsoft SQL Server for replication, 6-10
- settings, 18-1
 - configure log files, 18-18
 - default recipient list, 18-13
 - license, 18-14
 - log file actions, 18-19
 - logging (server), 18-17
 - mail parameters, 18-12
 - maintenance operationsl, 18-18
 - notifications, 18-2
 - create, 18-2
 - notificationsl, 18-1
 - scheduler (server), 18-23
 - server log level, 18-17
 - view server log, 18-19
- software requirements, 1-1, 1-2
- source database
 - file channel, 12-1
 - Microsoft SQL Server, 6-1
 - MySQL, 8-1
 - ODBC, 9-1
 - Oracle, 5-1
 - PostgreSQL, 11-1
 - SAP Sybase ASE, 7-1
- source databases
 - adding, 14-4
 - files, 12-1
- source endpoint
 - file channel, 12-6
- SQL Server source
 - security, 6-6
- SQLite
 - for transforms, 15-6
- SQLite syntax
 - for filters, 15-12
- supplemental logging, 5-11
- supported browsers, 1-2
- supported databases, 4-1
- supported Linux versions, 1-2
- supported operating systems, 1-1
- supported platforms, 1-1
- supported Windows versions, 1-1
- system architecture, 3-3

T

- table errors, 19-4
- table errors settings, 15-54
- table settings, 15-1
 - filters, 15-7

- general, 15-2
- transforms, 15-3
- target database
 - Amazon Redshift, 10-1
 - file channel, 12-1
 - Microsoft SQL Server, 6-1
 - PostgreSQL, 11-1
- target databases
 - adding, 14-4
 - files, 12-1
 - MySQL, 8-1
 - Oracle, 5-1
 - SAP Sybase ASE, 7-1
- target endpoint
 - file channel, 12-11
- task, 3-3
 - log, 16-7
- task configurations, 15-46
- task log file
 - viewing, 17-14
- task logging, 17-13
 - log file actions, 17-14
 - set logging level, 17-13
 - view log file, 17-14
- task logging level, 17-13
- task logging settings, 15-56
- task settings, 15-46
 - apply changes, 15-50, 15-51
 - apply changes settings sub-tab, 15-50
 - apply changes tuning sub-tab, 15-51
 - apply error sub-tab, 15-55
 - Control Tables tab, 15-47
 - data error sub-tab, 15-53
 - environmental errors sub-tab, 15-53
 - error handling, 15-52, 15-53, 15-54, 15-55
 - error handling settings sub-tab, 15-53
 - full load, 15-48, 15-50
 - full load settings sub-tab, 15-48
 - full load tuning sub-tab, 15-50
 - logging, 15-56
 - logging file sub-tab, 15-56
 - logging level sub-tab, 15-56
 - table error sub-tab, 15-54
 - Target Metadata tab, 15-47
- task status, 16-5
- tasks, 13-2, 14-1, 15-1, 16-1
 - add filters, 15-7
 - add tables, 14-5
 - creating, 14-1
 - customizing, 15-1
 - deleting, 14-9
 - designing, 14-1
 - editing, 14-9
 - filters, 15-7
 - range builder, 15-11
 - SQLite syntax, 15-12
 - global transformations, 15-13
 - importing and exporting, 14-9
 - messages, 16-5
 - new, 14-1

- running a task, 16-1
- selecting tables, 14-6
- setting up, 14-1
- table settings, 15-1, 15-2
- transforms, 15-3
 - creating expressions, 15-6
 - SQLite syntax, 15-6
- tasks view, 13-2
- topologies, 3-4
- Transform Tab
 - unique index, designating on target, 15-4
- transforms, 15-3
 - creating expressions, 15-6
 - expression builder, 15-30
 - SQLite syntax, 15-6
- tutorial, 2-1

U

- UI Console, 13-1
 - connecting to with multiple users, 1-14
 - design mode, 13-4
 - monitor mode, 13-6
 - server view, 13-8
 - tasks view, 13-2
 - viewing tasks, 13-4
- uninstall from Linux, 1-12
- unique index
 - designating on target, 15-5
- using HTTPS, 20-1

V

- variables, 18-7
- variables in messages, 18-7
- view database information, 14-3
- view server log, 18-19
- view task log file, 17-14
- viewing a license, 18-16
- viewing tasks in the UI Console, 13-4

W

- Web console, 13-1
- Web security, 20-1
- Windows
 - required software, 1-1
- working with file-channel files, 12-4