

Getting Started with Attunity Replicate

Version 6.1

May 29, 2018

Copyright © 2018 Attunity Ltd. All rights reserved.

Primary Author: Charlton Book

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.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, endpoints, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Attunity license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software—Restricted Rights (June 1987). Attunity Ltd., 70 Blanchard Road, Burlington, MA 01803

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Attunity is a registered trademark of Attunity Ltd and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Attunity is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Attunity is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Attunity is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Introduction

This tutorial will walk you through the steps required to successfully set up and run an Oracle to Microsoft SQL Server replication task in Attunity Replicate. Although the instructions in this tutorial refer specifically to Oracle and Microsoft SQL Server, the same steps apply when replicating between any of the databases supported by the Attunity Replicate Express edition. If you wish to use a different source and/or target, simply refer to the section that describes that database in the online Help, which can be accessed by clicking the Help link in the top right corner of the Attunity Replicate Console.

Note The term "endpoint" is used generically throughout this tutorial to refer to a data repository that can be used as a source or target in an Attunity Replicate task.

In this tutorial:

Step 1: What You Need	5
Step 2: Install Attunity Replicate	6
Step 3: Open the Attunity Replicate Console	13
Step 4: Define an Oracle Endpoint as the Replication Source	14
Step 5: Define a Microsoft SQL Server Endpoint as the Replication Target	16
Step 6: Define a Replication Task	18
Step 7: Run and Monitor the Replication Task	23
Step 8 (Optional): View the Replicated Tables in Microsoft SQL Server	24

Limitations of the Attunity Replicate Express Edition

Attunity Replicate Express Edition is intended for evaluation purposes only. You can upgrade to Attunity Replicate Full Edition at any time during the evaluation period by purchasing a production license from Attunity.

The limitations of the Attunity Replicate Express Edition are described below.

Design Limitations

Note: You must be an Attunity Replicate Administrator to use the Express Edition.

Replication from the following sources is supported:

- » IBM DB2 for LUW
- » Oracle
- » Microsoft SQL Server
- » MySQL
- » IBM Informix
- » SAP Sybase ASE

Replication to the following target is supported:

- » Amazon EMR (Preview)
- » Hadoop
- » Microsoft Azure HDInsight
- » Microsoft SQL Server
- » Oracle

The following task replication profiles are *not* available:

- » Bidirectional Replication

The following server functionality is *not* available:

- » Scheduler
- » FTS (File Transfer Service)

Runtime Limitations

The following runtime limitations apply:

- » Only one task can run at a time.
- » A maximum of 250 thousand rows per day can be replicated from source to target.
- » The license is valid for one year from the build date.

Step 1: What You Need

For this tutorial, you will need the following:

- » Attunity Replicate installed on a computer in your network
For additional installation information, see [Step 2: Install Attunity Replicate](#).
- » For the Oracle source:
 - » Access to the HR schema tables that are part of the Oracle database installation

Note If these tables are not available, contact your Oracle database administrator.

- » `system/<password>` for an admin user
- » For the Microsoft SQL Server target:
A 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 the Attunity Replicate Console, one of the following Internet browsers:
 - » Microsoft Internet Explorer version 11 and above
 - » Mozilla Firefox
 - » Google Chrome

Step 2: Install Attunity Replicate

This section explains how to install Attunity Replicate.

In this chapter:

[Recommended Hardware Configuration](#)

[Installing Attunity Replicate on Windows](#)

[Installing Attunity Replicate on Linux](#)

Recommended Hardware Configuration

This section describes the recommended hardware configurations for using Attunity Replicate.

The following table describes the recommended hardware configuration for installing Attunity Replicate on Windows and Linux operating systems. Note that the recommendations apply to mid-scale systems (i.e. hundreds of tasks) rather than large-scale systems (i.e. thousands of tasks).

Table 0.1 | 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	Additional cores are useful in any of the following situations: <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	More memory is useful in any of the following situations: <ul style="list-style-type: none">» Many tasks running in parallel» Long-running transactions on the source endpoint (for example, monthly batch processing)» Many active users on the source system
Disk requirements	320 GB 7200	500 GB 10,000	500 GB 15,000	A faster disk is useful in any of the following situations:

Table 0.1 | Recommended Hardware Configuration (Cont.)

	Basic System	Large System	Extra-Large System	Notes:
	RPM	RPM RAID	RPM RAID	<ul style="list-style-type: none">» Using a file-based target, such as Greenplum or Actian Vector» Long-running source transactions that may not fit into memory» Using tasks that are set up to continue processing during target outage A larger disk is required in any of the following situations: <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 RAID is recommended for system recoverability in case of disk failure for all configurations.
Network	1 Gb	10 Gb	Two 10 Gb	

Installing Attunity Replicate on Windows

This section describes the prerequisites and setup instructions for installing Attunity Replicate on Windows.

Windows Software Requirements

To install the Attunity Replicate Server and Console on a Windows computer, you must have the following installed on your system:

- » .NET Framework 4.5.2 or above
- » Visual C++ Redistributable for Visual Studio 2015. If it is not installed or if an older version is installed, it will be installed automatically during installation.
- » TLS v1.2 needs to be fully installed and configured prior to installing Replicate on a Windows 2016 Server.

Windows Permissions

Attunity Replicate needs to be installed as an Administrator.

The following privileges are required to start the Attunity Replicate **UI Server** service (which is run as a local system service), but are dropped as soon as the service is started:

- » SE_CREATE_GLOBAL_NAME
- » SE_CREATE_PAGEFILE_NAME
- » SE_CREATE_PERMANENT_NAME
- » SE_CREATE_SYMBOLIC_LINK_NAME
- » SE_CREATE_TOKEN_NAME
- » SE_DEBUG_NAME
- » SE_ENABLE_DELEGATION_NAME
- » SE_IMPERSONATE_NAME
- » SE_INC_BASE_PRIORITY_NAME
- » SE_INCREASE_QUOTA_NAME
- » SE_INC_WORKING_SET_NAME
- » SE_LOAD_DRIVER_NAME
- » SE_LOCK_MEMORY_NAME
- » SE_MACHINE_ACCOUNT_NAME
- » SE_MANAGE_VOLUME_NAME
- » SE_PROF_SINGLE_PROCESS_NAME
- » SE_RELABEL_NAME
- » SE_REMOTE_SHUTDOWN_NAME
- » SE_RESTORE_NAME
- » SE_SECURITY_NAME
- » SE_SHUTDOWN_NAME
- » SE_SYNC_AGENT_NAME
- » SE_SYSTEM_ENVIRONMENT_NAME
- » SE_SYSTEM_PROFILE_NAME
- » SE_SYSTEMTIME_NAME
- » SE_TAKE_OWNERSHIP_NAME
- » SE_TCB_NAME
- » SE_TIME_ZONE_NAME
- » SE_TRUSTED_CREDMAN_ACCESS_NAME
- » SE_UNDOCK_NAME

In addition, the account that runs Replicate needs to be granted access to the **Data** directory (~\Attunity\Replicate\Data) as well as any directory containing files (such as CSV files) that need to be used in a replication task.

The Installation Procedure

To install Attunity Replicate on Windows, run the Attunity Replicate Express setup wizard and follow the embedded instructions.

Note In the setup wizard's **Replication Server Location** screen, one of the options is **Connect to a remote Linux Attunity Replicate Server**. You should only select this option if you have already installed Attunity Replicate Server on a Linux machine. If you select this option, you will be prompted for the IP address and port number of the Linux machine in the following screen.

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

Installing Attunity Replicate on Linux

This section describes the prerequisites and setup instructions for installing Attunity Replicate on Linux.

Linux Software Requirements

Attunity Replicate can be installed on the following Linux platforms:

- » Red Hat Enterprise Linux 6.2, 6.3, 6.4, 6.5, 6.6, 6.7, 6.8, 7.0, 7.1, 7.2, and 7.3 (64-bit)
- » SUSE Linux 11.4 or 12.0 (64-bit)

The Installation Procedure

Before you carry out the following procedures, copy the Attunity Replicate RPM file to any location on the Linux computer.

The default installation directory for the Attunity Replicate is `/opt/attunity/replicate`. You can choose to install to a different directory as described below.

Note When installing or upgrading Replicate on SUSE Linux, the `--nodeps` parameter must be added to the command, as in the following example:

```
rpm -ivh --nodeps
```

To install Attunity Replicate:

Issue the following single command:

```
[user=user] [group=group] [verbose=true|debug=true] [nocredentials=true]  
[data=replicate_data_directory] rpm -ivh [--prefix dirname] areplicate-  
<version-build>.x86_64.rpm [--nodeps]
```

Example:

```
user=mike group=sales verbose=true rpm -ivh --prefix /opt/mydir/  
areplicate-6.1.0-102.x86_64.rpm
```

The optional command parameters are described in the [Optional Command Parameters](#) table below.

Table 0.2 | Optional Command Parameters

Parameter	Description
<code>[user=user]</code>	<p>The default user under which Attunity Replicate is installed is "attunity". You can choose to install the product under a different user by prefixing <code>user=user</code> to the command.</p> <p>See also the <code>[nocredentials=true]</code> parameter below.</p>
<code>[pass=password]</code>	<p>Sets the server password.</p> <p>Users will be prompted for this password when connecting to the Replicate server via the Replicate Web Console.</p> <p>You can either set this password when installing Replicate (recommended) or you can set it later by running the following command (from the Replicate "bin" directory):</p> <pre>./repctl -d data_dir setserverpassword password</pre> <p>where <code>data_dir</code> is the location of the Replicate "data" directory and <code>password</code> is the password for accessing the server.</p>
<code>[group=group]</code>	<p>The default group under which Attunity Replicate is installed is attunity. You can choose to install the product under a different group by prefixing <code>group=group</code> to the command.</p> <p>See also the <code>[nocredentials=true]</code> parameter below.</p>
<code>[nocredentials=true]</code>	<p>If the specified user or group is not defined locally (e.g. in Active Directory), you must include this parameter in the command. Otherwise, the installation will fail.</p>
<code>[data=replicate_data_directory]</code>	<p>The default Attunity Replicate "data" directory is:</p> <pre><product_dir>/data</pre> <p>Use this parameter to install the Replicate "data" directory in a non-default location.</p>
<code>[--prefix dirname]</code>	<p>Prefixes the attunity/replicate application directory with the path specified by <code>dirname</code>. So, for example, if <code>--prefix opt/mydir</code>, then Replicate will be installed to <code>opt/mydir/attunity/replicate</code>.</p> <p>Only required when installing Attunity Replicate to a non-default path.</p>
<code>[verbose=true debug=true]</code>	<p>Specify <code>verbose=true</code> for more information during the installation or <code>debug=true</code> for detailed debug messages</p>

Table 0.2 | Optional Command Parameters (Cont.)

Parameter	Description
	during the installation.
<code>--nodeps</code>	This parameter is only required when installing Replicate on Linux SUSE.

Setup will perform the following actions:

- » Create a new user and group named “attunity” (unless you chose to use a different user and group and/or a user named “attunity” already exists).
- » Change the Attunity Replicate installation directory owner to the “attunity” user and group or to your preferred user and group.
- » Install the application files.
- » Start the service (areplicate).

Note The Attunity Replicate/`lib` folder *must* precede the `/usr/lib64` folder in the `LD_LIBRARY_PATH` environment variable.

The environment variable is set in the `<product_dir>/bin/arep_login.sh` file.

Example:

```
LD_LIBRARY_PATH=/opt/attunity/replicate/lib:/usr/lib64
```

Step 3: Open the Attunity Replicate Console

From the Windows Start menu, select **All Programs > Attunity Replicate > Attunity Replicate Console**.

Note You can access Attunity Replicate from any computer in your system.

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

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

where *<computer name>* is the name or IP address of the computer (including the Windows domain name) on which the Attunity Replicate Server is installed.

Step 4: Define an Oracle Endpoint as the Replication Source

This section guides you through the steps required to add an Oracle endpoint as the source database. This is the database from where you want to replicate data.

To add an Oracle source database:

1. In Task view, click **Manage Endpoint Connections**.

The **Manage Endpoint Connections** dialog box opens.

2. Click **New Endpoint Connection**.

3. Provide the following information:

- » **Name:** Type **OracleSource**.

- » **Description:** Optionally, enter a description or leave blank.

- » **Role:** Select **Source**.

- » **Type:** Select **Oracle**.

- » **Connection string:** Enter the connect string to the Oracle database you work with, in any Oracle format.

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

```
tomato:1521/orcl
```

- » **User Name:** Enter the user name for the Oracle database you wish to access.

The default user name is **SYSTEM**.

- » **Password:** Enter the password for the Oracle database you wish to access.

The default password is **manager**.

The figure below shows the data that you need to enter in the **General** tab of the Oracle endpoint.

General

Advanced

Name:

OracleSource

Description:

Role:

Source

Type:

Oracle

Connection string:

tomato:1521/orcl

User name:

SYSTEM

Password:

....

Test Connection

Save

Discard Changes

Close

4. Click **Test Connection** to verify the information you entered and the availability of the database.
5. Click **Save** to save the endpoint settings.

Step 5: Define a Microsoft SQL Server Endpoint as the Replication Target

This section guides you through the steps required to add a Microsoft SQL Server endpoint as the target database. This is the database to where you want to replicate data.

To add a Microsoft SQL Server target endpoint:

1. In Tasks view, click **Manage Endpoint Connections**.

The Manage Endpoint Connections dialog box opens.

2. Click **New Endpoint Connection**.

3. Provide the following information:

- » **Name:** Type `sqlserver_target`.

- » **Description:** Optionally, enter a description or leave blank.

- » **Role:** Select **Target**.

- » **Server name:** Enter the name of the computer where your Microsoft SQL Server database is installed.

For example, if you connect to a Microsoft SQL Server database on a computer called `bee`, enter **bee**.

- » Select one of the following:

Windows authentication if your Microsoft SQL Server database is configured to accept Windows authentication.

SQL Server authentication if your Microsoft SQL Server database is not configured to accept Windows authentication. In this case, you also need to provide a valid user name and password.

- » **Database name:** Enter **tempdb**, which is the name of the database to where you are going to replicate data. If you created a new database for this purpose, enter the name of that database.

General

Advanced

Name:

sqlserver_target

Description:

Role:

☐ Source ☒ Target

Type:

Microsoft SQL Server

Server name:

bee01-win7

☒ Windows authentication

☐ SQL Server authentication

User name:

Password:

Database name:

tempdb

Browse...

Test Connection

Save

Discard Changes

Close

4. Click **Test Connection** to verify the information you entered and the availability of the database.
5. Click **Save** to save the endpoint settings.

Step 6: Define a Replication Task

This section guides you through the steps required to define a replication task that copies the data from the HR.EMPLOYEES and HR.JOBS tables.

Adding a replication task consists of the following sub-tasks:


- » [Add a Replication Task to the Attunity Replicate Console](#)
- » [Add the Source and Target Endpoints to the Task](#)
- » [Select Tables for the Replication Task](#)

Add a Replication Task to the Attunity Replicate Console

This task guides you through adding a replication task to the Attunity Replicate Console.

To add a replication task:

1. Make sure **Tasks** is selected in the upper left corner of the Attunity Replicate Console.



2. Click **New Task** to open the **New Task** dialog box.
3. In the **Name** field, type **My_Task**.

New Task...

×

Name

i

My_Task

Description

Replication Profile

Select the profile that best describes your replication use case

☒ Unidirectional


☐ Bidirectional

i


Bidirectional replication is not available in Replicate Express

Task Options


Select the replication options for this task



Full Load



Apply Changes



Store Changes

Perform a full load of target tables, then keep them up-to-date by applying changes.

OK

Cancel

- The **Apply Changes** replication option is currently enabled. Disable it by clicking the **Apply Changes** button at the bottom of the dialog box.

Note When this option is enabled, Replicate detects any changes to the source tables and automatically applies them to the target tables. Using the **Apply Changes** task option has special requirements such as granting additional permissions on the source database. Therefore, to keep things simple, in this tutorial we will be setting up a **Full Load** only task. For information on setting up a task with Apply Changes, please refer to the *Attunity Replicate Setup and User Guide*.

5. Click **OK** to create the task.

The task opens in new tab. You are now ready to [add your source and target endpoints to the task definition](#).

Add the Source and Target Endpoints to the Task

This section guides you through adding the source and target endpoints to the replication task, which is a simple drag-and-drop operation. In the **Endpoints** tab, the following icons help you distinguish between source and target endpoints:



Source endpoint, which is represented by a database, file, or NoSQL icon, depending on the endpoint type, with an orange arrow pointing away from the source (a database in this example).

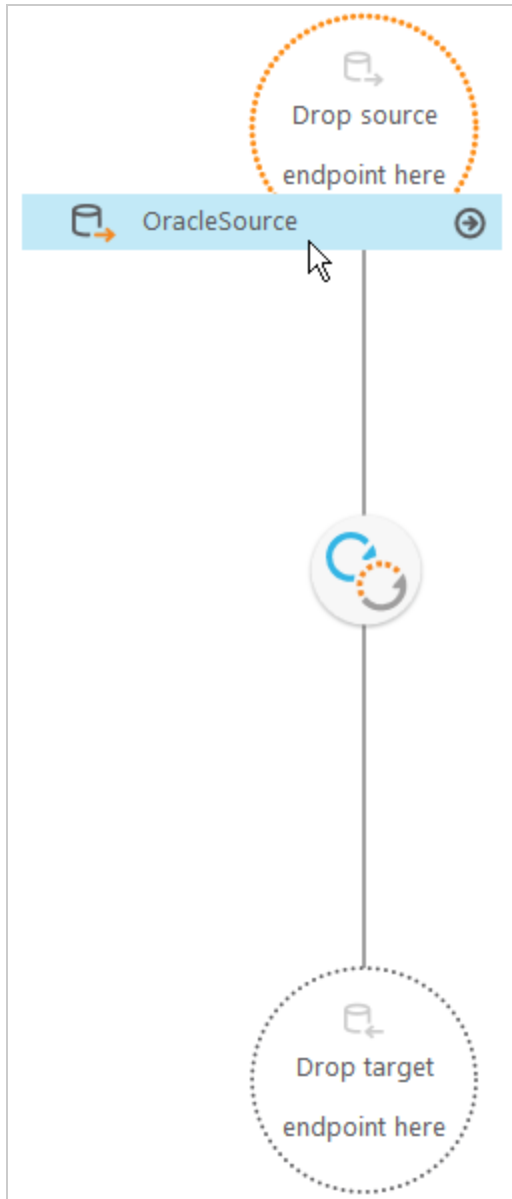


Target endpoint, which is represented by a database, file or NoSQL icon, depending on the endpoint type, with a blue arrow pointing toward the target (a database in this example).

The **Endpoints** pane consists of **All**, **Source**, and **Target** tabs. Because each database is named in a way that reflects whether it is a source or a target, click the **All** tab.

To add the source or target endpoints to the task:

1. In the **Endpoints** pane, click the **All** tab.
2. Drag the `OracleSource` database to the **Drop source endpoint here** area in the endpoints diagram.



3. Drag the `sqlserver_target` database to the **Drop target endpoint here** area.

Select Tables for the Replication Task

After adding the source and target databases, you now need to select which Oracle source tables you want to replicate to the Microsoft SQL Server target.

This task guides you through selecting specific tables (HR.EMPLOYEES and HR.JOBS) from the Oracle source. Replicate takes all of the data from these tables "as is" and copies it to the Microsoft SQL Server target.

To add tables to the replication task:

1. In the right pane of the Attunity Replicate Console, click **Table Selection**.
The **Select Tables** dialog box opens.
2. In the Select Tables dialog box, do the following:
 - » From the **Schema** list, select **HR**, and then click **Search**.
 - » From the **Table List**, select **EMPLOYEES**, and then click the right arrow to select that table.
 - » Repeat these steps for the **JOBS** table.
 - » Click **OK**.

Select Tables

Source Tables

Schema:

Table:

Table List

HR	COUNTRIES
HR	DEPARTMENTS
HR	EMPLOYEES
HR	JOBS
HR	JOB_HISTORY
HR	LOCATIONS
HR	QBCDC
HR	REGIONS

Tables: 8

Table Selection Patterns

Selected Tables

HR	EMPLOYEES
HR	JOBS

Tables: 2

3. On the task tab, click **Save**. The task is now **ready to run**.

Step 7: Run and Monitor the Replication Task

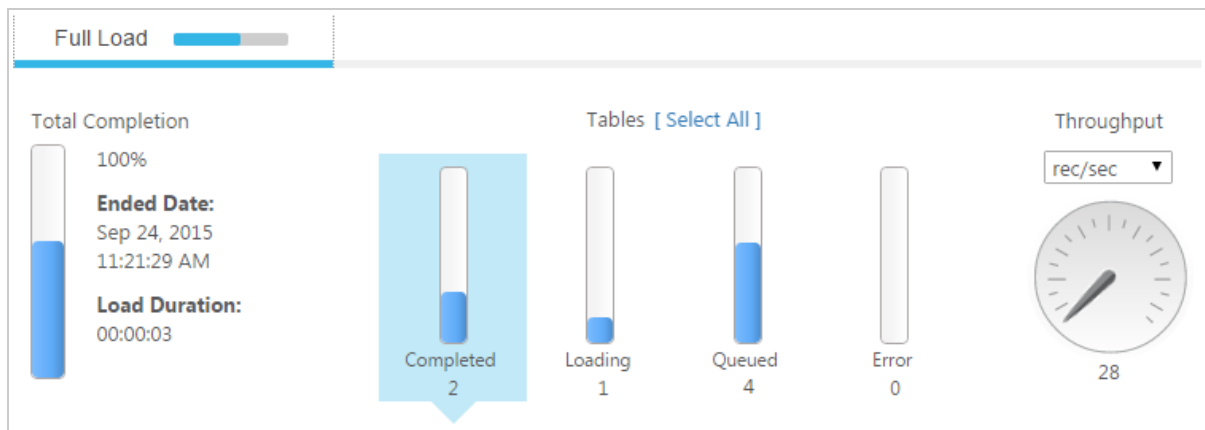
You can now run the replication task and see the results of the replication in real time. This section guides you through the steps required to run a Full Load replication task and view the progress in the Monitor.

To run and monitor the replication task:

1. On the task tab, click **Run**.

The **Starting task** message displays, and the console switches to Monitor view, which includes gauges and graphs on two tabs:

- » **Full Load** tab: Indicates the status progress during the Full Load process.
- » **Change Processing** tab: Monitors changes that occur after the Full Load completes. As we defined a Full Load only replication task, this tab is not relevant for this tutorial.



2. Click the **Select All** link above the Tables graphs. Replicate displays a table below the graphs with information about each of the tables being processed in the task.

Tables - All					
Filter By		Resume		Reload	
Table Name	Status	Estimated Count	Count	Elapsed Time	Progress
HR.EMPLOYEES	Completed	107	107	00:00:02	100%
HR.JOBS	Completed	19	19	00:00:02	100%

3. Click the individual bar graphs, such as the **Completed** graph and the **Loading** graph, to view additional information.

Step 8 (Optional): View the Replicated Tables in Microsoft SQL Server

After the replication completes, you can view the `tempdb` database in Microsoft SQL Server. You will see that this database now includes two new tables: `HR.EMPLOYEES` and `HR.JOBS`.

To view the replicated tables in Microsoft SQL Server:

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