



**Quick Start Guide for SAP HANA®
Database**

**SAP Replication Server® 15.7.1
SP200**

Linux, UNIX, and Microsoft Windows

DOCUMENT ID: DC01974-01-1571200-01

LAST REVISED: March 2014

Copyright © 2014 by SAP AG or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries. Please see <http://www.sap.com/corporate-en/legal/copyright/index.epx#trademark> for additional trademark information and notices.

Contents

Conventions	1
Before You Begin	5
Additional Replication Resources	5
Server Entries to the Interfaces Files	9
Adding Server Entries to the Interfaces Files	9
Creating the Sybase User Account	10
ASE-to-SAP HANA Database Replication Setup	11
System Architecture	11
Prerequisites	11
Planning Your Installation	12
Installation Checklist	12
System Requirements	12
Identify the Installation Directory	12
Downloading Products from SMP	13
Installing Replication Components	14
Installing Replication Server	14
Configuring Replication Components	18
Configuration Checklist	18
Configuring Replication Server	18
Testing Replication	23
Resetting the Primary ASE Database For Replication	26
Oracle-to-SAP HANA Database Replication Setup	29
System Architecture	29
Prerequisites	30
Planning Your Installation	30
Installation Checklist	30
System Requirements	31
Identify the Installation Directory	31
Downloading Products from SMP	31
Installing Replication Components	32

Installing Replication Server	33
Installing Replication Agent	36
Configuring Replication Components	38
Configuration Checklist	39
Configuring Replication Server	40
Configuring Replication Agent	43
Testing Replication	59
Resetting the Primary Oracle Database for Replication	61
Microsoft SQL Server-to-SAP HANA Database	
Replication Setup	63
System Architecture	63
Prerequisites	63
Planning Your Installation	64
Installation Checklist	64
System Requirements	65
Identify the Installation Directory	65
Downloading Products from SMP	65
Installing Replication Components	66
Installing Replication Server	67
Installing Replication Agent	70
Configuring Replication Components	72
Configuration Checklist	72
Configuring Replication Server	73
Configuring Replication Agent	76
Testing Replication	89
Resetting the Primary Microsoft SQL Server Database for Replication	91
DB2 UDB-to-SAP HANA Database Replication Setup	93
System Architecture	93
Prerequisites	93
Planning Your Installation	94
Installation Checklist	94
System Requirements	94
Identify the Installation Directory	95

Downloading Products from SMP	95
Installing Replication Components	96
Installing Replication Server	96
Installing Replication Agent	99
Configuring Replication Components	102
Configuration Checklist	102
Setting DB2 UDB Environment Variables	103
Node and Database Cataloging in a DB2 UDB Client	104
Configuring Replication Server	105
Configuring Replication Agent	108
Testing Replication	118
Resetting the Primary DB2 UDB Database for Replication	120
SAP Business Suite-to-SAP HANA Database	
Replication Setup	123
System Architecture	123
Prerequisites	124
Planning Your Installation	124
Downloading Products from SMP	125
Installing Replication Components	125
Configuring Replication Components	125
Configuring Replication Server	126
Configuring Replication Agent	131
SAP HANA-to-SAP HANA Database Replication Setup .	133
System Architecture	133
Installing Replication Agent for SAP HANA	135
User Management	137
Best Practices and Troubleshooting	141
Example Scenarios: Marking Tables for Replication ...	141
Troubleshooting Subscription Problems	142
Glossary	145
Index	153

Contents

Conventions

These style and syntax conventions are used in SAP® documentation.

Style conventions

Key	Definition
monospaced (fixed-width)	<ul style="list-style-type: none"> • SQL and program code • Commands to be entered exactly as shown • File names • Directory names
<i>italic monospaced</i>	In SQL or program code snippets, placeholders for user-specified values (see example below).
<i>italic</i>	<ul style="list-style-type: none"> • File and variable names • Cross-references to other topics or documents • In text, placeholders for user-specified values (see example below) • Glossary terms in text
bold sans serif	<ul style="list-style-type: none"> • Command, function, stored procedure, utility, class, and method names • Glossary entries (in the Glossary) • Menu option paths • In numbered task or procedure steps, user-interface (UI) elements that you click, such as buttons, check boxes, icons, and so on

If necessary, an explanation for a placeholder (system- or setup-specific values) follows in text. For example:

Run:

```
installation directory\start.bat
```

where *installation directory* is where the application is installed.

Syntax conventions

Key	Definition
{ }	Curly braces indicate that you must choose at least one of the enclosed options. Do not type the braces when you enter the command.
[]	Brackets mean that choosing one or more of the enclosed options is optional. Do not type the brackets when you enter the command.
()	Parentheses are to be typed as part of the command.
	The vertical bar means you can select only one of the options shown.
,	The comma means you can choose as many of the options shown as you like, separating your choices with commas that you type as part of the command.
...	An ellipsis (three dots) means you may repeat the last unit as many times as you need. Do not include ellipses in the command.

Case-sensitivity

- All command syntax and command examples are shown in lowercase. However, replication command names are not case-sensitive. For example, **RA_CONFIG**, **Ra_Config**, and **ra_config** are equivalent.
- Names of configuration parameters are case-sensitive. For example, **Scan_Sleep_Max** is not the same as **scan_sleep_max**, and the former would be interpreted as an invalid parameter name.
- Database object names are not case-sensitive in replication commands. However, to use a mixed-case object name in a replication command (to match a mixed-case object name in the primary database), delimit the object name with double quote characters. For example: **pdb_get_tables "TableName"**
- Identifiers and character data may be case-sensitive, depending on the sort order that is in effect.
 - If you are using a case-sensitive sort order, such as “binary,” you must enter identifiers and character data with the correct combination of uppercase and lowercase letters.
 - If you are using a sort order that is not case-sensitive, such as “nocase,” you can enter identifiers and character data with any combination of uppercase or lowercase letters.

Terminology

Replication Agent™ is a generic term used to describe the Replication Agents for SAP® Adaptive Server® Enterprise (SAP® ASE), Oracle, Microsoft SQL Server, and IBM DB2 for Linux, Unix and Windows. The specific names are:

- RepAgent – Replication Agent thread for Adaptive Server Enterprise
- Replication Agent for Oracle

- Replication Agent for Microsoft SQL Server
- Replication Agent for IBM DB2 UDB

Conventions

Before You Begin

Review these topics before you install and configure the replication system components.

Additional Replication Resources

Set up and begin using a basic replication system in which the SAP HANA® database is the replicate.

If your primary and replicate databases are not already participating in replication using SAP® Replication Server®, use SAP Replication Management Agent Designer to setup the replication environment. See *Replication Management Agent Configuration and Users Guide*.

You can find additional information about the SAP HANA database, including HANA Studio at http://help.sap.com/hana_platform.

While setting up the basic replication system, you may need to refer to additional documents, which are a part of the SAP® Replication Server® documentation set available at <http://infocenter.sybase.com>.

Table 1. Additional References

For Information About	Topic	See
Licensing	Download replication components from the SAP Service Marketplace (SMP).	<i>Downloading Products from SMP</i> on page 13
	Obtain a separate license for each SAP Replication Server and SAP® Replication Server® Options component.	<ul style="list-style-type: none"> • <i>Replication Server 15.7.1 SP200 > Release Bulletin Replication Server 15.7.1 SP200 for UNIX and Linux > Product Summary > Separately Licensed Products</i> • <i>Replication Server 15.7.1 SP200 > Release Bulletin Replication Server 15.7.1 SP200 for Windows > Product Summary > Separately Licensed Products</i>

For Information About	Topic	See
	Obtain a separate license for each Replication Server Options component.	<i>Replication Server 15.7.1 SP200 Release Bulletin for Windows > Product Summary > Separately Licensed Products > Replication Server Options</i>
Installing components from other sources	Install ODBC libraries for ExpressConnect for HANA database.	<i>Replication Server 15.7.1 SP200 > Release Bulletin Replication Server 15.7.1 SP200 for UNIX and Linux > Installation and Upgrade > Special Installation Instructions > Installing ODBC Libraries for ExpressConnect for HANA Database</i>
Connectivity to other servers	Learn about the interfaces file.	<i>Replication Server 15.7.1 SP200 > Administration Guide Volume 1 > Replication Server Technical Overview > Establish Replication Server Connections > Interfaces File</i>
	Learn how to edit an interfaces file.	<i>Editing the Interfaces File</i> on page 17
Roles, logins, and permissions	Learn about user roles and Replication Server tasks.	<i>Replication Server 15.7.1 SP200 > Administration Guide: Volume 1 15.7.1 SP200 > Introduction to Replication Server > Replication Server Roles and Responsibilities</i>
	Replication Server, Replication Agent, and database users.	<i>User Management</i> on page 137
Subscriptions	Subscriptions example.	<i>Replication Server 15.7.1 SP200 > Design Guide 15.7.1 SP200 > Implementation Strategies > Corporate Rollup > Subscriptions</i>

For Information About	Topic	See
	Learn about the materialization process and the types of subscription materialization.	<i>Replication Server 15.7.1 SP200 > Heterogeneous Replication Guide 15.7.1 SP200 > Replication System Overview > Heterogeneous Replication Limitations > Subscription Materialization</i>
	Troubleshoot problems that arise if a subscription is not in a valid state or if the check subscription command returns an error message.	<i>Troubleshooting Subscription Problems</i> on page 142
Configuration parameter reference	Replication Agent configuration parameter list.	<i>Replication Agent 15.7.1 SP200 Reference Manual for Linux, UNIX, and Microsoft Windows > Configuration Parameters > Configuration Parameter Reference</i>
Truncation log management	A RepAgent retrieves transactions from the Adaptive Server primary database log and sends them to the Replication Server.	<i>Replication Server 15.7.1 SP200 > Troubleshooting Guide 15.7.1 SP200 > Adaptive Server Log Problems</i>
	Replication Agent provides features for automatic and manual log truncation.	<i>Replication Agent 15.7.1 SP200 Administration Guide for Linux, UNIX, and Microsoft Windows > Administration > Replication Agent Management > Transaction Log Truncation</i>
Status monitoring	Learn how to verify Replication Server thread status.	<i>Replication Server 15.7.1 SP200 > Administration Guide: Volume 2 15.7.1 SP200 > Performance Tuning > Multiprocessor Platforms > Commands to Monitor Thread Status</i>

For Information About	Topic	See
IBM DB2-specific considerations	Review the general issues and considerations that are specific to using Replication Agent for IBM DB2 UDB.	<i>Replication Agent 15.7.1 SP200 Primary Database Guide for Linux, UNIX, and Microsoft Windows > Replication Agent for UDB > IBM DB2-Specific Considerations</i>
Troubleshooting	Review the procedures you can use to isolate replication failures.	<i>Replication Server 15.7.1 SP200 > Troubleshooting Guide 15.7.1 SP200 > Troubleshooting Overview > Troubleshooting Replication Failures</i>
	Troubleshoot replication failures if Replication Server is down.	<i>Replication Server 15.7.1 SP200 Troubleshooting Guide > Troubleshooting Overview > Troubleshooting Replication Failures > Replication Server Is Down</i>
	Troubleshoot errors that can occur when the Replication Agent starts up.	<i>Replication Server 15.7.1 SP200 > Troubleshooting Guide 15.7.1 SP200 > RepAgent Problems > Problems when Starting the RepAgent</i>
	Troubleshoot Replication Agent failure.	<i>Replication Agent 15.7.1 SP200 Administration Guide for Linux, UNIX, and Microsoft Windows > Troubleshooting > Replication Agent Failure</i>
	Review the basic troubleshooting procedures for Replication Agent and the replication system.	<i>Replication Agent 15.7.1 SP200 Administration Guide for Linux, UNIX, and Microsoft Windows > Troubleshooting</i>

Server Entries to the Interfaces Files

An interfaces file contains network definitions for each Replication Server and data server in a replication system.

The interfaces file is located in `$$SYBASE/interfaces` on UNIX platforms, or `$$SYBASE%\ini\sql.ini` on Microsoft Windows platforms.

When you install Replication Server in an existing installation directory, the installer adds Replication Server information to the existing interfaces file. When you install Replication Server in a unique installation directory, the installer creates a new interfaces file. In this situation, there are two interfaces files—one for your existing SAP applications and one for Replication Server.

The interfaces file requires entries for these components:

- Replication Server
- RSSD SAP® Adaptive Server® Enterprise (SAP® ASE) or ERSSD SAP® SQL Anywhere® for this Replication Server
- SAP HANA database server

If you use direct load materialization, the interfaces file requires an entry for the Replication Agent location.

For each server, you need:

- Server name
- Host name
- Port number

Adding Server Entries to the Interfaces Files

To add entries for the primary and replicate database server in the SAP Replication Server interfaces file, use **dsedit**.

1. Navigate to `$$SYBASE%\OCS-15_0\bin` (Windows), or `$$SYBASE/OCS-15_0/bin` (UNIX).
2. On Windows, double-click **dsedit.exe**.
On UNIX, enter:


```
dsedit
```
3. Select **InterfacesDriver** and click **OK**.
4. Select **Server Object > Add**.
5. Enter the name of the primary database server. Click **OK**.
6. Enter the host name and port number of the primary database server. For example:

Before You Begin

`chaucer, 35356`

7. Click **OK**.
8. Select **Server Object > Add**.
9. Enter the name of the replicate database server. Click **OK**.
10. Enter the host name and port number of the replicate database server. For example:
`johnson, 9756`
11. Click **OK** twice, then select **File > Exit**.

Creating the Sybase User Account

Create a Sybase® user account to ensure that Sybase product files and directories are created with consistent ownership and privileges.

You must create a Sybase user account before you install the replication system components. A user, such as the Sybase system administrator—who has read, write, and execute privileges—must perform all installation and configuration tasks. For Windows, the user must have the Administrator privilege.

1. To create a Sybase system administrator account, choose an existing account, or create a new account and assign a user ID, group ID, and password for it.
This account is sometimes called the “sybase” user account. See your operating system documentation for instructions on creating a new user account.
If you have already installed other Sybase software, the “sybase” user already exists.
2. Verify that you can log in to the machine using this account.

ASE-to-SAP HANA Database Replication Setup

Implement a sample replication system for replication from Adaptive Server Enterprise (ASE) to SAP HANA database.

Use the tasks in this document to set up a proof-of-concept or test replication system.


The tasks use the sample Replication Server instance SAMPLE_RS. If you have already installed Replication Server, skip to *Configuring Replication Components* on page 18.

For the most current information about Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

Multiple components are required to implement an ASE-to-SAP HANA database replication system: a primary Adaptive Server data server, a replicate SAP HANA database, and a Replication Server with ExpressConnect for HANA DB.

Figure 1: ASE-to-SAP HANA Database Replication Setup



Prerequisites

Review the prerequisites before you begin setting up replication for Adaptive Server.

- Familiarity with Adaptive Server and SAP HANA database.
- A valid license for Replication Server with ExpressConnect for HANA DB.
- Install and configure Adaptive Server to serve as the primary (source) database from which Replication Agent replicates transactions, and SAP HANA database as the replicate (target) database.
- TCP/IP connectivity between all hosts.

- Adjust Solaris directories, files, executable commands, and examples accordingly for Microsoft Windows and other UNIX platforms.

Do not use Replication Server reserved words for object names and connection names. The *Replication Server Reference Manual* includes a complete list of reserved words.

Planning Your Installation

Before installing the replication components, verify the system requirements, identify the installation directories, and obtain the installation software.

Installation Checklist

Use the checklist to navigate through the installation when you set up your replication environment.

Table 2. Installation Checklist

	Action	Topic
1	Make sure the target computer meets the minimum memory and disk space requirements.	<i>System Requirements</i>
2	Determine the directory in which to install the replication components.	<i>Identify the Installation Directory</i>
3	Download replication components from the SAP Service Marketplace (SMP).	<i>Downloading Products from SMP</i>
4	Install Replication Server.	<i>Installing Replication Server</i>

System Requirements

Make sure the target computer on which you are installing the replication components meets the minimum memory and disk space requirements.

See the *Replication Server Installation Guide > Planning Your Installation > System Requirements*.

To improve performance by avoiding multiple network hops, install Replication Server and the SAP HANA database on the same machine.

Note: The only versions of Linux that ExpressConnect for HANA DB supports are Red Hat Enterprise Linux 6 and SuSE Linux Enterprise Server (SLES) 11.

Identify the Installation Directory

Determine the directory in which to install the replication components.

The quick-start procedures described here assume you have installed the SAP software on the same host as the Adaptive Server data server. Replication Server and ExpressConnect for

HANA DB are installed in the same base directory, which is identified by the SYBASE environment variable. Select a path on a host drive to be the recipient of the installation and configuration activities.

Downloading Products from SMP

Download replication components from the SAP Service Marketplace (SMP).

Prerequisites

An SMP login ID and password.

Task

Note: If you are installing SAP® Replication Server® Options (RSO), there are several steps you must perform twice, because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/patches>, and enter your login ID and password.
2. Select **SAP Software Download Center > Installations and Upgrades > A-Z Index** and click **R** for SAP Replication Server products.
3. From the product listing, select:
 - **Sybase Replication Server, Option for HANA** or
 - **Sybase Replication Server, HANA Edition**
4. Select the software you need by version and platform.
5. Generate a license based on the license model—served or unserved—and license type:
 - Application Deployment CPU License (AC)
 - Application Deployment Other License (AO)
 - Application Deployment Standby CPU License (BC)
 - CPU License (CP)
 - Development and Testing License (DT)
 - Other License (OT)
 - Standby CPU License (SF)
 - Server License (SR)
 - Standalone Seat License (SS)

Note: To generate a license, you must provide some information, such as the host name, MAC address, and number of CPUs.

6. Download your product licenses, and place them in your product license directory after installation:
 - `SYBASE/SYSAM-2_0/licenses` (UNIX or Linux)
 - `%SYBASE%\SYSAM-2_0\licenses` (Windows)
 where *SYBASE* or *SYBASE%* is where you installed your product.

7. Download your installation files, and uncompress and extract all the installation images to your local drive.

Installing Replication Components

Review the procedures for installing the replication components.

Installing Replication Server

Install Replication Server using the setup program.

Prerequisites

- Allocate a disk partition of at least 20MB for each Replication Server you are installing. You can add more partitions later, if necessary. Make sure each partition is available and has write permissions.
- Allocate the entire partition to the Replication Server. If you allocate only a portion of the partition to Replication Server, you cannot use the remainder for any other purpose.
- Create the sample Replication Server instance `SAMPLE_RS`, which is used throughout the quick-start procedures, when you install Replication Server. If you have already installed Replication Server, see the instructions for setting up a Replication Server instance after installation in *Replication Server ASE-to-ASE Replication Quick Start Guide* > `SAMPLE_RS`.

Task

1. Verify that the drive on which you install the Replication Server has enough disk space for the components being installed, and at least 100MB of extra disk space for the installation program.
2. Download and extract the Replication Server installation image from the SAP® Service Marketplace (SMP).
3. Start the installation.
 - On Windows, launch the **setup** program.
If the installer does not start automatically, double-click **setup.exe** or select **Start** > **Run** and enter:

```
setup.exe
```
 - If there is not enough disk space in the temporary disk space directory, set the environment variable `TMP` to *directory_name* before running it again, where *directory_name* is the full path to and name of the temporary directory to which the installation program writes the temporary installation files.
 - (UNIX or Linux) If you downloaded the product, go to the directory where you extracted the installation image and start the installer:

```
./setup.bin
```

where `setup.bin` is the executable file name for installing Replication Server. If there is not enough disk space in the temporary disk space directory, set the environment variable `IATEMPDIR` to `tmp_dir` before running the installer again, where `tmp_dir` is the full path to the temporary directory to which the installation program writes the temporary installation files.

4. In the Introduction window, click **Next**.
5. Indicate where to install Replication Server:
 - To select an installation directory, click **Choose**, and browse for and select the directory.
 - To create a new directory, enter a new directory path.
 - To restore to the default directory, click **Restore Default Folder**.

If the installation directory you chose does not exist, you see:

```
The directory directory name does not exist. Do you want to create it?
```

Click **Yes**.

If the installation directory you selected exists, and already contains an Replication Server installation, you see:

```
Warning: You have chosen to install into an existing directory. If you proceed with this installation, any older versions of the products you choose to install that are detected in this directory will be replaced.
```

Click **Next**.

6. Select the type of installation:
 - **Typical** – installs the default components. This is recommended for most users.
 - **Full** – installs every Replication Server component, including all the supported language modules.
 - **Custom** – lets you select the components to install. Some components are automatically installed if they are required to run your selected components.

Click **Next**.

7. Select the geographic location, agree to the license agreement, then click **Next**.
8. On the SySAM License Server window, select one of:
 - **Specify License Keys** – click **Browse** to select the license file. To select multiple license files, use **Shift+Click** or **Ctrl+Click**. The license pane displays the license information.
Alternatively, copy and paste the license information directly in the license pane. Click **Next**.
If you specify a served license key, the installer prompts you to install a new SySAM license server. Click:

- **Next** to install the new SySAM license server and follow the installation prompts.
- **Previous** and select **Use previously deployed license server** if you have an existing SySAM license server on the same host.
- **Use previously deployed license server** – for an existing license server, enter the host name of the machine where the license server is running, and the port number if the port number you are using is not the default. Click **Next**.
- **Continue installation without a license key** – click **Next** to proceed without any license keys. The installer allows you to install and use the Replication Server components without a license for 30 days. To continue using these components after the end of the grace period, obtain valid licenses and install these licenses using the Replication Server License Installer.

See the *SySAM Users Guide*.

9. Configuring your server for e-mail notification enables designated users to receive information about license management events requiring attention.

Provide:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipient e-mail addresses
- Message severity that triggers e-mail messages

Click **Next**.

10. The installation summary window displays the selections you have made. Review the information, and click **Install**.

11. On the Start Sample Replication Server window, select:

- **Yes** to configure and start a sample Replication Server. The installer displays the configuration information for the sample Replication Server. Record this information. You must enter and confirm a password from 6 to 15 characters in length.

Note: You cannot start the sample Replication Server on Linux on POWER because the ERSSD requires SAP SQL Anywhere® Server, which is not available on Linux on POWER.

- **No** to manually configure a full-featured Replication Server and start a sample Replication Server after installation. The installer prompts you to either create a sample Replication Server directory or to proceed with the installation. Select either:
 - **Yes** – to create a sample Replication Server directory without starting the sample Replication Server, or,
 - **No** – to continue with the installation.

See the *Replication Server Configuration Guide* for Windows or UNIX.

Note: For a sample Replication Server instance, the `rs_charset` parameter must be set to `utf8`.

Click **Next**.

- When the final window appears, indicating a successful installation, click **Done** to exit the installer.

Editing the Interfaces File

Edit your Replication Server interfaces file.

- Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name instead of `SAMPLE_RS`.

- Shut down Replication Server:

```
1>shutdown
2>go
```

- Change "localhost" in the Replication Server interfaces file to the host name for your primary database.
- To use direct load materialization to materialize a table on the SAP HANA database, add an entry for the primary data server in the interfaces file.
- In the Replication Server interfaces file, add an entry identifying the replicate SAP HANA database:

```
[dataservername]
master tcp ether hostname port
query tcp ether hostname port
```

where *hostname* and *port* are the host and port number of the SAP HANA database, and *dataservername* is a label that identifies the host and port number.

Note: Each SAP HANA database installs with a unique instance number, and the port number is the instance number prefixed with 3 and suffixed with 15:

```
3in15
```

where *in* is the two-digit instance number. For example, the port number for an SAP HANA database with instance number 1 is 30115.

- Navigate to `$$SYBASE/$$SYBASE_REP/samp_repserver`, and execute `RUN_SAMPLE_RS`.

See also

- *Server Entries to the Interfaces Files* on page 9
- *Adding Server Entries to the Interfaces Files* on page 9

Configuring Replication Components

Review the procedures for configuring each of the replication components in the replication system.

Configuration Checklist

Use the table as a checklist to navigate through the configuration when you set up your replication environment.

Table 3. Configuration Checklist

	Action	User	Topic
1	Replication Server: Set up ExpressConnect for HANA DB.		<i>Obtaining the SAP HANA Database ODBC Drivers</i>
2	Replication Server: Create a Replication Server connection to SAP HANA database.	Maintenance user	<i>Configuring Replication Server for Replication to SAP HANA Database</i>
3	Replication Server: Configure Replication Server for replication from the primary database.	rs_username	<i>Configuring Replication Server for Replication from the Primary Database</i>
4	Test that replication works as intended.		<i>Testing Replication on page 23</i>

See also

- *User Management* on page 137

Configuring Replication Server

Configure Replication Server for ASE-to-SAP HANA database replication.

Obtaining the SAP HANA Database ODBC Drivers

Set up ExpressConnect for HANA DB for connections between SAP HANA database and Replication Server.

You must add the SAP HANA database ODBC driver to the ExpressConnect for HANA DB installation. Download these libraries from the SAP Service Marketplace and install them after you have installed Replication Server. See the *Replication Server Release Bulletin > Special Installation Instructions > Installing ODBC Libraries for ExpressConnect for HANA DB*.

Configuring Replication Server for Replication to SAP HANA Database

Create a Replication Server connection to SAP HANA database.

If you are not using the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

1. If you are using SAP Secure User Store, create a user store of encrypted credentials:

```
hdbuserstore set rds myhost:xxxxx my_securestore_user
my_securestore_pwd
```

where:

- *rds* is the key for the secure store entry.
- *myhost.xxxxx* is the connection environment host name and port number.
- *my_securestore_user* and *my_securestore_pwd* are SAP Secure User Store credentials.

2. On SAP HANA database, create a maintenance user:

- a) As the SAP HANA database admin user, create a user ID with a temporary password:

```
CREATE USER muser PASSWORD mpwd_temp
```

- b) Log in to SAP HANA database with the temporary password, and then update the password:

```
ALTER USER muser PASSWORD mpwd_new
```

- c) Log out of SAP HANA database, and log in again with the new password.

3. Log in to SAP HANA database:

```
./hdbsql -u user -p password -i id
```

where:

- *user* is the SAP HANA database user who has the authority to grant access to the maintenance user (*muser*).
- *password* is the *user* password.
- *id* is the instance number.

4. On SAP HANA database, grant these authorities to the maintenance user:

```
GRANT CREATE ANY, DELETE, DROP, EXECUTE,
INDEX, SELECT, UPDATE ON SCHEMA myschema TO muser
```

where *myschema* is your SAP HANA database schema, and *muser* is the maintenance user you just created.

5. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name instead of SAMPLE_RS.

6. Create a connection to the replicate SAP HANA database instance using ExpressConnect for HANA DB.

ASE-to-SAP HANA Database Replication Setup

For a standard connection:

```
create connection to rds.rdb
using profile rs_ase_to_hanadb;ech
set username muser
set password mpwd
go
```

For SAP Secure User Store:

```
create connection to rds.rdb
using profile rs_ase_to_hanadb;ech
set username auser
set password apwd
set dsi_connector_sec_mech to "hdbuserstore"
go
```

where:

- *rds* is the replicate SAP HANA database. For a standard connection, this must match the data server name in the Replication Server interfaces file entry. See the *Adaptive Server Enterprise Configuration Guide for UNIX > Set Up Communications Across the Network > Contents of the interfaces File* and the *Adaptive Server Enterprise Configuration Guide for Windows > Network Communications Using sql.ini*. For an SAP Secure User Store connection, this must match what you used as the key to create a user store of encrypted credentials with the **hdbuserstore** utility.
- *rdb* is a placeholder; you must provide a value, but it is not used.
- *muser* is the maintenance user for the replicate SAP HANA database instance you have already created.
- *mpwd* is the replicate SAP HANA database maintenance user password.
- *auser* and *apwd* are unused values supplied only to satisfy the syntax of the **create connection** command.

Note: For an SAP Secure User Store connection, the same user who ran the **hdbuserstore** utility must run the **create connection** command.

If you have trouble connecting to the replicate SAP HANA database instance using ExpressConnect for HANA DB, see the *Replication Server Troubleshooting Guide > Common Error Messages > Connector Error Messages*.

7. Use **admin show_connections, 'replicate'** to display the replicate connection you created.

Configuring Replication Server for Replication from the Primary Database

Configure Replication Server for replication from the primary database.

Adding the Primary Database to Replication Server

Add the primary Adaptive Server database to Replication Server.

1. Add the Replication Server instance to the interfaces file of the data server that hosts the primary database.

Note: Restart the Adaptive Server instance if the interfaces file has changed.

2. Go to `$SYBASE/$SYBASE_REP/init/rs`.
3. Make a copy of the `setupdb.rs` file and rename it as `myprimary.rs`.
4. Edit the `myprimary.rs` file.

Note: Modify the parameters described in this table.

Table 4. Sample Values for the myprimary.rs File

Parameter	Description	Value
sybinit.release_directory	Valid path of the Replication Server software (\$SYBASE).	/opt/software
rs.rs_name	Name of the Replication Server.	SAMPLE_RS
rs.rs_rs_sa_user	User ID that has “sa” privileges on Replication Server.	sa
rs.rs_rs_sa_pass	Password of the “sa” user.	sa_pass
rs.rs_ds_name	Name of the data server that hosts the primary database.	sunak1502i
rs.rs_ds_sa_user	User ID that has “sa” privileges on data server.	sa
rs.rs_ds_sa_password	Password of the “sa” user for the data server.	password
rs.rs_db_name	Name of the primary database.	pubs2
rs.rs_needs_repagent	Whether you plan to replicate from the specified primary database.	yes
rs.rs_db_maint_user	ID for the user whose work is not replicated when logged in to the primary database (the maintenance user.)	<database>_maint If the user ID does not exist, the script creates the user ID on the database. The user ID cannot be the name of an alias
rs.rs_db_maint_password	Password for the maintenance user.	<database>_maint_ps

Parameter	Description	Value
rs.rs_ltm_rs_user	User that the Replication Agent uses to log in to the Replication Server. The name must exist, and typically comes from values that were set up during Replication Server creation.	<i>SAMPLE_RS_ra</i>
rs.rs_ltm_rs_pass	Password of the rs.rs_ltm_rs_user .	sa_pass
rs.rs_db_physical_for_logical	Whether this is a warm standby database.	no

5. Save the file.
6. Go to `$SYBASE/$SYBASE_REP/install`.
7. Create the connection from the primary database to the Replication Server by running the resource file:

```
./rs_init -r ../init/rs/myprimary.rs
```

If the **rs_init** command fails, check the log file in `$SYBASE/$SYBASE_REP/init/logs`, correct the issue, then disable the RepAgent as follows:

- a. Log in to the primary Adaptive Server using an “sa” user role and access the primary database.
- b. Disable the RepAgent thread in the primary database:

```
sp_config_rep_agent pubs2, 'disable'
go
```

Re-run **rs_init**.

8. Validate the primary connection:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

9. Enter:

```
admin who
go
```

The output from the command looks similar to:

```
Spid  Name  State          Info
----  -
32 DSI EXEC Awaiting Command 101(1)
SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD

18 DSI          Awaiting Message 101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
23 DIST        Awaiting Wakeup  101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD

29 SQT          Awaiting Wakeup  101:1 DIST
SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
```

```

11 SQM      Awaiting Message 101:1 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
9   SQM      Awaiting Message 101:0 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
36 REP AGENT Awaiting Command SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
156 DSI EXEC  Awaiting Command 270(1) sunak1502i.pubs2
140 DSI      Awaiting Message 270 sunak1502i.pubs2
136 DIST    Awaiting Wakeup 270 sunak1502i.pubs2
134 SQT      Awaiting Wakeup 270:1 DIST sunak1502i.pubs2
142 SQM      Awaiting Message 270:1 sunak1502i.pubs2
155 SQM      Awaiting Message 270:0 sunak1502i.pubs2
135 REP AGENT Awaiting Command sunak1502i.pubs2
162 DSI EXEC  Awaiting Command 268(1) rds.rdb
157 DSI      Awaiting Message 268 rds.rdb
133 SQM      Awaiting Message 268:0 rds.rdb
21  dSUB      Sleeping
6   dCM      Awaiting Message
7   dAIO     Awaiting Message
26  dREC      Sleeping dREC
8   dDELSEG  Awaiting Message
163 USER     Active sa
5   dALARM   Awaiting Wakeup
27  dSYSAM   Sleeping

```

10. Exit the `isql` session.

Testing Replication

When you finish setting up the replication system, test that replication works as intended.

1. Connect to the primary Adaptive Server instance as a non-maintenance user with the ability to create tables.
2. At the primary Adaptive Server database, create a table named `ptab1` to replicate:

```

use pubs2
go
create table ptab1
(idno int not null,
name varchar(20) null)
go
sp_primarykey ptab1, idno
go

```

Note: By default, this DDL creates a table with the owner `dbo`.

3. Grant permissions to any new or existing object to be replicated in the primary database:

```

grant all on ptab1 to public
go

```

4. Mark the `ptab1` table for replication:

```

sp_setreptable ptab1, 'true'
go
The replication status for 'ptab1' is set to true, owner_off.

```

5. Add rows to the table, ptab1:

```
insert into ptab1 values (1, "Burt Gringlesby")
insert into ptab1 values (2, "Douglas Wong")
insert into ptab1 values (3, "Tom Servo")
insert into ptab1 values (4, "Innes del Castillo")
insert into ptab1 values (5, "Akiko Yokomoto")
go
```

6. At the replicate SAP HANA database instance, create a table named PTAB1:

```
CREATE TABLE PTAB1
(IDNO INT PRIMARY KEY,
NAME VARCHAR(20))
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

If an owner is not specified with the table name, the owner of the table—*<tableowner>*—is the user who is logged in when the table is created.

7. Grant permissions to any new or existing object to be replicated in the replicate SAP HANA database so that the Replication Server maintenance user can update this table:

```
grant all privileges on PTAB1 to public
```

8. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

9. Create a replication definition against the primary Adaptive Server database.

To create the *ptab1_repdef* replication definition for the *ptab1* table on the primary database, enter:

```
create replication definition ptab1_repdef
with primary at pds.pdb
with primary table named ptab1
with replicate table named <tableowner>.ptab1
(idno integer,
"name" varchar(20))
primary key (idno)
go
```

where the replicate table on SAP HANA database—*ptab1*—is owned by *<tableowner>*, and *sunak1502i.pubs2* is the name of the primary connection created earlier.

10. Create a subscription against the replicate connection to materialize the PTAB1 table on the SAP HANA database.

Create the *ptab1_sub* subscription with **direct_load** option for the *ptab1_repdef* replication definition:

```
create subscription ptab1_sub
for ptab1_repdef
with replicate at rds.rdb
without holdlock
```

```
direct_load
user puser password ppwd
go
```

where:

- *rds.rdb* is the replicate SAP HANA database connection created earlier.
- *puser* is the user who selects the data from the table in the primary Adaptive Server database.
- *ppwd* is the password of the *puser*. You must enter a password.

Note: The interfaces file used by the replicate Replication Server must contain an entry for the primary Adaptive Server where the initial data resides.

To create the subscriptions at the replicate Replication Server using the no-materialization method, use the **create subscription** command with the **without materialization** clause.

For information about using the **create subscription** command in **direct_load** materialization, see *Replication Server Reference Manual*.

11. Check the subscription status to ensure there are no errors:

```
check subscription ptabl_sub
for ptabl_repdef
with replicate at rds.rdb
go
```

If there are errors, see *Troubleshooting Subscription Problems*.

12. Connect to the primary Adaptive Server database as a non-maintenance user with the ability to insert rows into the table:

```
insert into ptabl values (10, "Michel DeFrance")
insert into ptabl values (11, "Dirk Stringer")
go
```

13. Check the status of the subscription created with the **direct_load** option by using the **check subscription** command. When the subscription is in the **VALID** state, materialization is complete:

```
check subscription ptabl_sub
for ptabl_repdef
with replicate at rds.rdb
go
Subscription ptabl_sub is VALID at the replicate.
```

14. After the status is **VALID** at the primary and replicate database, on the SAP HANA database, verify that the data is replicated to the replicate PTAB1 table:

```
SELECT * FROM PTAB1
```

See also

- *Troubleshooting Subscription Problems* on page 142

Resetting the Primary ASE Database For Replication

In a test environment, you may occasionally need to reset the replication environment.

1. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you did not set up the sample Replication Server instance, enter your Replication Server instance name instead of SAMPLE_RS.

2. Connect to the primary database:

```
1>connect to pds.pdb
2>go
```

where:

- *pds.pdb* is the primary connection.
- *pds* is the Adaptive Server data server of the primary connection.
- *pdb* is the Adaptive Server database of the primary connection.

3. Stop the Replication Agent instance:

```
1>sp_stop_rep_agent pdb
2>go
```

4. Set the database to release the secondary truncation point:

```
1>dbcc settrunc ('ltm','ignore')
2>go
```

5. Disconnect from the primary database:

```
1>disconnect
2>go
```

6. Connect to the RSSD:

```
1>connect to rssid
2>go
```

7. Reset the locator stored in Replication Server:

```
1> rs_zeroltm pds, pdb
1> go
```

8. Disconnect from the RSSD:

```
1>disconnect
2>go
```

9. Connect again to the primary database:

```
1>connect to pds.pdb
2>go
```

10. Reset the truncation point to the end of the database transaction log:

```
1>dbcc settrunc ('ltm','valid')
2>go
```

11. Start the Replication Agent instance:


```
1>sp_start_rep_agent pdb
2>go
```

12. Disconnect from the primary database:

```
1>disconnect
2>go
```

13. Verify that the Replication Agent instance has connected to Replication Server:

```
1>admin who
2>go
```

If the Replication Agent instance has connected, Replication Server shows the connection status. For example:

```
27 REP AGENT  Awaiting Command
                pds.pdb
```


Oracle-to-SAP HANA Database Replication Setup

Implement a sample replication system for replication from Oracle to SAP HANA database.

Use the tasks in this document to set up a proof-of-concept or test replication system. Only basic Oracle features are addressed in the example scenario. In particular, Oracle Automatic Storage Management (ASM) and Real Application Clusters (RAC) are not addressed in this guide, although Replication Agent does support these features. See the *Replication Agent Primary Database Guide*

The tasks use the sample Replication Server instance SAMPLE_RS. If you have already installed Replication Server, skip to the configuration procedures.


The replication of sequences is not discussed in this document. To set up sequence replication, see the *Replication Agent Installation Guide* and the *Replication Agent Primary Database Guide*.

For the most current information about Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

Multiple components are required to implement an Oracle-to-SAP HANA database replication system: a primary Oracle data server, a replicate SAP HANA database, a Replication Server with ExpressConnect for HANA DB, and Replication Agents for Oracle.

Figure 2: Oracle-to-SAP HANA Database Replication Setup



Prerequisites

Review the prerequisites before you begin setting up Oracle-to-SAP HANA database replication.

- Familiarity with Oracle and SAP HANA database.
- Valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- Install and configure Oracle Database 10g or 11g Enterprise Edition to serve as the primary (source) database from which Replication Agent replicates transactions, and SAP HANA database as the replicate (target) database.
- If Oracle 10g or 11g Release 1 is the primary database, install the Oracle JDBC™ thin driver for Oracle 10g or 11g and for JDK 1.4 and 1.5. If the primary database is Oracle 11g Release 2, install the Oracle JDBC thin driver 11.2 for JDK 1.6.
- TCP/IP connectivity between all hosts.
- Install any operating system patches required for Java.
- The Replication Agent has direct access to the Oracle online and archived redo logs.

Do not use Replication Server reserved words for object names and connection names. The *Replication Server Reference Manual* includes a complete list of reserved words.

Planning Your Installation

Before installing the replication components, verify the system requirements, identify the installation directories, and obtain the installation software.

Installation Checklist

Use the checklist to navigate through the installation when you set up your replication environment.

Table 5. Installation Checklist

	Action	Topic
1	Make sure the target computer meets the minimum memory and disk space requirements.	<i>System Requirements</i>
2	Determine the directory in which to install the replication components.	<i>Identify the Installation Directory</i>
3	Download replication components from the SAP Service Marketplace (SMP).	<i>Downloading Products from SMP</i>

	Action	Topic
4	Install Replication Server.	<i>Installing Replication Server</i>
5	Install Replication Agent.	<i>Installing Replication Agent</i>

System Requirements

Make sure the target computer on which you are installing the replication components meets the minimum memory and disk space requirements.

See the *Replication Server Installation Guide > Planning Your Installation > System Requirements* and the *Replication Agent Installation Guide > Planning Your Installation > System Requirements*.

To improve performance by avoiding multiple network hops, install Replication Server and the SAP HANA database on the same machine. Replication Agent must have direct access to the Oracle online and redo logs and must be installed on the same machine as the primary database.

Note: The only versions of Linux that ExpressConnect for HANA DB supports are Red Hat Enterprise Linux 6 and SuSE Linux Enterprise Server (SLES) 11.

Identify the Installation Directory

Determine the directory in which to install the replication components.

The quick-start procedures described here assume you have installed the SAP software on the same host as the Oracle data server. Replication Agent, Replication Server, and ExpressConnect for HANA DB database are installed in the same base directory, which is identified by the SYBASE environment variable. Select a path on a host drive to be the recipient of the installation and configuration activities. Replication Agent must be installed on the same type of host as the primary Oracle instance and must be able to directly access the primary Oracle redo and archive logs.

Replication Agent Installation Directory

Replication Agent versions 15.5 and later use the RAX-15_5 directory name. If you have an existing Replication Agent installation of version 15.5 or later and do not want to overwrite it, install the latest version of Replication Agent in a different location.

Downloading Products from SMP

Download replication components from the SAP Service Marketplace (SMP).

Prerequisites

An SMP login ID and password.

Task

Note: If you are installing SAP® Replication Server® Options (RSO), there are several steps you must perform twice, because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/patches>, and enter your login ID and password.
2. Select **SAP Software Download Center > Installations and Upgrades > A-Z Index** and click **R** for SAP Replication Server products.
3. From the product listing, select:
 - **Sybase Replication Server, Option for HANA** or
 - **Sybase Replication Server, HANA Edition**
4. Select the software you need by version and platform.
5. Generate a license based on the license model—served or unserved—and license type:
 - Application Deployment CPU License (AC)
 - Application Deployment Other License (AO)
 - Application Deployment Standby CPU License (BC)
 - CPU License (CP)
 - Development and Testing License (DT)
 - Other License (OT)
 - Standby CPU License (SF)
 - Server License (SR)
 - Standalone Seat License (SS)

Note: To generate a license, you must provide some information, such as the host name, MAC address, and number of CPUs.

6. Download your product licenses, and place them in your product license directory after installation:
 - `$SYBASE/SYSAM-2_0/licenses` (UNIX or Linux)
 - `%SYBASE%\SYSAM-2_0\licenses` (Windows)where `$SYBASE` or `%SYBASE%` is where you installed your product.
7. Download your installation files, and uncompress and extract all the installation images to your local drive.

Installing Replication Components

Review the procedures for installing the replication components.

See also

- *Configuring Replication Components* on page 18

Installing Replication Server

Install Replication Server using the setup program.

Prerequisites

- Allocate a disk partition of at least 20MB for each Replication Server you are installing. You can add more partitions later, if necessary. Make sure each partition is available and has write permissions.
- Allocate the entire partition to the Replication Server. If you allocate only a portion of the partition to Replication Server, you cannot use the remainder for any other purpose.
- The sample Replication Server instance `SAMPLE_RS` used in these procedures must be created when you install Replication Server. If you have already installed Replication Server, see the instructions for how to set up a Replication Server instance after installation in *Replication Server ASE-to-ASE Replication Quick Start Guide > SAMPLE_RS*.

Task

1. Verify that the drive on which you install the Replication Server has enough disk space for the components being installed, and at least 100MB of extra disk space for the installation program.
2. Download and extract the Replication Server installation image from the SAP® Service Marketplace (SMP).
3. Start the installation.

- On Windows, launch the **setup** program.

If the installer does not start automatically, double-click **setup.exe** or select **Start > Run** and enter:

```
setup.exe
```

If there is not enough disk space in the temporary disk space directory, set the environment variable `TMP` to *directory_name* before running it again, where *directory_name* is the full path to and name of the temporary directory to which the installation program writes the temporary installation files.

- (UNIX or Linux) If you downloaded the product, go to the directory where you extracted the installation image and start the installer:

```
./setup.bin
```

where `setup.bin` is the executable file name for installing Replication Server.

If there is not enough disk space in the temporary disk space directory, set the environment variable `IATEMPDIR` to *tmp_dir* before running the installer again, where *tmp_dir* is the full path to the temporary directory to which the installation program writes the temporary installation files.

4. In the Introduction window, click **Next**.
5. Indicate where to install Replication Server:
 - To select an installation directory, click **Choose**, and browse for and select the directory.
 - To create a new directory, enter a new directory path.
 - To restore to the default directory, click **Restore Default Folder**.

If the installation directory you chose does not exist, you see:

```
The directory directory name does not exist. Do you want to create it?
```

Click **Yes**.

If the installation directory you selected exists, and already contains an Replication Server installation, you see:

```
Warning: You have chosen to install into an existing directory. If you proceed with this installation, any older versions of the products you choose to install that are detected in this directory will be replaced.
```

Click **Next**.

6. Select the type of installation:
 - **Typical** – installs the default components. This is recommended for most users.
 - **Full** – installs every Replication Server component, including all the supported language modules.
 - **Custom** – lets you select the components to install. Some components are automatically installed if they are required to run your selected components.

Click **Next**.

7. Select the geographic location, agree to the license agreement, then click **Next**.

8. On the SySAM License Server window, select one of:

- **Specify License Keys** – click **Browse** to select the license file. To select multiple license files, use **Shift+Click** or **Ctrl+Click**. The license pane displays the license information.

Alternatively, copy and paste the license information directly in the license pane. Click **Next**.

If you specify a served license key, the installer prompts you to install a new SySAM license server. Click:

- **Next** to install the new SySAM license server and follow the installation prompts.
- **Previous** and select **Use previously deployed license server** if you have an existing SySAM license server on the same host.

- **Use previously deployed license server** – for an existing license server, enter the host name of the machine where the license server is running, and the port number if the port number you are using is not the default. Click **Next**.
- **Continue installation without a license key** – click **Next** to proceed without any license keys. The installer allows you to install and use the Replication Server components without a license for 30 days. To continue using these components after the end of the grace period, obtain valid licenses and install these licenses using the Replication Server License Installer.

See the *SySAM Users Guide*.

9. Configuring your server for e-mail notification enables designated users to receive information about license management events requiring attention.

Provide:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipient e-mail addresses
- Message severity that triggers e-mail messages

Click **Next**.

10. The installation summary window displays the selections you have made. Review the information, and click **Install**.
11. On the Start Sample Replication Server window, select:

- **Yes** to configure and start a sample Replication Server. The installer displays the configuration information for the sample Replication Server. Record this information. You must enter and confirm a password from 6 to 15 characters in length.

Note: You cannot start the sample Replication Server on Linux on POWER because the ERSSD requires SAP SQL Anywhere Server, which is not available on Linux on POWER.

- **No** to manually configure a full-featured Replication Server and start a sample Replication Server after installation. The installer prompts you to either create a sample Replication Server directory or to proceed with the installation. Select either:
 - **Yes** – to create a sample Replication Server directory without starting the sample Replication Server, or,
 - **No** – to continue with the installation.

See the *Replication Server Configuration Guide* for Windows or UNIX.

Note: For a sample Replication Server instance, the **rs_charset** parameter must be set to utf8.

Click **Next**.

12. When the final window appears, indicating a successful installation, click **Done** to exit the installer.

Installing Replication Agent

Install Replication Agent using the GUI wizard.

Note: If there is not enough disk space in your default temporary directory, set the IATEMPDIR (Linux or UNIX) or TMP (Windows) environment variable to a directory that has enough space. Include the full path to this directory.

1. Log in to the Replication Agent host machine using an operating system user account with authority to start, stop, and administer the Replication Agent instance (for example, the “sybase” user).
2. Close all nonessential applications, and minimize any open windows.
3. Download and extract the Replication Agent installation image from the SAP Service Marketplace (SMP).
4. Start the installation program:
 - On Microsoft Windows platforms, the installation program should start automatically. If it does not, start the installation program from Microsoft Windows Explorer by double-clicking the **setup.exe** file.
 - On UNIX platforms, enter:

```
./setup.bin
```
5. Click **Next**.

Note: You can click **Cancel** to stop the installation at any point.

6. Select the geographic location, agree to the license agreement, then click **Next**.
7. Specify an installation directory. The default installation directory is:
 - Existing `%SYBASE%` or `c:\sybase` on Microsoft Windows platforms
 - Existing `$SYBASE` or `/opt/sybase` on UNIX platforms

Click **Next** to accept the default installation directory, or:

- Click **Choose** to select an installation directory in the file browser. Click **Next**.
- Enter a directory name. Click **Next**.

If you continue and the older products were not installed with the current version of Replication Agent, the installation program overwrites any common files.

If the directory name does not exist, click **Yes** when prompted to create it:

```
The directory does not exist.  
Do you want to create it?
```

If the directory you specified exists, you see:

```
Warning: You have chosen to install into an existing  
directory. If you proceed with this installation,
```

any older versions of the products you choose to install that are detected in this directory will be replaced.

Note: On Microsoft Windows platforms, if you are prompted to overwrite a DLL, click **Yes** only if the version of the new DLL is newer than the one the installation program is attempting to overwrite.

8. Select the type of database the Replication Agent will connect to.

Click **Next**.

9. Select one of these options to enter the license:

- **Specify license keys** – browse to or specify the license file.
- **Use previously deployed license server** – use a previously deployed license server. Enter the host name of the machine where the license server is running and the port number if the port number you are using is not the default.
- **Continue installation without a license key** – install and use Replication Agent without a license for a grace period of 30 days. To continue using Replication Agent after the end of the grace period, obtain a valid license from the SAP Service Marketplace (SMP), and install it.

Note: Replication Agent is licensed statically when installed with RSHE. If you are installing Replication Agent with RSHE, locate the file that contains a valid SySAM license for the Replication Agent component in RSHE. Specify that license file here.

See the *SySAM Users Guide*.

Click **Next** until you see the SySAM Notification window.

10. On the SySAM Notification window, configure your server for e-mail notification. When configuration is enabled, you receive information about license management events that require attention. Select **Yes** and either accept the default values that are supplied, or enter values for:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipients of the notification
- Message severity level of an event that triggers e-mail notification:
 - Informational
 - Warning
 - Error

If you choose not to have e-mail alerts or severity messages logged, select **No**.

Click **Next**.

11. Review the product features or components listed on the Preinstallation Summary window. Click **Install**.

Oracle-to-SAP HANA Database Replication Setup

The installation program installs the components in the installation directory you specified, and displays an installation progress indicator.

If errors occur during the installation, the installation program displays error messages. Exit the installation program wizard to correct the cause of the error, then restart the installation program.

If the software is installed successfully, you see a window confirming the successful installation.

12. Click **Done** to complete the installation and close the installation program.

Verifying the Installation

Verify that Replication Agent has been successfully installed.

1. In a command window, change to the directory where you installed Replication Agent.
2. Set the environment variables by sourcing the `SYBASE.csh` file (UNIX or Linux) or by executing the `SYBASE.bat` file (Windows).
3. Change to `$SYBASE/RAX-15_5/bin` (UNIX or Linux) or `%SYBASE%\RAX-15_5\bin` (Windows).
4. Obtain the Replication Agent version string:

- For Windows:

```
ra.bat -v
```

```
rao -v
```

- For UNIX or Linux:

```
./ra.sh -v
```

```
./rao.sh -v
```

Note: Make sure permission for the `ra.sh` file is set no broader than 700 (read/write/execute for user, no permissions for group and other).

If Replication Agent installs successfully, you see the SAP copyright and the Replication Agent version string.

Configuring Replication Components

Review the procedures for configuring each of the replication components in the replication system.

Configuration Checklist

Use the table as a checklist to navigate through the configuration when you set up your replication environment.

Table 6. Configuration Checklist

	Action	User	Topic
1	Create a Replication Server connection to the primary database.	Maintenance user	<i>Creating a Replication Server Connection to the Primary Database</i>
2	Replication Server: Set up ExpressConnect for HANA DB.		<i>Obtaining the SAP HANA Database ODBC Drivers</i>
3	Replication Server: Create a Replication Server connection to SAP HANA database.	Maintenance user	<i>Configuring Replication Server for Replication to SAP HANA Database</i>
4	Replication Server: Configure Replication Server for replication from the primary database.	rs_username	<i>Configuring Replication Server for Replication from the Primary Database</i>
5	Oracle database: Verify that redo logging is turned on.	Primary database administrative user	<i>Verifying the Current Archive Setting of the Redo Logs on page 44</i>
6	Oracle database: Verify that Oracle logs the values of the primary keys and unique indexes.	Primary database administrative user	<i>Verifying the Supplemental Logging of Primary Key Data</i>
7	Oracle database: Configure Oracle LogMiner	Primary database administrative user	<i>Configuring Oracle LogMiner</i>
8	Oracle database: Create an Oracle database user for use by Replication Agent.	Primary database administrative user	<i>Creating Oracle User and Grant Permissions</i>
9	Oracle database: Validate the required Oracle settings.	Primary database administrative user	<i>Verifying the Oracle User Roles Created by the Replication Agent</i>
10	Prepare and create a Replication Agent instance for each Oracle database being replicated.		<i>Creating the Replication Agent Instance</i>
11	Verify that the Replication Agent instance has been successfully installed.		<i>Verifying the Replication Agent Instance Installation</i>

	Action	User	Topic
12	Initialize the Replication Agent instance.		<i>Initializing the Replication Agent Instance</i>

See also

- *User Management* on page 137

Configuring Replication Server

Configure Replication Server for Oracle-to-SAP HANA database replication.

Creating a Replication Server Connection to the Primary Database

Create a Replication Server connection to the primary database.

1. In Replication Server, create a connection to the primary Oracle database:

```
create connection to pds.pdb
using profile rs_rs_to_oracle_ra;standard
set username muser
set password mpwd
with log transfer on, dsi_suspended
go
```

where:

- *pds* is the value of the **rs_source_ds** parameter specified in Replication Agent.
- *pdb* is the value of **rs_source_db** specified in Replication Agent.
- *muser* is the maintenance user for the primary Oracle database. This user must already exist and have the necessary select permissions in the primary database. See the *Replication Server Options Primary Database Guide*.
- *mpwd* is the maintenance user password.

For information about the maintenance user, see the *Replication Server Heterogeneous Replication Guide > SAP Replication Products > Replication Server > Database Connections > Maintenance User Purpose*.

2. Use **admin show_connections, 'primary'** to display the primary connection you created.

Obtaining the SAP HANA Database ODBC Drivers

Set up ExpressConnect for HANA DB for connections between SAP HANA database and Replication Server.

You must add the SAP HANA database ODBC driver to the ExpressConnect for HANA DB installation. Download these libraries from the SAP Service Marketplace and install them after you have installed Replication Server. See the *Replication Server Release Bulletin > Special Installation Instructions > Installing ODBC Libraries for ExpressConnect for HANA DB*.

Configuring Replication Server for Replication to SAP HANA Database

Create a Replication Server connection to SAP HANA database.

If you are not using the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

1. If you are connecting to SAP HANA database with a standard connection and not using SAP Secure User Store, add an entry to your Replication Server interfaces file identifying the replicate SAP HANA database, and then stop and start Replication Server:

```
[dataservername]
master tcp ether hostname port
query tcp ether hostname port
```

where *hostname* and *port* are the host and port number of the SAP HANA database, and *dataservername* is a label used to identify the host and port number. For information about stopping and starting Replication Server, see the *Replication Server Administration Guide Volume 1 > Manage a Replication System*.

Note: Each SAP HANA database installs with a unique instance number, and the port number is the instance number prefixed with 3 and suffixed with 15:

```
3in15
```

where *in* is the two-digit instance number. For example, the port number for an SAP HANA database with instance number 1 is 30115.

If you are using SAP Secure User Store, create a user store of encrypted credentials:

```
hdbuserstore set rds myhost:xxxxx my_securestore_user
my_securestore_pwd
```

where:

- *rds* is the key for the secure store entry.
- *myhost.xxxxx* is the connection environment host name and port number.
- *my_securestore_user* and *my_securestore_pwd* are SAP Secure User Store credentials.

2. On SAP HANA database, create a maintenance user:
 - a) As the SAP HANA database admin user, create a user ID with a temporary password:


```
CREATE USER muser PASSWORD mpwd_temp
```
 - b) Log in to SAP HANA database with the temporary password, and then update the password:


```
ALTER USER muser PASSWORD mpwd_new
```
 - c) Log out of SAP HANA database, and log in again with the new password.
3. Log in to SAP HANA database:


```
./hdbsql -u user -p password -i id
```

Oracle-to-SAP HANA Database Replication Setup

where *user* is the SAP HANA database user, *password* is the user password, and *id* is the instance number.

4. On SAP HANA database, grant these authorities to the maintenance user:

```
GRANT CREATE ANY, DELETE, DROP, EXECUTE,  
INDEX, SELECT, UPDATE ON SCHEMA myschema TO muser
```

where *myschema* is your SAP HANA database schema, and *muser* is the maintenance user you just created.

5. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name instead of SAMPLE_RS.

6. Create a connection to the replicate SAP HANA database instance using ExpressConnect for HANA DB.

For a standard connection:

```
create connection to rds.rdb  
using profile rs_oracle_to_hanadb;ech  
set username muser  
set password mpwd  
go
```

For SAP Secure User Store:

```
create connection to rds.rdb  
using profile rs_oracle_to_hanadb;ech  
set username auser  
set password apwd  
set dsi_connector_sec_mech to "hdbuserstore"  
go
```

where:

- *rds* is the replicate SAP HANA database. For a standard connection, this must match the data server name in the interfaces file entry. See the *Adaptive Server Enterprise Configuration Guide for UNIX > Set Up Communications Across the Network > Contents of the interfaces File* and the *Adaptive Server Enterprise Configuration Guide for Windows > Network Communications Using sql.ini*. For an SAP Secure User Store connection, this must match what you used as the key to create a user store of encrypted credentials with the **hdbuserstore** utility.
- *rdb* is a placeholder; you must provide a value, but it is not used.
- *muser* is the maintenance user for the replicate SAP HANA database instance you have already created.
- *mpwd* is the replicate SAP HANA database maintenance user password.
- *auser* and *apwd* are unused values supplied only to satisfy the syntax of the **create connection** command.

Note: For an SAP Secure User Store connection, the same user who ran the **hdbuserstore** utility must run the **create connection** command.

If you have trouble connecting to the replicate SAP HANA database instance using ExpressConnect for HANA DB, see the *Replication Server Troubleshooting Guide > Common Error Messages > Connector Error Messages*.

7. Use **admin show_connections, 'replicate'** to display the replicate connection you created.

Configuring Replication Server for Replication from the Primary Database

Configure Replication Server for replication from the primary database by granting create object permission to the **rs_username** user.

The **rs_username** user, which is used for communication between Replication Agent and Replication Server, must have **create object** permission before you can use it to create a replication definition. You must grant this permission manually from Replication Server.

1. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

```
isql -Uusername -Ppassword -SSAMPLE_RS
```

2. Enter:

```
grant create object to SAMPLE_RS_ra
go
```

where *SAMPLE_RS_ra* is the value for the **rs_username** parameter you entered in the resource file.

Configuring Replication Agent

Review the procedures for configuring Replication Agent for Oracle.

Primary Oracle Database Configuration for Replication

Configure the primary Oracle database for replication and to ensure that the Oracle redo log data is not overwritten before it is read by Replication Agent.

In addition to Oracle LogMiner, you can now alternatively configure the Replication Agent for Oracle Log Reader component to use Oracle XStream APIs to read from the transaction logs. The XStream APIs enable you to share data changes between Oracle and non-Oracle data sources, such as SAP Replication Server.

Note: You can configure the Replication Agent for Oracle Log Reader component to use XStream APIs from version 15.7.1 SP120 and later.

Verifying the Current Archive Setting of the Redo Logs

Verify that redo logging, which Oracle uses to maintain a log infrastructure, is turned on in the primary database.

1. Use SQLPLUS to connect to Oracle as a system administrator.
2. From SQLPLUS, run:

```
select log_mode from v$database;
```

- If the archive log is on, you should see:

```
LOG_MODE  
-----  
ARCHIVELOG
```

3. To turn on log archiving:

```
shutdown;  
startup mount;  
alter database archivelog;  
alter database open;
```

See the *Replication Agent Primary Database Guide*.

Verifying the Supplemental Logging of Primary Key Data

By default, Oracle does not log primary keys and unique indexes. For successful replication of all table values, you must log these values.

1. Use SQLPLUS to connect to Oracle as a system administrator:

```
SELECT SUPPLEMENTAL_LOG_DATA_MIN,  
SUPPLEMENTAL_LOG_DATA_PK,  
SUPPLEMENTAL_LOG_DATA_UI  
FROM V$DATABASE;
```

If logging of primary key and unique index values is enabled, the return values are:

```
SUP SUP SUP  
--- --- ---  
YES YES YES
```

2. If the result is different, turn on supplemental logging by executing:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA (PRIMARY  
KEY, UNIQUE INDEX) COLUMNS;  
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Configuring Oracle LogMiner

Configure Oracle LogMiner on the primary Oracle database.

1. If you have not already installed Oracle LogMiner, go to \$ORACLE_HOME/rdbms/admin, log in as a “sys as sysdba” user, and execute the Oracle LogMiner installation script:

```
@dbmslm.sql
```

2. If the public synonym DBMS_LOGMNR does not exist, create a public synonym so that you do not have to log in as the owner to execute LogMiner functions:

```
CREATE PUBLIC SYNONYM DBMS_LOGMNR FOR
SYS.DBMS_LOGMNR;
```

Setting Up an Oracle XStream JDBC Driver

To use the Replication Agent for Oracle with XStream APIs, use the **setenv** function to set or modify ORACLE_HOME, LD_LIBRARY_PATH, and CLASSPATH environment variables.

Prerequisites

- Download and install the appropriate XStream JDBC driver from the Oracle Web site for your primary Oracle database environment in which Replication Agent resides.
- Make sure your Java Runtime Environment (JRE) is compatible with your Oracle database for using XStream APIs. For example, if you have installed 64-bit Oracle, use the 64-bit JRE 1.7.x. See the Oracle documentation.

By default, Replication Agent uses 64bit JRE when creating an instance with the XStream APIs Log Reader component.

Task

1. Set the ORACLE_HOME environment variable to the Oracle installation directory for your platform:

```
setenv $ORACLE_HOME <install_dir>/oracle/<version>/
```

where *install_dir* is the installation directory and *version* is the version number of the Oracle database.

The Oracle installation directory might be different, based on the database version installed on your platform. Set the ORACLE_HOME variable before setting the LD_LIBRARY_PATH and CLASSPATH environment variables.

2. Add the path of the Oracle XStream JDBC driver to the CLASSPATH environment variable for your platform.

Add only one version of a JDBC driver in the CLASSPATH; otherwise, Replication Agent will have problems connecting to the primary database.

On Linux:

```
setenv LD_LIBRARY_PATH $ORACLE_HOME/lib:$LD_LIBRARY_PATH
```

```
setenv CLASSPATH $ORACLE_HOME/jdbc/lib/ojdbc6.jar:$ORACLE_HOME/
rdbms/jlib/xstreams.jar:$CLASSPATH
```

On HP-UX64:

```
setenv SHLIB_PATH $ORACLE_HOME/lib:$SHLIB_PATH
```

Oracle-to-SAP HANA Database Replication Setup

```
setenv CLASSPATH $ORACLE_HOME/jdbc/lib/ojdbc6.jar:$ORACLE_HOME/  
rdbms/jlib/xstreams.jar:$CLASSPATH
```

On AIX 64:

```
setenv LIBPATH $ORACLE_HOME/lib:$LIBPATH
```

```
setenv CLASSPATH $ORACLE_HOME/jdbc/lib/ojdbc6.jar:$ORACLE_HOME/  
rdbms/jlib/xstreams.jar:$CLASSPATH
```

On SUNOS-i86pc:

```
setenv LD_LIBRARY_PATH $ORACLE_HOME/lib:$LD_LIBRARY_PATH
```

```
setenv CLASSPATH $ORACLE_HOME/jdbc/lib/ojdbc6.jar:$ORACLE_HOME/  
rdbms/jlib/xstreams.jar:$CLASSPATH
```

On SUNOS-64:

```
setenv LD_LIBRARY_PATH_64 $ORACLE_HOME/lib:$LD_LIBRARY_PATH_64
```

```
setenv CLASSPATH $ORACLE_HOME/jdbc/lib/ojdbc6.jar:$ORACLE_HOME/  
rdbms/jlib/xstreams.jar:$CLASSPATH
```

where:

- *Oracle_Home* – is the Oracle database installation directory.
- *path_name* – is the library path where you installed the JDBC driver:
 - *LD_LIBRARY_PATH* on UNIX
 - *SHLIB_PATH* on HP-UX64
 - *LIBPATH* on AIX
- *driver* – is the name of the JDBC driver; for Oracle 11.2.0.3, *ojdbc6.jar*. The *xstreams.jar* is the name of the XStream library path.

3. (On Windows) Set Oracle environment variables in any of these ways:

Use the Control Panel	<ul style="list-style-type: none">• Select Start > Settings > Control Panel > System > Advanced System Settings > Environment Variables.• Click New in the User variables pane.• In the New User Variable dialog, enter:<ul style="list-style-type: none">• Variable name – is the environment variable that you want to set. For example, <i>ORACLE_HOME</i>, <i>PATH</i>, and <i>CLASSPATH</i>.• Variable value – is the directory path that you want to set for the environment variable. For example, <i>c:\oracle\version</i> (installation directory), <i>C:\ORACLE_HOME\bin</i> (the library path), and <i>C:\ORACLE_HOME\jdbc\lib\ojdbc6.jar</i> (the JDBC driver path).
------------------------------	--

Edit the existing variables	<ul style="list-style-type: none"> • Select the PATH and CLASSPATH environment variables in the System variables pane, and click Edit. • In the Edit System Variable dialog, using the semicolon (;) path separator, enter the variable value for: <ul style="list-style-type: none"> • Path – %ORACLE_HOME%\bin;%ORACLE_HOME%\LIB;%path% • CLASSPATH – %ORACLE_HOME%\jdbc\lib\ojdbc6.jar;%ORACLE_HOME%\rdbms\jlib\xstreams.jar;%CLASSPATH%
In the command line prompt	Enter: <pre>set CLASSPATH=ORACLE_HOME\jdbc\lib \ojdbc6.jar;ORACLE_HOME\rdbms\jlib\xstreams.jar</pre>

Creating Oracle User and Grant Permissions

Create an Oracle database user for use by Replication Agent when connected to the primary database.

Note: Permission to grant access to objects owned by “SYS” may require the command to be executed by an Oracle user with **sysdba** privileges.

1. Use SQLPLUS to connect to the primary database as a system administrator and run these commands to create an Oracle user named “RA_USER,” specified by **pds_username**, with the password “sybase,” and grant permissions to the user:

```
CREATE USER RA_USER PROFILE "DEFAULT" IDENTIFIED
BY "sybase" DEFAULT TABLESPACE "USERS" ACCOUNT
UNLOCK;
GRANT ALTER ANY PROCEDURE TO RA_USER;
GRANT ALTER ANY TABLE TO RA_USER;

GRANT ALTER SESSION TO RA_USER;
GRANT ALTER SYSTEM TO RA_USER;
GRANT CONNECT TO RA_USER;
GRANT CREATE PROCEDURE TO RA_USER;
GRANT CREATE PUBLIC SYNONYM TO RA_USER;
GRANT CREATE SESSION TO RA_USER;
GRANT CREATE TABLE TO RA_USER;
GRANT DROP ANY PROCEDURE TO RA_USER;
GRANT DROP ANY TABLE TO RA_USER;
GRANT DROP PUBLIC SYNONYM TO RA_USER;
GRANT EXECUTE CATALOG_ROLE TO RA_USER;
GRANT EXECUTE ON SYS.DBMS_FLASHBACK TO RA_USER;
GRANT EXECUTE ON SYS.DBMS_LOCK TO RA_USER;
GRANT RESOURCE TO RA_USER;
GRANT SELECT ANY TRANSACTION TO RA_USER;
GRANT SELECT_CATALOG_ROLE TO RA_USER;
GRANT SELECT ON SYS.ARGUMENT$ TO RA_USER;
GRANT SELECT ON SYS.ATTRIBUTE$ TO RA_USER;
GRANT SELECT ON SYS.CCOL$ TO RA_USER;
```

Oracle-to-SAP HANA Database Replication Setup

```
GRANT SELECT ON SYS.CDEF$ TO RA_USER;
GRANT SELECT ON SYS.COL$ TO RA_USER;
GRANT SELECT ON SYS.COLLECTION$ TO RA_USER;
GRANT SELECT ON SYS.COLTYPE$ TO RA_USER;
GRANT SELECT ON SYS.CON$ TO RA_USER;
GRANT SELECT ON SYS.DEFERRED_STG$ TO RA_USER;
GRANT SELECT ON SYS.ICOL$ TO RA_USER;
GRANT SELECT ON SYS.IND$ TO RA_USER;
GRANT SELECT ON SYS.INDCOMPART$ TO RA_USER;
GRANT SELECT ON SYS.INDPART$ TO RA_USER;
GRANT SELECT ON SYS.INDSUBPART$ TO RA_USER;
GRANT SELECT ON SYS.LOB$ TO RA_USER;
GRANT SELECT ON SYS.LOBCOMPPART$ TO RA_USER;
GRANT SELECT ON SYS.LOBFRA$ TO RA_USER;
GRANT SELECT ON SYS.MLOG$ TO RA_USER;
GRANT SELECT ON SYS.NTAB$ TO RA_USER;
GRANT SELECT ON SYS.OBJ$ TO RA_USER;
GRANT SELECT ON SYS.OPQTYPE$ TO RA_USER;
GRANT SELECT ON SYS.PARTOBJ$ TO RA_USER;
GRANT SELECT ON SYS.PROCEDUREINFO$ TO RA_USER;
GRANT SELECT ON SYS.RECYCLEBIN$ TO RA_USER;
GRANT SELECT ON SYS.SEG$ TO RA_USER;
GRANT SELECT ON SYS.SEQ$ TO RA_USER;
GRANT SELECT ON SYS.SNAP$ TO RA_USER;
GRANT SELECT ON SYS.TAB$ TO RA_USER;
GRANT SELECT ON SYS.TABCOMPART$ TO RA_USER;
GRANT SELECT ON SYS.TABPART$ TO RA_USER;
GRANT SELECT ON SYS.TABSUBPART$ TO RA_USER;
GRANT SELECT ON SYS.TS$ TO RA_USER;
GRANT SELECT ON SYS.TYPE$ TO RA_USER;
GRANT SELECT ON SYS.USER$ TO RA_USER;
GRANT SELECT ON V_$LOGMNR_CONTENTS TO RA_USER;
GRANT SELECT ON V_$LOGMNR_LOGS TO RA_USER;
```

For Oracle 12c:

```
GRANT LOGMINING TO RA_USER
```

The Replication Agent for Oracle Log Reader that uses XStream APIs to read from Oracle 11g and 12c transaction logs requires

DBMS_XSTREAM_AUTH.GRANT_ADMIN_PRIVILEGE.

To grant the XStream administrator privilege to the **pds_username**, issue:

```
BEGIN
DBMS_XSTREAM_AUTH.GRANT_ADMIN_PRIVILEGE( grantee =>
'pds_username' );
END;
/
```

See the *Replication Agent Primary Database Guide*.

2. If you configure the **ra_admin_owner** user, make sure that these permissions are also granted:

```
GRANT CREATE SESSION TO RA_USER;
GRANT CREATE TABLE TO RA_USER;
```

```
GRANT CREATE SEQUENCE TO RA_USER;
GRANT CREATE ANY PROCEDURE TO RA_USER;
GRANT SELECT_CATALOG_ROLE TO RA_USER;
```

3. If you intend to replicate Oracle partitioned tables or partitioned large objects (LOBs), grant select privileges to the user ID specified by **pds_username** on certain system tables.

See the *Replication Agent Primary Database Guide*.

Verifying the Oracle User Roles Created by the Replication Agent

Validate the required Oracle settings.

1. Use SQLPLUS to run this command as the new Oracle user:

```
select GRANTED_ROLE from USER_ROLE_PRIVS;
GRANTED_ROLE
-----
CONNECT
RESOURCE
SELECT_CATALOG_ROLE
```

The user who starts the Replication Agent instance must also have read access to the Oracle redo log files and the Oracle archive directory that contains the archive log files to be accessed for replication. If the Replication Agent is configured to remove old archive files, the user must have update authority to the directory and the archive log files.

2. Add the JDBC driver to the CLASSPATH environment variable. For example, you can append the existing CLASSPATH to the new JDBC driver and add this command to the .login script of a UNIX and Linux user:

```
setenv CLASSPATH /path_name/driver.jar:$CLASSPATH
```

where *driver.jar* is your JDBC driver (ojdbc14.jar, ojdbc5.jar, ojdbc6.jar, or ojdbc7.jar).

For Windows, add this to the CLASSPATH environment variable:

```
set CLASSPATH=path_name\jdbc\lib\driver.jar;%CLASSPATH%
```

where *driver.jar* is your JDBC driver (ojdbc14.jar, ojdbc5.jar, ojdbc6.jar, or ojdbc7.jar).

Creating the Replication Agent Instance

A single installation of the Replication Agent can support replication from multiple databases, however, one Replication Agent instance is needed for each Oracle database that is to be replicated.

Prepare and create a Replication Agent instance for each Oracle database being replicated.

1. Obtain the Oracle JDBC driver from Oracle, and update the CLASSPATH environment variable.

To find the JDBC connection that communicates with the primary Oracle instance, locate the JDBC driver that was available when the Oracle release was originally created.

Oracle-to-SAP HANA Database Replication Setup

- For Oracle 12c Release 1:
`$ORACLE_HOME/jdbc/lib/ojdbc7.jar`
- For Oracle 11g Release 2:
`$ORACLE_HOME/jdbc/lib/ojdbc6.jar`

Note: Install the `xstreams.jar` for Oracle XStream JDBC driver in the `CLASSPATH` environment variable.

- For Oracle 11g Release 1:
`$ORACLE_HOME/jdbc/lib/ojdbc5.jar`
- For Oracle 10g:
`$ORACLE_HOME/jdbc/lib/ojdbc14.jar`

Replication Agent may require a newer version of the JDBC driver, which can be found at <http://www.oracle.com/technetwork/indexes/downloads/index.html>.

Replication Agent may require a newer version of the JDBC driver, which can be found on the Oracle Web site.

2. Add the JDBC driver to the `CLASSPATH` environment variable. For example, you can append the existing `CLASSPATH` to the new JDBC driver and add this command to the `.login` script of a UNIX and Linux user:

```
setenv CLASSPATH /path_name/driver.jar:$CLASSPATH
```

where *driver.jar* is your JDBC driver (`ojdbc14.jar`, `ojdbc5.jar`, or `ojdbc6.jar`).

For Windows, add this to the `CLASSPATH` environment variable:

```
set CLASSPATH=path_name\jdbc\lib\driver.jar;%CLASSPATH%
```

where *driver.jar* is your JDBC driver (`ojdbc14.jar`, `ojdbc5.jar`, or `ojdbc6.jar`).

3. Record the location of the `tnsnames.ora` file, and record the connection name from that file for the Replication Agent to use to connect to the Oracle primary database. You can find these values from any machine from which you can use an Oracle `SQLPLUS` session to connect to the desired Oracle instance. The `ORACLE_SID` is from the environment variable, `$ORACLE_SID`.

Record the `ORACLE_SID` value for the instance you want to connect to.

Obtain host and port information from the file called `tnsnames.ora`, available at `$ORACLE_HOME/network/admin` (UNIX or Linux) or `%ORACLE_HOME%\network\admin` (Windows).

Record the:

- Host name of the Oracle database on which the TNS listener is operating
- Port number the TNS listener is monitoring

- ORACLE_SID value for the instance you want to connect to
4. (Optional) If your operating system has a Java 6.0.x-compatible JRE or JDK installed, you can use the Oracle **iSQL** (the browser-based interface to SQL*PLUS) demo items to connect to Oracle using the JDBC driver and to verify the connection information.
See the Oracle JDBC documentation.
 5. Obtain a local copy of the Oracle `timezone` file, so Replication Agent can correctly process the Oracle `timestamp with timezone` datatype.

Note: This step is required only if the Replication Agent is on a machine (host) other than the Oracle host and does not have access to the Oracle `timezone` file.

There are two `timezone` files in the `$ORACLE_HOME/oracore/zoneinfo` (UNIX or Linux) or `%ORACLE_HOME%\oracore\zoneinfo` (Windows) directory:

- The `timezone*.dat` binary file is the default file that contains the most commonly used time zones.
- The `timezlg*.dat` binary file contains a larger set of defined time zones.

The Oracle instance uses the `timezone` file from the `ORA_TZFILE` environment variable. If `ORA_TZFILE` is not set, use the default `timezone` file.

Make a copy of the `timezone` file available on the machine hosting the Replication Agent. To determine which `timezone` file is used by the primary Oracle instance, connect to the Oracle instance and execute:

```
select * from V$TIMEZONE_FILE
```

Note: These files are Oracle version- and platform-dependent. You cannot use a `timezone` file from a little-endian platform on a big-endian platform, nor can you use a `timezone` file from a version of Oracle that is different than the version of the primary Oracle.

6. Locate the Replication Agent for Oracle resource file template.

You can record, and store in a resource file, most of the configuration values required to create and initialize a Replication Agent. Using a resource file allows an instance to be easily removed and re-created.

The resource file template (`oracle.rs`) is in the Replication Agent directory `$SYBASE/RAX-15_5/init` (UNIX or Linux) or `%SYBASE%\RAX-15_5\init` (Windows).

The resource file template (`oracle.rs`) is in the Replication Agent directory `$SYBASE/RAX-15_2/init` (UNIX or Linux) or `%SYBASE%\RAX-15_2\init` (Windows).

7. Create an instance resource file:

Oracle-to-SAP HANA Database Replication Setup

Copy the resource file template `$SYBASE/RAX-15_5/init/oracle.rs` (UNIX or Linux) or `%SYBASE%\RAX-15_5\init\oracle.rs` (Windows) to another file that contains the configuration information for a new instance, for example, `myra.rs`.

Copy the resource file template `$SYBASE/RAX-15_5/init/rao_oracle.rs` (UNIX or Linux) or `%SYBASE%\RAX-15_5\init\rao_oracle.rs` (Windows) to another file that contains the configuration information for a new instance, for example, `myra.rs`.

Copy the resource file template `$SYBASE/RAX-15_2/init/oracle.rs` (UNIX or Linux) or `%SYBASE%\RAX-15_2\init\oracle.rs` (Windows) to another file that contains the configuration information for a new instance, for example, `myra.rs`.

Change the values that are assigned to properties in the resource file to match the values required for your environment. Most of these values are the location, user, and password information listed in the `tnsname.ora` file for the Oracle, Replication Server, and RSSD connections.

- Set the value of **rs_source_ds** to the value you chose for *pds* when you created a Replication Server connection from the primary database.
- Set the value of **rs_source_db** to the value you chose for *pdb* when you created a Replication Server connection from the primary database.
- When you create a Replication Agent instance with a resource file:
 - The **asa_password** configuration parameter value cannot contain single quotes, double quotes, or a semicolon.
 - The **pds_username** and **pds_password** configuration parameter values cannot contain single or double quotes if the **create_pds_username** parameter is set to `yes`.

The resource file is self-documenting, with each parameter described. See the *Replication Agent Administration Guide*. For a complete list of configuration parameters, see the *Replication Agent Reference Manual*.

Warning! The **rs_source_ds** and the **rs_source_db** values must match the `{pds.pdb}` values of your Replication Server primary connection name that you configured when you created a Replication Server connection from the primary database.

Table 7. Resource File Parameters

Parameter	Description	Example Values
instance name	Any valid name.	myra
admin_port	Port number that Replication Agent uses.	9030 (if in use, select a different port number)
itm_admin_user	User name for administering the Replication Agent instance.	sa

Parameter	Description	Example Values
ltm_admin_pw	Password for the user administering the Replication Agent instance.	The default password policy is that the password must be between 6 and 255 characters. For example: sa_pass
pds_tns_connection	Connection name found in the <code>tnsnames.ora</code> file that identifies the connection information for the primary database.	ORA102.JDOE_HOST.COM
pds_tns_filename	File name identifying the Oracle <code>tnsnames.ora</code> file to be used to identify the connection information for the primary database.	A valid Oracle <code>tnsnames.ora</code> file. For example: <code>/opt/oracle/network/admin/tnsnames.ora</code>
pds_username	User ID that Replication Agent uses to access primary data server.	RA_USER
pds_password	Password for pds_username .	sybase
rs_host_name	Machine where Replication Server is installed.	jdoe_host1
rs_port_number	Port number where Replication Server is installed.	11752
rs_username	Replication Server user with connect source and create object permissions. This is the user used to create a Replication Server connection from the primary database.	SAMPLE_RS_ra
rs_password	Password for rs_username .	sa_pass
rs_charset	Character set that Replication Server is using. The value defined for the rs_charset configuration parameter must match the RS_charset value in the Replication Server configuration file, <code>\$SYBASE/REP-15_5/install/<server>.cfg</code> (UNIX or Linux) or <code>%SYBASE%\REP-15_5\install\<server>.cfg</code> (Windows).	<ul style="list-style-type: none"> Windows: cp850 UNIX: iso_1 For a sample Replication Server instance, the rs_charset parameter must be set to utf8.

Parameter	Description	Example Values
rs_source_ds	Valid name representing the data server of the Oracle primary database. This value is used to create a Replication Server connection from the primary database. If you use direct load materialization, add an entry with this name in the interfaces file.	pds
rs_source_db	Valid name representing the Oracle primary database. This value is used to create a Replication Server connection from the primary database.	pdb
rssd_host_name	Machine where RSSD resides.	jdoe_host1
rssd_port_number	Port number where RSSD resides.	11751
rssd_database_name	Database name for RSSD.	SAMPLE_RS_ERSSD
rssd_username	Valid user for RSSD.	SAMPLE_RS_RSSD_maint
rssd_password	Password for rssd_username .	sa_pass
pdb_timezone_file	Path of the Oracle timezone file.	/software/oracle/Oral0g/oracore/zoneinfo/timezone.dat
start_instance	Start the instance that was created.	no
initialize_instance	Initialize the Replication Agent instance.	no
pdb_include_archives	Enables or disables the use of Oracle archive log files.	USE_DEFAULT
pdb_archive_path	Identifies directory path where Replication Agent expects to find archived Oracle redo log files.	A valid directory path on the machine hosting Replication Agent that points to a location where Oracle puts the archived redo log files.

8. Create and execute the new instance resource file.

Note: Be sure your CLASSPATH environment variable points to the correct Oracle JDBC driver before proceeding.

a) Validate the settings in the resource file:

- On Windows:

```
%SYBASE%\RAX-15_5\bin\ra_admin.bat -vr myra.rs
```

```
%SYBASE%\RAX-15_5\bin\rao_admin.bat -vr myra.rs
```

```
%SYBASE%\RAX-15_2\bin\ra_admin.bat -vr myra.rs
```

- On UNIX or Linux:

```
$SYBASE/RAX-15_5/bin/ra_admin.sh -vr myra.rs
```

```
$SYBASE/RAX-15_5/bin/rao_admin.sh -vr myra.rs
```

```
$SYBASE/RAX-15_2/bin/ra_admin.sh -vr myra.rs
```

Note: Make sure permission for the `ra_admin.sh` file is set no broader than 700 (read/write/execute for user, no permissions for group and other).

where *myra.rs* is the path and name of the resource file. Validation results are returned as either:

- Response-file processing completed
- or
- Response-file processing completed with errors

If any validation fails, the `ra_admin` utility returns an error message and information about the failure. You can repeat the validation process as many times as necessary until it executes without error. No entities are changed or created.

- b) Create and configure the Replication Agent instance:

- On Windows:

```
%SYBASE%\RAX-15_5\bin\ra_admin.bat -r myra.rs
```

```
%SYBASE%\RAX-15_5\bin\rao_admin.bat -r myra.rs
```

```
%SYBASE%\RAX-15_2\bin\ra_admin.bat -r myra.rs
```

- On UNIX or Linux:

```
$SYBASE/RAX-15_5/bin/ra_admin.sh -r myra.rs
```

```
$SYBASE/RAX-15_5/bin/rao_admin.sh -r myra.rs
```

```
$SYBASE/RAX-15_2/bin/ra_admin.sh -r myra.rs
```

Note: Make sure permission for the `ra_admin.sh` file is set no broader than 700 (read/write/execute for user, no permissions for group and other).

where *myra.rs* is the path and name of the resource file.

- c) Results are returned as either:

- Response-file processing completed
- Response-file processing completed with errors

See the *Replication Agent Administration Guide*.

9. If you set **start_instance** to yes in your resource file, and if the log indicates that the Replication Agent instance was started, skip this step, and continue with the next section.

Change to the instance directory and run Replication Agent in the background.

- On Windows:

```
cd %SYBASE%\RAX-15_5\myra
```

Oracle-to-SAP HANA Database Replication Setup

```
cd %SYBASE%\RAX-15_2\myra
```

- On UNIX or Linux:

```
cd $SYBASE/RAX-15_5/myra
```

```
cd $SYBASE/RAX-15_2/myra
```

On UNIX or Linux, execute the RUN file in the background:

```
bg pid
```

where *pid* is the process ID.

Your Replication Agent for Oracle is now running.

See also

- *Setting Up an Oracle XStream JDBC Driver* on page 45

Verifying the Replication Agent Instance Installation

Verify that the Replication Agent instance has been successfully installed.

1. Use **dsedit** to update the Replication Server `sql.ini` (Windows) or `interfaces` (UNIX or Linux) file to include an entry for the Replication Agent location.

Note: You can use any Tabular Data Stream™ (TDS) client utility (**isql**, **isqlApp**, or **SQLAdvantage**) that you prefer.

If you use direct load materialization, also add the server name part of the connection name to the `interfaces` file. This entry should point to Replication Agent. For example:

```
pds
master tcp ether hostname 9030
query tcp ether hostname 9030
```

Note: The server name is the same value as the **rs_source_ds** value in the resource file used to create a Replication Agent instance.

2. (Optional) Verify the connection to Replication Agent:
 - a) Open a command window in the `%SYBASE%` (Windows) or `$SYBASE` (UNIX or Linux) directory of your Replication Agent installation.
 - b) Set the environment variables by executing the `SYBASE.bat` file (Windows) or by sourcing the `SYBASE.csh` file (UNIX or Linux).
 - c) Log in to Replication Agent:

```
isql -Usa -Psa_pass -Smyra
```

Note: These verification steps are optional because they were performed when you verified the resource file.

```
isql -Username -Ppassword -Smyra
```

3. Verify the Replication Agent connection to Replication Server:

a) Enter:

```
test_connection RS
go
```

b) This result is returned:

```
Type Connection
-----
RS      succeeded
(1 row affected)
```

c) If the result indicates a failure, either the server is not responding or the connection properties (*host*, *port*, *user*, or *password*) are incorrect. Verify the host and port configuration values, and manually log in to the Replication Server as the configured user to determine which property is incorrectly defined.

See the *Replication Agent Administration Guide*.

4. Verify the Replication Agent connection to the primary database:

a) Enter:

```
test_connection PDS
go
```

You see:

```
Type Connection
-----
PDS      succeeded
(1 row affected)
```

b) If the result indicates a failure:

- The server is not responding, or
- One of the connection properties is incorrect.

Check the `tnsnames.ora` file and `tnsnames` configuration values, and manually log in to the primary Oracle database as the configured user to find which property is incorrectly defined.

See *Testing Network Connectivity* in the *Replication Agent Administration Guide*.

See also

- *Initializing the Replication Agent Instance* on page 57

Initializing the Replication Agent Instance

Initialize the Replication Agent instance.

Note: If you set `initialize_instance` to yes in your resource file, and if you received no errors while the instance was created, skip to step 2.

1. The `ra_admin init` command verifies that the primary Oracle database is correctly configured to provide archived logging and supplemental logging, and that the Oracle user ID used by the Replication Agent has the necessary permissions. In addition to creating

objects in the database to support stored procedure replication, the **ra_admin init** command also initializes the Replication Agent System Database (RASD) by reading schema information, and redo log location information from the primary Oracle database.

To initialize the Replication Agent instance, run:

```
ra_admin init
go
```

2. Enter:

```
resume
go
```

If the Replication Agent successfully transfers to a replicating state, you see:

State	Action
REPLICATING	Ready to replicate data.

The Replication Agent goes to the REPLICATION DOWN state if an error occurs:

```
isql -Usa -Psa_pass -Smyra
1> resume
2> go
Msg 32000, Level 20, State 0:
Server 'myra', Procedure 'resume', Line 1:
Command <resume> failed - Desired state <REPLICATE>
could not be achieved. Current state: <REPLICATION DOWN>
```

3. The **ra_status** command returns the state of the Replication Agent. It is a good practice to verify that the Replication Agent remains in replication state, even after the resume command executes successfully.

To detect an error that occurred after replication start-up, execute:

```
ra_status
go
```

If the Replication Agent is in replicating state, **ra_status** returns:

State	Action
REPLICATING	Ready to replicate data.

The Replication Agent goes to the REPLICATION DOWN state if an error occurs:

```
1> ra_status
2> go
State          Action
-----
REPLICATION DOWN  Waiting for operator command.
```

4. Validate that both primary and replicate connections are active:

```
isql -Usa -Psa_pass -SSAMPLE_RS
admin who
go
```



```
isql -Username -Ppassword -SSAMPLE_RS
admin who
go
```

Note: Be aware that the DSI connection for the primary database connection is usually suspended because you are not replicating data back to the primary database.

Do not proceed until **admin who** returns status for threads similar to this:

Spid	Name	State	Info
13	DSI EXEC	Awaiting Command	101(1) SAMPLE_RS_ERSSD.
9	DSI	Awaiting Message	101 SAMPLE_RS_ERSSD.
8	SQM	Awaiting Message	101:0 SAMPLE_RS_ERSSD.
54	DSI EXEC	Awaiting Command	102(1) rds.rdb
53	DSI	Awaiting Message	102 rds.rdb
17	SQM	Awaiting Message	102:0 rds.rdb
20	DSI EXEC	Suspended	103(1) pds.pdb
21	DSI	Suspended	103 pds.pdb
24	DIST	Awaiting Wakeup	103 pds.pdb
25	SQT	Awaiting Wakeup	103:1 DIST pds.pdb
23	SQM	Awaiting Message	103:1 pds.pdb
22	SQM	Awaiting Message	103:0 pds.pdb
62	REP AGENT	Awaiting Command	pds.pdb

Testing Replication

When you finish setting up the replication system, test that replication works as intended.

1. Connect to the primary Oracle instance as a regular user, rather than as the maintenance user.
2. At the primary Oracle database, create a table named `ptab1` to replicate:

```
CREATE TABLE ptab1
(idno NUMBER PRIMARY KEY,
name VARCHAR2(20));
```

3. Grant permission to the Replication Agent to select data from the `ptab1` table for the **direct_load** subscription.

```
grant all on ptab1 to public;
```

4. Connect to Replication Agent through **isql**, and mark the `ptab1` table for replication:

```
pdb_setreptable ptab1, mark
go
```

5. Add rows to the `ptab1` table:

```
insert into ptab1 values (1, 'Burt Gringlesby');
insert into ptab1 values (2, 'Douglas Wong');
insert into ptab1 values (3, 'Tom Servo');
insert into ptab1 values (4, 'Innes del Castillo');
```

```
insert into ptab1 values (5, 'Akiko Yokomoto');
commit;
```

- In Replication Agent, create a replication definition against the primary Oracle database:

```
rs_create_repdef ptab1
go
Table/Procedure Name      RepDef Name      Status
-----
<OracleTableOwner>.PTAB1  "<repdefname>"  Created
(1 row affected)
```

Record the name of the replication definition for use later.

- At the replicate SAP HANA database instance, create a table named PTAB1:

```
CREATE TABLE <tableowner>.PTAB1
(IDNO INT PRIMARY KEY,
NAME VARCHAR(20));
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

If an owner is not specified with the table name, the owner of the table—*<tableowner>*—is the user who is logged in when the table is created.

- Grant permissions to any new or existing object to be replicated in the replicate database so that the Replication Server maintenance user can update this table:

```
grant all privileges on <tableowner>.PTAB1 to public
```

- Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

- Alter the replication definition that you created earlier to point to the correct replicate table:

```
alter replication definition <repdefname>
with replicate table named <tableowner>.ptab1
go
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

- Create a subscription against the replicate connection to materialize the PTAB1 table on the SAP HANA database.

Create the ptab1_sub subscription with **direct_load** option for the *<repdefname>* replication definition. Replication Server logs in to Replication Agent for Oracle, and Replication Agent issues a **SELECT** command to select data from the Oracle database, enter:

```
create subscription ptab1_sub
for <repdefname>
```

```
with replicate at rds.rdb
without holdlock
direct_load
user puser password ppwd
go
```

where:

- *rds.rdb* is the replicate SAP HANA database connection created earlier.
- *puser* is the user who connects to Replication Agent for Oracle.
- *ppwd* is the password of the *puser*. You must enter a password.

Note: The interfaces file used by the replicate Replication Server must contain an entry for the Replication Agent for Oracle that is used to select data from the Oracle database.

To create the subscriptions at the replicate Replication Server using the no materialization method, use the **create subscription** command with the **without materialization** clause.

For information about using the **create subscription** command in **direct_load** materialization, see *Replication Server Reference Manual*.

12. Check the subscription status to ensure there are no errors:

```
check subscription ptabl_sub
for <repdefname>
with replicate at rds.rdb
go
```

If there are errors, see *Troubleshooting Subscription Problems*.

13. Connect to the Oracle database as a non-maintenance user with the ability to insert rows into the table:

```
insert into ptabl values (10, 'Michel DeFrance');
insert into ptabl values (11, 'Dirk Stringer');
commit;
```

14. On the Replication Server, check the status of the subscription created with the **direct_load** option by using the **check subscription** command. When the subscription is in the VALID state, materialization is complete:

```
check subscription ptabl_sub
for <repdefname>
with replicate at rds.rdb
go
Subscription ptabl_sub is VALID at the replicate.
```

15. After the status is VALID at the primary and replicate database, on the replicate database, verify that the data is replicated to the replicate PTAB1 table:

```
SELECT * FROM <tableowner>.PTAB1
```

Resetting the Primary Oracle Database for Replication

In a test environment, you may occasionally need to reset the replication environment.

1. Log in to the Replication Agent and suspend replication:

Oracle-to-SAP HANA Database Replication Setup

```
isql -Usa -Psa_pass -Smyra
suspend
go
```

2. Set the transaction log truncation point to the end of the current transaction log.

```
ra_locator move_truncpt
```

3. Refresh the Replication Agent repository rather than overwriting it:

```
ra_admin refresh
```

Note: If you prefer to delete and replace all the information in the Replication Agent repository, issue the **ra_admin deinit, force** command followed by a normal **ra_admin init** command (without the **force** option).

4. Reset the locator stored in Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
connect
go
rs_zeroltm pds, pdb
go
exit
go
```

You see:

```
Locator has been reset to zero.
(return status = 0)
```

5. Resume replication with the Replication Agent **resume** command:

```
isql -Usa -Psa_pass -Smyra
resume
go
```

Microsoft SQL Server-to-SAP HANA Database Replication Setup

Implement a sample replication system for replication from Microsoft SQL Server to SAP HANA database.

Use the tasks in this document to set up a proof-of-concept or test replication system.


The tasks use the sample Replication Server instance SAMPLE_RS. If you have already installed Replication Server, skip to the configuration procedures.

For the most current information about Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

Multiple components are required to implement an Microsoft SQL Server-to-SAP HANA database replication system: a primary Microsoft SQL Server data server, a replicate SAP HANA database, a Replication Server with ExpressConnect for HANA DB, and Replication Agents for Microsoft SQL Server.

Figure 3: Microsoft SQL Server-to-SAP HANA Database Replication Setup



Prerequisites

Review the prerequisites before you begin setting up Microsoft SQL Server-to-SAP HANA database replication.

- Familiarity with Microsoft SQL Server and SAP HANA database.
- Valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- Valid licenses for Replication Server and Replication Server Options.
- Valid licenses for the Replication Server and Replication Server Heterogeneous Edition.

Microsoft SQL Server-to-SAP HANA Database Replication Setup

- Install and configure Microsoft SQL Server 2008 or 2012 to serve as the primary (source) database from which Replication Agent replicates transactions, and SAP HANA database as the replicate (target) database.
- Install Microsoft SQL Server JDBC driver 4.0.
- TCP/IP connectivity between all hosts.

Because Replication Agent for Microsoft SQL Server is limited to the Microsoft Windows platform, all examples shown are for Windows.

Note: On Windows Vista and Windows Server 2008, you must be logged in as an Administrator.

Do not use Replication Server reserved words for object names or connection names. The *Replication Server Reference Manual* includes a complete list of reserved words.

Planning Your Installation

Before installing the replication components, verify the system requirements, identify the installation directories, and obtain the installation software.

Installation Checklist

Use the checklist to navigate through the installation when you set up your replication environment.

Table 8. Installation Checklist

	Action	Topic
1	Make sure the target computer meets the minimum memory and disk space requirements.	<i>System Requirements</i>
2	Determine the directory in which to install the replication components.	<i>Identify the Installation Directory</i>
3	Download replication components from the SAP Service Marketplace (SMP).	<i>Downloading Products from SMP</i>
4	Install Replication Server.	<i>Installing Replication Server</i>
5	Install Replication Agent.	<i>Installing Replication Agent</i>

System Requirements

Make sure the target computer on which you are installing the replication components meets the minimum memory and disk space requirements.

See the *Replication Server Installation Guide > Planning Your Installation > System Requirements* and the *Replication Agent Installation Guide > Planning Your Installation > System Requirements*.

To improve performance by avoiding multiple network hops, install Replication Server and the SAP HANA database on the same machine.

Note: The only versions of Linux that ExpressConnect for HANA DB supports are Red Hat Enterprise Linux 6 and SuSE Linux Enterprise Server (SLES) 11.

Identify the Installation Directory

Determine the directory in which to install the replication components.

The quick-start procedures described here assume you have installed the SAP software on the same host as the Microsoft SQL Server data server. Replication Agent, Replication Server, and ExpressConnect for HANA DB are installed in the same base directory, which is identified by the SYBASE environment variable. Select a path on a host drive to be the recipient of the installation and configuration activities. Replication Agent must be installed on a host from which it can directly access the primary Microsoft SQL Server instance transaction log.

Replication Agent Installation Directory

Replication Agent versions 15.5 and later use the RAX-15_5 directory name. If you have an existing Replication Agent installation of version 15.5 or later and do not want to overwrite it, install the latest version of Replication Agent in a different location.

Downloading Products from SMP

Download replication components from the SAP Service Marketplace (SMP).

Prerequisites

An SMP login ID and password.

Task

Note: If you are installing SAP® Replication Server® Options (RSO), there are several steps you must perform twice, because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/patches>, and enter your login ID and password.

2. Select **SAP Software Download Center > Installations and Upgrades > A-Z Index** and click **R** for SAP Replication Server products.
3. From the product listing, select:
 - **Sybase Replication Server, Option for HANA** or
 - **Sybase Replication Server, HANA Edition**
4. Select the software you need by version and platform.
5. Generate a license based on the license model—served or unserved—and license type:
 - Application Deployment CPU License (AC)
 - Application Deployment Other License (AO)
 - Application Deployment Standby CPU License (BC)
 - CPU License (CP)
 - Development and Testing License (DT)
 - Other License (OT)
 - Standby CPU License (SF)
 - Server License (SR)
 - Standalone Seat License (SS)

Note: To generate a license, you must provide some information, such as the host name, MAC address, and number of CPUs.

6. Download your product licenses, and place them in your product license directory after installation:
 - `$SYBASE/SYSAM-2_0/licenses` (UNIX or Linux)
 - `%SYBASE%\SYSAM-2_0\licenses` (Windows)where `$SYBASE` or `%SYBASE%` is where you installed your product.
7. Download your installation files, and uncompress and extract all the installation images to your local drive.

Installing Replication Components

Review the procedures for installing the replication components.

Install all replication components on the same host where a Microsoft SQL Server data server has already been installed and is running with the primary database. This simplifies the quick-start process.

Installing Replication Server

Install Replication Server using the setup program.

Prerequisites

- Allocate a disk partition of at least 20MB for each Replication Server you are installing. You can add more partitions later, if necessary. Make sure each partition is available and has write permissions.
- Allocate the entire partition to the Replication Server. If you allocate only a portion of the partition to Replication Server, you cannot use the remainder for any other purpose.
- The sample Replication Server instance `SAMPLE_RS` used in these procedures must be created when you install Replication Server. If you have already installed Replication Server, see the instructions for how to set up a Replication Server instance after installation in *Replication Server ASE-to-ASE Replication Quick Start Guide > SAMPLE_RS*.

Task

1. Verify that the drive on which you install the Replication Server has enough disk space for the components being installed, and at least 100MB of extra disk space for the installation program.
2. Download and extract the Replication Server installation image from the SAP® Service Marketplace (SMP).
3. Start the installation.

- On Windows, launch the **setup** program.

If the installer does not start automatically, double-click **setup.exe** or select **Start > Run** and enter:

```
setup.exe
```

If there is not enough disk space in the temporary disk space directory, set the environment variable `TMP` to `directory_name` before running the installer again, where `directory_name` is where the installation program writes the temporary installation files. When specifying the `directory_name`, include its full path.

- (UNIX or Linux) If you downloaded the product from SMP, go to the directory where you extracted the installation image and start the installer:

```
./setup.bin
```

where `setup.bin` is the executable file name for installing Replication Server.

If there is not enough disk space in the temporary disk space directory, set the environment variable `IATEMPDIR` to `tmp_dir` before running the installer again, where `tmp_dir` is where the installation program writes the temporary installation files. When specifying the `tmp_dir`, include its full path.

4. In the Introduction window, click **Next**.
5. Specify where to install Replication Server:

Option	Description
Click Choose .	To select an installation directory. A window appears, and from that window browse for and select the directory.
Enter a new directory path.	To create a new directory.
Click Restore Default Folder .	To restore to the default directory if you do not want to use the directory that you have entered.

- If the directory you chose does not exist, click **Yes** to create it.
 - If it does exist, and already contains a Replication Server installation, you are warned that you will overwrite the older version. Click **Next**.
6. Select the type of installation:
 - **Typical** – installs the default components. This is recommended for most users.
 - **Full** – installs every Replication Server component, including all the supported language modules.
 - **Custom** – lets you select the components to install. Some components are automatically installed if they are required to run your selected components.

Click **Next**.

7. Select the geographic location, agree to the license agreement, then click **Next**.
8. On the SySAM License Server window, select one of:
 - **Specify License Keys** – click **Browse** to select the license file. To select multiple license files, use **Shift+Click** or **Ctrl+Click**. The license pane displays the license information.
Alternatively, copy and paste the license information directly in the license pane. Click **Next**.
If you specify a served license key, the installer prompts you to install a new SySAM license server. Click:
 - **Next** to install the new SySAM license server and follow the installation prompts.
 - **Previous** and select **Use previously deployed license server** if you have an existing SySAM license server on the same host.
 - **Use previously deployed license server** – for an existing license server, enter the host name of the machine where the license server is running, and the port number if the port number you are using is not the default. Click **Next**.
 - **Continue installation without a license key** – click **Next** to proceed without any license keys. The installer allows you to install and use the Replication Server components without a license for 30 days. To continue using these components after

the end of the grace period, obtain valid licenses and install these licenses using the Replication Server License Installer.

See the *SySAM Users Guide*.

9. Configuring your server for e-mail notification enables designated users to receive information about license management events requiring attention.

Provide:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipient e-mail addresses
- Message severity that triggers e-mail messages

Click **Next**.

10. The installation summary window displays the selections you have made. Review the information, and click **Install**.

11. On the Start Sample Replication Server window, select:

- **Yes** to configure and start a sample Replication Server. The installer displays the configuration information for the sample Replication Server. Record this information. You must enter and confirm a password from 6 to 15 characters in length.

Note: You cannot start the sample Replication Server on Linux on POWER because the ERSSD requires SAP SQL Anywhere Server, which is not available on Linux on POWER.

- **No** to manually configure a full-featured Replication Server and start a sample Replication Server after installation. The installer prompts you to either create a sample Replication Server directory or to proceed with the installation. Select either:
 - **Yes** – to create a sample Replication Server directory without starting the sample Replication Server, or,
 - **No** – to continue with the installation.

See the *Replication Server Configuration Guide* for Windows or UNIX.

Note: For a sample Replication Server instance, the **rs_charset** parameter must be set to utf8.

Click **Next**.

12. When the final window appears, indicating a successful installation, click **Done** to exit the installer.

Installing Replication Agent

Install Replication Agent using the GUI wizard.

Note: If there is not enough disk space in your default temporary directory, set the TMP Windows environment variable to a directory that has enough space. Include the full path to this directory.

1. Log in to the Replication Agent host machine using an operating system user account with authority to start, stop, and administer the Replication Agent instance (for example, the “sybase” user).
2. Close all nonessential applications, and minimize any open windows.
3. Download and extract the Replication Agent installation image from the SAP Service Marketplace (SMP).
4. Start the installation program. On Microsoft Windows platforms, the installation program should start automatically. If it does not, double-click the **setup.exe** file.
5. Click **Next**.
6. Select the geographic location, agree to the license agreement, then click **Next**.
7. Specify an installation directory. The default installation directory is the existing `%SYBASE%` or `c:\sybase`.

Click **Next** to accept the default installation directory, or:

- Click **Choose** to select an installation directory in the file browser. Click **Next**.
- Enter a directory name. Click **Next**.

8. Select the type of database the Replication Agent will connect to.

Click **Next**.

9. Select one of these options to enter the license:

- **Specify license keys** – browse to or specify the license file.
- **Use previously deployed license server** – use a previously deployed license server. Enter the host name of the machine where the license server is running and the port number if the port number you are using is not the default.
- **Continue installation without a license key** – install and use Replication Agent without a license for a grace period of 30 days. To continue using Replication Agent after the end of the grace period, obtain a valid license from the SAP Service Marketplace (SMP), and install it.

Note: Replication Agent is licensed statically when installed with RSHE. If you are installing Replication Agent with RSHE, locate the file that contains a valid SySAM license for the Replication Agent component in RSHE. Specify that license file [here](#).

See the *SySAM Users Guide*.

Click **Next** until you see the SySAM Notification window.

10. On the SySAM Notification window, configure your server for e-mail notification. When configuration is enabled, you receive information about license management events that require attention. Select **Yes** and either accept the default values that are supplied, or enter values for:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipients of the notification
- Message severity level of an event that triggers e-mail notification:
 - Informational
 - Warning
 - Error

If you choose not to have e-mail alerts or severity messages logged, select **No**.

Click **Next**.

11. Review the product features or components listed on the Preinstallation Summary window. Click **Install**.

12. Click **Done**.

Verifying the Installation

Verify that Replication Agent has been successfully installed.

1. In a command window, change to the directory where you installed Replication Agent.
2. Change to %SYBASE%\RAX-15_5\bin.
3. Obtain the Replication Agent version string:

```
ra.bat -v
```

If Replication Agent installs successfully, you see the SAP copyright and the Replication Agent version string.

Configuring Replication Components

Review the procedures for configuring each of the replication components in the replication system.

Configuration Checklist

Use the table as a checklist to navigate through the configuration when you set up your replication environment.

Table 9. Configuration Checklist

	Action	User	Topic
1	Replication Server: Create a Replication Server connection to the primary database.	Maintenance user	<i>Creating a Replication Server Connection to the Primary Database</i>
2	Replication Server: Set up ExpressConnect for HANA DB.		<i>Obtaining the SAP HANA Database ODBC Drivers</i>
3	Replication Server: Create a Replication Server connection to SAP HANA database.	Maintenance user	<i>Configuring Replication Server for Replication to SAP HANA Database</i>
4	Replication Server: Configure Replication Server for replication from the primary database.	rs_username	<i>Configuring Replication Server for Replication from the Primary Database</i>
5	Replication Agent: Configure and verify the primary Microsoft SQL Server database for replication.		<i>Primary Microsoft SQL Server Configuration and Verification</i>
6	Replication Agent: Configure and validate the Replication Agent resource file.		<i>Configuring and Validating the Replication Agent Resource File</i>
7	Replication Agent: Initialize the primary data server.		<i>First-Time Initialization</i>
8	Replication Agent: Create another Replication Agent instance, after you have already initialized the primary data server.		<i>Subsequent Initialization</i>

	Action	User	Topic
9	Verify the replicating state of Replication Agent.		<i>Verifying the Replication Agent Replicating State</i>
10	Test that replication works as intended.		<i>Testing Replication</i>

See also

- *User Management* on page 137

Configuring Replication Server

Configure Replication Server for Microsoft SQL Server-to-SAP HANA database replication.

Creating a Replication Server Connection to the Primary Database

Create a Replication Server connection to the primary database.

1. In Replication Server, create a connection to the primary Microsoft SQL Server database:

```
create connection to pds.pdb
using profile rs_rs_to_msss_ra;standard
set username muser
set password mpwd
with log transfer on, dsi_suspended
go
```

where:

- *pds* is the value of the **rs_source_ds** parameter specified in Replication Agent.
- *pdb* is the value of **rs_source_db** specified in Replication Agent.
- *muser* is the maintenance user for the primary Microsoft SQL Server database.
- *mpwd* is the maintenance user password.

For information about the maintenance user, see the *Replication Server Heterogeneous Replication Guide > SAP Replication Products > Replication Server > Database Connections > Maintenance User Purpose*.

2. Use **admin show_connections, 'primary'** to display the primary connection you created.

Obtaining the SAP HANA Database ODBC Drivers

Set up ExpressConnect for HANA DB for connections between SAP HANA database and Replication Server.

You must add the SAP HANA database ODBC driver to the ExpressConnect for HANA DB installation. Download these libraries from the SAP Service Marketplace and install them after you have installed Replication Server. See the *Replication Server Release Bulletin > Special Installation Instructions > Installing ODBC Libraries for ExpressConnect for HANA DB*.

Configuring Replication Server for Replication to SAP HANA Database

Create a Replication Server connection to SAP HANA database.

If you are not using the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

1. If you are connecting to SAP HANA database with a standard connection and not using SAP Secure User Store, add an entry to your Replication Server interfaces file identifying the replicate SAP HANA database, and then stop and start Replication Server:

```
[dataservername]
master tcp ether hostname port
query tcp ether hostname port
```

where *hostname* and *port* are the host and port number of the SAP HANA database, and *dataservername* is a label used to identify the host and port number. For information about stopping and starting Replication Server, see the *Replication Server Administration Guide Volume 1 > Manage a Replication System*.

Note: Each SAP HANA database installs with a unique instance number, and the port number is the instance number prefixed with 3 and suffixed with 15:

```
3in15
```

where *in* is the two-digit instance number. For example, the port number for an SAP HANA database with instance number 1 is 30115.

If you are using SAP Secure User Store, create a user store of encrypted credentials:

```
hdbuserstore set rds myhost:xxxxx my_securestore_user
my_securestore_pwd
```

where:

- *rds* is the key for the secure store entry.
- *myhost.xxxxx* is the connection environment host name and port number.
- *my_securestore_user* and *my_securestore_pwd* are SAP Secure User Store credentials.

2. On SAP HANA database, create a maintenance user:
 - a) As the SAP HANA database admin user, create a user ID with a temporary password:

```
CREATE USER muser PASSWORD mpwd_temp
```

- b) Log in to SAP HANA database with the temporary password, and then update the password:

```
ALTER USER muser PASSWORD mpwd_new
```

- c) Log out of SAP HANA database, and log in again with the new password.

3. Log in to SAP HANA database:

```
./hdbsql -u user -p password -i id
```


where *user* is the SAP HANA database user, *password* is the user password, and *id* is the instance number.

4. On SAP HANA database, grant these authorities to the maintenance user:

```
GRANT CREATE ANY, DELETE, DROP, EXECUTE,
INDEX, SELECT, UPDATE ON SCHEMA myschema TO muser
```

where *myschema* is your SAP HANA database schema, and *muser* is the maintenance user you just created.

5. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name instead of SAMPLE_RS.

6. Create a connection to the replicate SAP HANA database instance using ExpressConnect for HANA DB.

For a standard connection:

```
create connection to rds.rdb
using profile rs_msss_to_hanadb;ech
set username muser
set password mpwd
go
```

For SAP Secure User Store:

```
create connection to rds.rdb
using profile rs_msss_to_hanadb;ech
set username auser
set password apwd
set dsi_connector_sec_mech to "hdbuserstore"
go
```

where:

- *rds* is the replicate SAP HANA database. For a standard connection, this must match the data server name in the interfaces file entry. See the *Adaptive Server Enterprise Configuration Guide for UNIX > Set Up Communications Across the Network > Contents of the interfaces File* and the *Adaptive Server Enterprise Configuration Guide for Windows > Network Communications Using sql.ini*. For an SAP Secure User Store connection, this must match what you used as the key to create a user store of encrypted credentials with the **hdbuserstore** utility.
- *rdb* is a placeholder; you must provide a value, but it is not used.
- *muser* is the maintenance user for the replicate SAP HANA database instance you have already created.
- *mpwd* is the replicate SAP HANA database maintenance user password.
- *auser* and *apwd* are unused values supplied only to satisfy the syntax of the **create connection** command.

Note: For an SAP Secure User Store connection, the same user who ran the **hdbuserstore** utility must run the **create connection** command.

Microsoft SQL Server-to-SAP HANA Database Replication Setup

If you have trouble connecting to the replicate SAP HANA database instance using ExpressConnect for HANA DB, see the *Replication Server Troubleshooting Guide > Common Error Messages > Connector Error Messages*.

7. Use **admin show_connections, 'replicate'** to display the replicate connection you created.

Configuring Replication Server for Replication from the Primary Database

Configure Replication Server for replication from the primary database by granting create object permission to the **rs_username** user.

The **rs_username** user, which is used for communication between Replication Agent and Replication Server, must have **create object** permission before you can use it to create a replication definition. You must grant this permission manually from Replication Server.

1. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

```
isql -Uusername -Ppassword -SSAMPLE_RS
```

2. Enter:

```
grant create object to SAMPLE_RS_ra  
go
```

where *SAMPLE_RS_ra* is the value for the **rs_username** parameter you entered in the resource file.

Configuring Replication Agent

Review the procedures for configuring Replication Server for Microsoft SQL Server.

Primary Microsoft SQL Server Configuration and Verification

Configure and verify the primary Microsoft SQL Server database for replication.

Verifying the Compatibility Level of Microsoft SQL Server

Verify that the database compatibility level is set to Microsoft SQL Server 2008, Microsoft SQL Server 2008 R2, or Microsoft SQL Server 2012.

To verify the compatibility level:

1. In the Microsoft SQL Server Management Studio, right-click your primary database and select **Properties**.
2. In Database Properties, click **Options**, and verify that the compatibility level is "SQL Server 2008 (100) or SQL Server 2012 (110)." If it is not, select that compatibility and click **OK**.

Disabling Microsoft SQL Server Replication

A Microsoft SQL Server publication cannot be created in the primary database used by Replication Agent for Microsoft SQL Server, and you cannot simultaneously use Microsoft replication and Replication Agent on the same Microsoft SQL Server database.

If a Microsoft SQL Server publication already exists, remove the publication before using Replication Agent for Microsoft SQL Server.

Creating a Microsoft SQL Server User and Granting Permissions

Create a Microsoft SQL Server user named ra_user, and grant permissions to it.

Log in to the primary Microsoft SQL Server as a system administrator, and run:

```
use master;
create login ra_user
    with password = 'my_pass';
use <primary_database>;
create user ra_user for login ra_user;
EXEC sp_addsrvrolemember 'ra_user', 'sysadmin';
```

where *my_pass* is a password that complies with the Microsoft SQL Server complexity policy.

Enabling Remote DAC

Enable remote DAC.

1. Log in to the server:

Microsoft SQL Server 2008	"C:\Program Files\Microsoft SQL Server \100\Tools\Binn\SQLCMD.EXE" -U <i>username</i> -P <i>password</i> -S <i>serverName</i>
Microsoft SQL Server 2012	"C:\Program Files\Microsoft SQL Server \110\Tools\Binn\SQLCMD.EXE" -U <i>username</i> -P <i>password</i> -S <i>serverName</i>

where *username*, *password*, and *serverName* are your user ID, password, and Microsoft SQL Server name.

2. Enable DAC to allow remote connections:

```
1> sp_configure 'remote admin connections', 1
2> go
```

3. Update the **remote admin connections** setting:

```
1> reconfigure
2> go
```

4. Verify the **remote admin connections** setting:

```
1> sp_configure 'remote admin connections'
2> go
```

The result returned should be:

name	minimum	maximum	config_value	run_value
-----	-----	-----	-----	-----

```
-----  
remote admin connections 0          1          1          1
```

Obtaining the Microsoft SQL Server DAC Port Number

Obtain the Microsoft SQL Server DAC port number.

1. Use a text editor to open the ERRORLOG file. This file is located in the log directory of your Microsoft SQL Server. For example:

```
C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG\ERRORLOG
```

2. Search for "Dedicated admin" to find an entry similar to this:

```
2007-11-09 13:40:02.40 Server Dedicated admin  
connection support was established for listening  
locally on port 1348.
```

3. Record the port number specified in this entry for use in a later step.

Verifying the Version of the Microsoft Filter Manager Library

Confirm that the Microsoft Filter Manager Library version is 5.1.2600.2978 or later, to make the primary transaction log files readable for Replication Agent.

To determine the version of the library:

1. Navigate to `c:\windows\system32`.
2. Right-click **fltlib.dll** and select **Properties**.
You see the **Properties** dialog box.
3. On the **Properties** dialog box, click the **Details** tab.
4. If the version is earlier than 5.1.2600.2978, go to the Microsoft Web site at <http://www.update.microsoft.com/windowsupdate/v6/vistadefault.aspx?ln=en-us>, and update your Windows system.

Making the Primary Transaction Log Files Readable for Replication Agent

Install and set up the `sybfilter` driver so that Replication Agent can read the primary transaction log files.

1. In Windows Explorer, navigate to the `sybfilter` driver installation directory. This directory is located at `%SYBASE%\RAX-15_5\system\<platform>`, where `<platform>` is:
 - `winx64` – for 64-bit Windows Server 2008, Windows Server 2008 R2, and Windows 7.
 - `winx86` – for 32-bit Windows Server 2008, Windows Server 2008 R2, and Windows 7.
 - `winvistax64` – for all other supported 64-bit Windows versions.

2. In Windows Explorer, navigate to the `sybfilter` driver installation directory. This directory is located at `%SYBASE%\RAX-15_2\system\, where <platform> is winx86, winx64, or winvistax64.`

3. To install the driver, right-click `sybfilter.inf`, and click **Install**.

Note: There can be only one installation of the `sybfilter` driver on a Windows machine. After the driver is installed, it works for all Replication Agent for Microsoft SQL Server instances running on the same machine.

4. Under any directory, create a configuration file to store all log file paths for primary databases. The configuration file must have a `.cfg` suffix. For example, under the directory `%SYBASE%\RAX-15_5\system\, create a file named LogPath.cfg.`

5. Under any directory, create a configuration file to store all log file paths for primary databases. The configuration file must have a `.cfg` suffix. For example, under the directory `%SYBASE%\RAX-15_2\system\, create a file named LogPath.cfg.`

6. Add a system environment variable named `RACFGFilePath`, and set its value to the path of the configuration file.

a) Open the **Control Panel**, click **System**, click the **Advanced** tab, and choose **Environment Variables**.

b) Click **New** to add a new system variable.

c) Name the variable `RACFGFilePath`, and set its value to the location of your configuration file.

7. In Windows Explorer, navigate to `%SYBASE%\RAX-15_5\bin`, and right-click `sybfiltermgr.exe` and select **Run as administrator** to start the `sybfilter` driver management console.

8. In Windows Explorer, navigate to `%SYBASE%\RAX-15_52bin`, and double-click the `sybfiltermgr.exe` file to start the `sybfilter` driver management console.

9. To start the `sybfilter` driver, enter **start** at the management console.

10. Add the log file path to the `sybfilter` driver with the user manager or by modifying the configuration file:

- User manager – use the **add** command in the management console. The syntax for this command is:

```
add serverName dbName logFilePath
```

For example:

```
add myserverName dbName C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb2_log.ldf
```

The user manager refreshes all log paths in the `sybfilter` driver automatically after adding the log path into the configuration file.

Microsoft SQL Server-to-SAP HANA Database Replication Setup

- Configuration file – to add the log file path directly to the configuration file, open and manually edit the configuration file. This an example of log file path entries:

```
[myserver, pdb1]
log_file_path=C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb11_log.ldf
log_file_path=C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb12_log.ldf
[myserver, pdb2]
log_file_path=C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb2_log.ldf
```

After you have added the log file paths to the configuration file, use the **refresh** command in the management console.

11. If you added a log file, restart Microsoft SQL Server to make the log file readable.
12. At the management console, verify that log files are readable. If some log files are unreadable, make sure the files have been created and that Microsoft SQL Server has been restarted, if necessary.

See the *Replication Agent Primary Database Guide*.

Stopping the Analysis Service

Terminate the SQL Server Analysis Service if it is running.

1. From the Windows Control Panel, select **Administrative Tools > Services**.
2. Find the service named SQL Server Analysis Service (*SERVER*), where *SERVER* is the name of your Microsoft SQL Server data server.
3. Stop this service.

Enabling TCP/IP

Enable TCP/IP.

1. Go to **Microsoft SQL Server > Configuration Tools**, and choose **SQL Server Configuration Manager > SQL Server Network Configuration > Protocols for [SQLInstanceName]**, where *SQLInstanceName* is your SQL Server instance.
2. Right-click **TCP/IP**, and choose **Enable**.

Installing the Microsoft SQL Server JDBC Driver

Install the Microsoft SQL Server JDBC driver.

1. Go to the Microsoft download site at <http://www.microsoft.com/downloads> and search for:

```
Microsoft JDBC Driver 4.0
```

2. Click the link for the 4.0 driver, and follow the instructions to download and install the driver.

3. Open **Control Panel > System**, click the **Advanced** tab, and click **Environment Variables**.
4. In the System Variables window, if the CLASSPATH variable is not listed, click **New**. Otherwise, select it, and click **Edit**.
5. Enter the full path of the JDBC driver named `sqljdbc4.jar`, using a semicolon (;) to separate it from any other drivers.
6. Click **OK** three times.

Configuring and Validating the Replication Agent Resource File

A single installation of the Replication Agent can support replication from multiple databases. However, one Replication Agent instance is needed for each Microsoft SQL Server database that is to be replicated.

1. Locate the resource file template.

The majority of configuration values required to create and initialize a Replication Agent. Using a resource file allows an instance to be easily removed and re-created.

The Microsoft SQL Server resource file template is at `%SYBASE%\RAX-15_5\init\mssql.rs`.

The Microsoft SQL Server resource file template is at `%SYBASE%\RAX-15_2\init\mssql.rs`.

2. Create an instance resource file:

Copy the resource file template `%SYBASE%\RAX-15_5\init\mssql.rs` to another file, for example, `myra.rs`.

Copy the resource file template `%SYBASE%\RAX-15_2\init\mssql.rs` to another file, for example, `myra.rs`.

Change the values that are assigned to properties in the resource file to match the values required for your environment. Most of these values are host, port, user, and password information for the Microsoft SQL Server, Replication Server, and RSSD connections.

- Set the value of **rs_source_ds** to the value you chose for *rax* and *pds* in previous procedures.
- Set the value of **rs_source_db** to the value you chose for *test* and *pdb* in previous procedures.
- When you create a Replication Agent instance with a resource file:
 - The **asa_password** configuration parameter value cannot contain single quotes, double quotes, or a semicolon.
 - The **pds_username** and **pds_password** configuration parameter values cannot contain single or double quotes if the **create_pds_username** parameter is set to yes.

The resource file is self-documenting, with each parameter described. See the *Replication Agent Administration Guide*.

Warning! The combined values of **rs_source_ds** and **rs_source_db** must match the *{pds.pdb}* values of the Replication Server primary connection name.

Table 10. Resource File Parameters

Parameter	Description	Example Values
instance name	Any valid name.	myra
admin_port	Port number that Replication Agent uses.	9030 (if in use, select a different port number)
ltm_admin_user	User name for administering the Replication Agent instance.	sa
ltm_admin_pw	Password for the user administering the Replication Agent instance.	The default password policy is that the password must be between 6 and 255 characters. For example: sa_pass
pds_server_name	Name of the primary Microsoft SQL Server.	TEAMSTER
pds_port_number	Port number for the primary Microsoft SQL Server.	1433
pds_dac_port_number	Microsoft SQL Server dedicated administrator connection (DAC) port number.	1348
pds_database_name	Name of the primary database.	test_db
pds_username	User ID that Replication Agent uses to access the primary data server.	ra_user
pds_password	Password for pds_username .	sybase
rs_host_name	Machine where Replication Server is installed.	teamster
rs_port_number	Port number where Replication Server is installed.	11752
rs_username	Replication Server user with connect source and create object permissions.	<i>SAMPLE_RS_ra</i>
rs_password	Password for rs_username .	SAMPLE_RS_ra_ps

Parameter	Description	Example Values
rs_source_ds	Valid name representing data server of primary Microsoft SQL Server database. Note: If you use direct load materialization, add an entry with this name in the interfaces file.	pds
rs_source_db	Valid name representing primary Microsoft SQL Server database.	pdb
rs_charset	Character set that Replication Server is using. The value defined for the rs_charset configuration parameter must match the RS_charset value in the Replication Server configuration file, %SYBASE%\REP-15_5\install\<server>.cfg. The value defined for the rs_charset configuration parameter must match the RS_charset value in the Replication Server configuration file, %SYBASE%\REP-15_2\install\<server>.cfg.	cp850
rssd_host_name	Machine where RSSD resides.	teamster
rssd_port_number	Port number where RSSD resides.	11751
rssd_database_name	Database name for RSSD.	SAMPLE_RS_ERSSD
rssd_username	Valid user for RSSD.	SAMPLE_RS_RSSD_maint
rssd_password	Password for rssd_username .	SAMPLE_RS_RSSD_maint_ps
start_instance	Automatically start the instance that was created.	yes
initialize_instance	Automatically initialize the new Replication Agent instance.	no

3. Be sure your CLASSPATH environment variable points to the Microsoft SQL Server JDBC driver, and use the Replication Agent **ra_admin** utility to validate the settings in the new instance resource file using the **-vr** parameter. For example:

```
%SYBASE%\RAX-15_5\bin\ra_admin -vr myra.rs
```

```
%SYBASE%\RAX-15_2\bin\ra_admin -vr myra.rs
```

Validation results are returned as either:

Microsoft SQL Server-to-SAP HANA Database Replication Setup

- Response-file processing completed, or,
- Response-file processing completed with errors

If any validation fails, the **ra_admin** utility returns an error message and information about the failure. You can repeat the validation process as many times as necessary until it executes without error. No entities are changed or created.

Initialization

To initialize the primary data server, Replication Agent for Microsoft SQL Server installs objects at the data-server level and at the database level.

Data-server level modifications are required only once.

If you have previously initialized the primary data server and a Replication Agent instance and want to initialize another Replication Agent instance, see *Subsequent Initialization*.

See also

- *Subsequent Initialization* on page 86
- *Verifying the Replication Agent Replicating State* on page 87

Configuring the Primary Data Server and a Replication Agent Instance

Configure the primary data server and a Replication Agent instance.

1. Stop the Microsoft SQL Server service:

- a) In **Control Panel > Administrative Tools > Services**, find the service named Microsoft SQL Server (*SERVER*), where *SERVER* is the name of your Microsoft SQL Server data server. For example:

```
Microsoft SQL Server (TEAMSTER)
```

- b) Stop the service.

2. Restart Microsoft SQL Server in single-user mode:

- a) Click **Start > Control Panel > Administrative Tools > Services**.
- b) Right-click your Microsoft SQL Server instance, and choose **Properties**.
- c) In the **General** tab, click **Stop**.
- d) Under Start parameters, enter:

```
-m
```

Click **Start**.

3. Execute the resource file to create the Replication Agent instance. In your resource file, if you set **start_instance** to **yes**, this step also starts the Replication Agent instance. In your resource file, if you also configured **initialize_instance** to **yes**, this step also initializes the Microsoft SQL Server and the Replication Agent instance.

After validating the resource file, create the Replication Agent instance:

```
%SYBASE%\RAX-15_5\bin\ra_admin -r myra.rs
```

Execution results are returned as either:

- Response-file processing completed, or,
- Response-file processing completed with errors

See the *Replication Agent Administration Guide*.

4. If **start_instance** is **no** in your resource file, change to the SYBASE directory and start the Replication Agent instance:

```
cd %SYBASE%\RAX-15_5\myra
```

Execute the RUN file, for example:

```
RUN_myra
```

5. Use **dsedit** to update the Replication Server `sql.ini` to include an entry for the Replication Agent location.

To use direct load materialization, also add the server name part of the connection name to the interfaces file. This entry should point to Replication Agent. For example:

```
pds
master tcp ether hostname 9030
query tcp ether hostname 9030
```

6. Verify the connection to Replication Agent.

- a) Open a command window in the %SYBASE% directory of your Replication Agent installation.
- b) Set the environment variables by executing the SYBASE.bat file.
- c) Log in to Replication Agent:

```
isql -Usa -Psa_pass -Smyra
```

7. If your resource file has **initialize_instance** set to **no**, initialize the primary data server and the Replication Agent instance:

- a) Open a command window, and log in to the Replication Agent instance.
- b) Create server xlog objects:

```
server_xlog init
```

- c) Shut down Replication Agent:

```
shutdown
```

8. Stop the Microsoft SQL Server in single-user mode:

- a) Log in to the server:

```
"C:\Program Files\Microsoft SQL Server\90\Tools\Binn\SQLCMD.EXE" -U username -P password -S serverName
```

where *username*, *password*, and *serverName* are your user ID, password, and Microsoft SQL Server name.

- b) Issue the **shutdown** command.

9. Restart Microsoft SQL Server in multiuser mode (normal start):

- a) In **Control Panel > Administrative Tools > Services**, find the service named Microsoft SQL Server (*SERVER*), where *SERVER* is the name of your Microsoft SQL Server data server. For example:

```
Microsoft SQL Server (TEAMSTER)
```

- b) Start the service.

10. Restart Replication Agent.

- a) Change to the SYBASE directory and start the Replication Agent instance:

```
cd %SYBASE%\RAX-15_5\myra
```

- b) Execute the RUN file, for example:

```
RUN_myra
```

11. Initialize Replication Agent.

- a) Log in to Replication Agent using **isql**:

```
isql -Usa -P -Smyra
```

where *myra* is the name of Replication Agent.

- b) Issue:

```
ra_admin init
```

The Replication Agent instance is now running, and the primary data server and the Replication Agent instance have been initialized.

Subsequent Initialization

Create another Replication Agent instance, after you have already initialized the primary data server.

See also

- *Verifying the Replication Agent Replicating State* on page 87

Creating and Initializing a Replication Agent Instance

Execute the resource file to create the Replication Agent instance.

1. After validating the resource file, create the Replication Agent instance:

```
%SYBASE%\RAX-15_5\bin\ra_admin -r myra.rs
```

Execution results are returned as either:

- Response-file processing completed
- or:
- Response-file processing completed with errors

See the *Replication Agent Administration Guide*.

2. If **initialize_instance** is **no** in your resource file, start the Replication Agent instance.

- a) Change to the directory containing your Replication Agent instance RUN file.
- b) Execute the RUN file, for example:

```
RUN_myra
```

3. Use **dsedit** to update the Replication Server `sql.ini` to include an entry for the Replication Agent location.

To use direct load materialization, also add the server name part of the connection name to the interfaces file. This entry should point to Replication Agent. For example:

```
pds
master tcp ether hostname 9030
query tcp ether hostname 9030
```

Note: The server name is the same value as the **rs_source_ds** value in the resource file used to create a Replication Agent instance.

4. Verify the connection to Replication Agent.
 - a) Open a command window in the `%SYBASE%` directory of your Replication Agent installation.
 - b) Set the environment variables by executing the `SYBASE.bat` file.
 - c) Log in to Replication Agent:

```
isql -Usa -Psa_pass -Smyra
```

5. If **initialize_instance** in **no** in your resource file, open a new command window, log in to the Replication Agent instance, and initialize the Replication Agent instance and the primary Microsoft SQL Server database using this Replication Agent command:

```
ra_admin init
```

Your Replication Agent instance is now running and has been initialized.

Verifying the Replication Agent Replicating State

Verify the replicating state of Replication Agent.

1. Enter:

```
resume
go
```

If the Replication Agent successfully transfers to a replicating state, you see:

```
State          Action
-----
REPLICATING Ready to replicate data.
```

The Replication Agent goes to the **REPLICATION DOWN** state if an error occurs:

```
isql -Usa -Psa_pass -SSAMPLE_RS
1> resume
2> go
Msg 32000, Level 20, State 0:
Server 'myra', Procedure 'resume', Line 1:
Command <resume> failed - Desired state <REPLICATE>
```

Microsoft SQL Server-to-SAP HANA Database Replication Setup

```
could not be achieved. Current state: <REPLICATION  
DOWN>
```

2. The **ra_status** command returns the state of the Replication Agent. Verify that the Replication Agent remains in replicating state, even after the resume command executes successfully.

To detect an error that occurred after replication start-up, execute:

```
ra_status  
go
```

If the Replication Agent is in replicating state, **ra_status** returns:

```
State           Action  
-----  
REPLICATING    Ready to replicate data.
```

Replication Agent goes to the REPLICATION DOWN state if an error occurs.

3. Validate that both primary and replicate connections are active:

```
isql -Usa -Psa_pass -SSAMPLE_RS  
admin who  
go
```

Note: Be aware that:

- The DSI connection for the primary database connection is usually suspended because you are not replicating data back to the primary database.
 - The Replication Agent connection, if established for the replicate database connection, is usually down, because you are not replicating data from the replicate database.
-

Do not proceed until **admin who** returns status for threads similar to this:

```
Spid Name           State           Info  
-----  
13  DSI EXEC  Awaiting Command 101(1) SAMPLE_RS_ERSSD.  
SAMPLE_RS_ERSSD  
9   DSI      Awaiting Message 101 SAMPLE_RS_ERSSD.  
SAMPLE_RS_ERSSD  
8   SQM      Awaiting Message 101:0 SAMPLE_RS_ERSSD.  
SAMPLE_RS_ERSSD  
8   SQM      Awaiting Message 101:0 SAMPLE_RS_ERSSD.  
SAMPLE_RS_ERSSD  
54  DSI EXEC  Awaiting Command 102(1)  
my_mssql_access_service.rdb  
53  DSI      Awaiting Message 102  
my_mssql_access_service.rdb  
17  SQM      Awaiting Message 102:0  
my_mssql_access_service.rdb DSI  
EXEC      Suspended      103(1) pds.pdb  
DSI      Suspended      103 pds.pdb  
24  DIST     Awaiting Wakeup 103 pds.pdb  
25  SQT      Awaiting Wakeup 103:1 DIST pds.pdb  
23  SQM      Awaiting Message 103:1 pds.pdb
```

```
22 SQM Awaiting Message 103:0 pds.pdb
62 REP AGENT Awaiting Command pds.pdb
```

Testing Replication

When you finish setting up the replication system, test that replication works as intended.

1. Connect to the primary Microsoft SQL Server database as a regular user, rather than as the maintenance user. Make sure the regular user also exists in the replicate database.
2. At the primary Microsoft SQL Server database, create a table named `ptabl` to replicate:

```
CREATE TABLE ptabl
(idno int PRIMARY KEY,
name varchar(20));
```

3. Grant permissions to any new or existing object to be replicated in the primary database:

```
grant all on ptabl to public;
```

4. Connect to Replication Agent through `isql`, and mark the `ptabl` table for replication:

```
pdb_setreptable ptabl, mark
go
```

5. Add rows to the `ptabl` table:

```
insert into ptabl values (1, 'Burt Gringlesby');
insert into ptabl values (2, 'Douglas Wong');
insert into ptabl values (3, 'Tom Servo');
insert into ptabl values (4, 'Innes del Castillo');
insert into ptabl values (5, 'Akiko Yokomoto');
```

6. In Replication Agent, create a replication definition against the primary Microsoft SQL Server database:

```
rs_create_repdef ptabl
go
Table/Procedure Name      RepDef Name      Status
-----
<MSSQLTableOwner>.PTAB1 "<repdefname>" Created
(1 row affected)
```

Record the name of the replication definition for use later.

7. At the replicate SAP HANA database instance, create a table named `PTAB1`:

```
CREATE TABLE <tableowner>.PTAB1
(IDNO INT PRIMARY KEY,
NAME VARCHAR(20));
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

If an owner is not specified with the table name, the owner of the table—`<tableowner>`—is the user who is logged in when the table is created.

- Grant permissions to any new or existing object to be replicated in the replicate database so that the Replication Server maintenance user can update this table:

```
grant all privileges on <tableowner>.PTAB1 to public
```

- Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

- Alter the replication definition that you created earlier to point to the correct replicate table:

```
alter replication definition <repdefname>  
with replicate table named <tableowner>.ptab1  
go
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

- Create a subscription against the replicate connection to materialize the PTAB1 table on the SAP HANA database.

Create the `ptab1_sub` subscription with **direct_load** option for the `<repdefname>` replication definition. Replication Server logs in to Replication Agent for Microsoft SQL Server, and Replication Agent issues a **SELECT** command to select data from the Microsoft SQL Server database, enter:

```
create subscription ptab1_sub  
for <repdefname>  
with replicate at rds.rdb  
without holdlock  
direct_load  
user puser password ppwd  
go
```

where:

- `rds.rdb` is the replicate SAP HANA database connection created earlier.
- `puser` is the user who connects to Replication Agent for Microsoft SQL Server.
- `ppwd` is the password of the `puser`. You must enter a password.

Note: The interfaces file used by the replicate Replication Server must contain an entry for the Replication Agent for Microsoft SQL Server that is used to select data from the Microsoft SQL Server database.

To create the subscriptions at the replicate Replication Server using the no materialization method, use the **create subscription** command with the **without materialization** clause.

For information about using the **create subscription** command in **direct_load** materialization, see *Replication Server Reference Manual*.

- Check the subscription status to ensure there are no errors:


```
check subscription ptab1_sub
for <repdefname>
with replicate at rds.rdb
go
```

If there are errors, see *Troubleshooting Subscription Problems*.

13. Connect to the Microsoft SQL Server database as a non-maintenance user with the ability to insert rows into the table:

```
insert into ptab1 values (10, 'Michel DeFrance');
insert into ptab1 values (11, 'Dirk Stringer');
commit;
```

14. On the Replication Server, check the status of the subscription created with the **direct_load** option by using the **check subscription** command. When the subscription is in the **VALID** state, materialization is complete:

```
check subscription ptab1_sub
for <repdefname>
with replicate at rds.rdb
go
Subscription ptab1_sub is VALID at the replicate.
```

15. After the status is **VALID** at the primary and replicate database, on the replicate database, verify that the data is replicated to the replicate **PTAB1** table:

```
SELECT * FROM <tableowner>.PTAB1
```

See also

- *Configuring Replication Server for Replication from the Primary Database* on page 76

Resetting the Primary Microsoft SQL Server Database for Replication

In a test environment, you may occasionally need to reset the replication environment.

1. Protect your new environment from old log information by archiving all current redo log data. In the Replication Agent, run:

```
ra_locator move_truncpt
```

2. Reinitialize Replication Agent, which refreshes the Replication Agent repository rather than overwriting it:

```
ra_admin refresh
```

Note: If you prefer to delete and replace all the information in the Replication Agent repository, issue **ra_admin deinit, force** followed by **ra_admin init** (without the **force** option).

3. Reset the locator stored in Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
connect
go
rs_zeroltm pds, pdb
go
```

Microsoft SQL Server-to-SAP HANA Database Replication Setup

```
exit  
go
```

You see:

```
Locator has been reset to zero.  
(return status = 0)
```

DB2 UDB-to-SAP HANA Database Replication Setup

Implement a sample replication system from DB2 UDB to SAP HANA database.

Use the tasks in this document to set up a proof-of-concept or test replication system.


The tasks use the sample Replication Server instance SAMPLE_RS. If you have already installed Replication Server, skip to the configuration procedures.

For the most current information about Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

Multiple components are required to implement a DB2 UDB-to-SAP HANA database replication system: a primary DB2 UDB data server, a replicate SAP HANA database, a Replication Server with ExpressConnect for HANA DB, and Replication Agents for DB2 UDB.

Figure 4: DB2 UDB-to-SAP HANA Database Replication Setup



Prerequisites

Review the prerequisites before you begin setting up DB2 UDB-to-SAP HANA database replication.

- Familiarity with DB2 UDB and SAP HANA database.
- Valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- Valid licenses for the Replication Server and the Replication Server Heterogeneous Edition.

DB2 UDB-to-SAP HANA Database Replication Setup

- Install and configure DB2 UDB 9.5, 9.7, 10.1, or 10.5 to serve as the primary (source) database from which Replication Agent replicates transactions, and SAP HANA database as the replicate (target) database.
- TCP/IP connectivity between all hosts.
- Install any operating system patches required for Java.

Do not use Replication Server reserved words for object names or connection names. The *Replication Server Reference Manual* includes a complete list of reserved words.

Planning Your Installation

Before installing the replication components, verify the system requirements, identify the installation directories, and obtain the installation software.

Installation Checklist

Use the checklist to navigate through the installation when you set up your replication environment.

Table 11. Installation Checklist

	Action	Topic
1	Make sure the target computer meets the minimum memory and disk space requirements.	<i>System Requirements</i>
2	Determine the directory in which to install the replication components.	<i>Identify the Installation Directory</i>
3	Download replication components from the SAP Service Marketplace (SMP).	<i>Downloading Products from SMP</i>
4	Install Replication Server.	<i>Installing Replication Server</i>
5	Install Replication Agent.	<i>Installing Replication Agent</i>

System Requirements

Make sure the target computer on which you are installing the replication components meets the minimum memory and disk space requirements.

See the *Replication Server Installation Guide > Planning Your Installation > System Requirements* and the *Replication Agent Installation Guide > Planning Your Installation > System Requirements*.

To improve performance by avoiding multiple network hops, install Replication Server and the SAP HANA database on the same machine.

Note: The only versions of Linux that ExpressConnect for HANA DB supports are Red Hat Enterprise Linux 6 and SuSE Linux Enterprise Server (SLES) 11.

Identify the Installation Directory

Determine the directory in which to install the replication components.

The quick-start procedures described here assume you have installed the SAP software on the same host as the DB2 UDB server. Replication Agent, Replication Server, and ExpressConnect for HANA DB are installed in the same base directory, which is identified by the SYBASE environment variable. Select a path on the host drive to be the recipient of the installation and configuration activities. Replication Agent must be installed on the same host as either the DB2 UDB server or the DB2 Administration Client.

Replication Agent Installation Directory

Replication Agent versions 15.5 and later use the RAX-15_5 directory name. If you have an existing Replication Agent installation of version 15.5 or later and do not want to overwrite it, install the latest version of Replication Agent in a different location.

Downloading Products from SMP

Download replication components from the SAP Service Marketplace (SMP).

Prerequisites

An SMP login ID and password.

Task

Note: If you are installing SAP® Replication Server® Options (RSO), there are several steps you must perform twice, because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/patches>, and enter your login ID and password.
2. Select **SAP Software Download Center > Installations and Upgrades > A-Z Index** and click **R** for SAP Replication Server products.
3. From the product listing, select:
 - **Sybase Replication Server, Option for HANA** or
 - **Sybase Replication Server, HANA Edition**
4. Select the software you need by version and platform.
5. Generate a license based on the license model—served or unserved—and license type:
 - Application Deployment CPU License (AC)
 - Application Deployment Other License (AO)
 - Application Deployment Standby CPU License (BC)
 - CPU License (CP)

- Development and Testing License (DT)
- Other License (OT)
- Standby CPU License (SF)
- Server License (SR)
- Standalone Seat License (SS)

Note: To generate a license, you must provide some information, such as the host name, MAC address, and number of CPUs.

6. Download your product licenses, and place them in your product license directory after installation:
 - `$SYBASE/SYSAM-2_0/licenses` (UNIX or Linux)
 - `%SYBASE%\SYSAM-2_0\licenses` (Windows)where `$SYBASE` or `%SYBASE%` is where you installed your product.
7. Download your installation files, and uncompress and extract all the installation images to your local drive.

Installing Replication Components

Review the procedures for installing the replication components.

Install all replication components on the same host where a DB2 UDB data server has already been installed and is running with the primary database. This simplifies the quick-start process.

Installing Replication Server

Install Replication Server using the setup program.

Prerequisites

- Allocate a disk partition of at least 20MB for each Replication Server you are installing. You can add more partitions later, if necessary. Make sure each partition is available and has write permissions.
- Allocate the entire partition to the Replication Server. If you allocate only a portion of the partition to Replication Server, you cannot use the remainder for any other purpose.
- The sample Replication Server instance `SAMPLE_RS` used in these procedures must be created when you install Replication Server. If you have already installed Replication Server, see the instructions for how to set up a Replication Server instance after installation in *Replication Server ASE-to-ASE Replication Quick Start Guide > SAMPLe_RS*.

Task

1. Verify that the drive on which you install the Replication Server has enough disk space for the components being installed, and at least 100MB of extra disk space for the installation program.
2. Download and extract the Replication Server installation image from the SAP® Service Marketplace (SMP).
3. Start the installation.

- On Windows, launch the **setup** program.

If the installer does not start automatically, double-click **setup.exe** or select **Start > Run** and enter:

```
setup.exe
```

If there is not enough disk space in the temporary disk space directory, set the environment variable `TMP` to *directory_name* before running it again, where *directory_name* is the full path to and name of the temporary directory to which the installation program writes the temporary installation files.

- (UNIX or Linux) If you downloaded the product, go to the directory where you extracted the installation image and start the installer:

```
./setup.bin
```

where `setup.bin` is the executable file name for installing Replication Server.

If there is not enough disk space in the temporary disk space directory, set the environment variable `IATEMPDIR` to *tmp_dir* before running the installer again, where *tmp_dir* is the full path to the temporary directory to which the installation program writes the temporary installation files.

4. In the Introduction window, click **Next**.
5. Indicate where to install Replication Server:
 - To select an installation directory, click **Choose**, and browse for and select the directory.
 - To create a new directory, enter a new directory path.
 - To restore to the default directory, click **Restore Default Folder**.

If the installation directory you chose does not exist, you see:

```
The directory directory name does not exist. Do you want to create it?
```

Click **Yes**.

If the installation directory you selected exists, and already contains an Replication Server installation, you see:

```
Warning: You have chosen to install into an existing directory. If you proceed with this installation, any older versions of the products you choose to
```

install that are detected in this directory will be replaced.

Click **Next**.

6. Select the type of installation:

- **Typical** – installs the default components. This is recommended for most users.
- **Full** – installs every Replication Server component, including all the supported language modules.
- **Custom** – lets you select the components to install. Some components are automatically installed if they are required to run your selected components.

Click **Next**.

7. Select the geographic location, agree to the license agreement, then click **Next**.

8. On the SySAM License Server window, select one of:

- **Specify License Keys** – click **Browse** to select the license file. To select multiple license files, use **Shift+Click** or **Ctrl+Click**. The license pane displays the license information.

Alternatively, copy and paste the license information directly in the license pane. Click **Next**.

If you specify a served license key, the installer prompts you to install a new SySAM license server. Click:

- **Next** to install the new SySAM license server and follow the installation prompts.
- **Previous** and select **Use previously deployed license server** if you have an existing SySAM license server on the same host.
- **Use previously deployed license server** – for an existing license server, enter the host name of the machine where the license server is running, and the port number if the port number you are using is not the default. Click **Next**.
- **Continue installation without a license key** – click **Next** to proceed without any license keys. The installer allows you to install and use the Replication Server components without a license for 30 days. To continue using these components after the end of the grace period, obtain valid licenses and install these licenses using the Replication Server License Installer.

See the *SySAM Users Guide*.

9. Configuring your server for e-mail notification enables designated users to receive information about license management events requiring attention.

Provide:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipient e-mail addresses
- Message severity that triggers e-mail messages

Click **Next**.

10. The installation summary window displays the selections you have made. Review the information, and click **Install**.

11. On the Start Sample Replication Server window, select:

- **Yes** to configure and start a sample Replication Server. The installer displays the configuration information for the sample Replication Server. Record this information. You must enter and confirm a password from 6 to 15 characters in length.

Note: You cannot start the sample Replication Server on Linux on POWER because the ERSSD requires SAP SQL Anywhere Server, which is not available on Linux on POWER.

- **No** to manually configure a full-featured Replication Server and start a sample Replication Server after installation. The installer prompts you to either create a sample Replication Server directory or to proceed with the installation. Select either:
 - **Yes** – to create a sample Replication Server directory without starting the sample Replication Server, or,
 - **No** – to continue with the installation.

See the *Replication Server Configuration Guide* for Windows or UNIX.

Note: For a sample Replication Server instance, the **rs_charset** parameter must be set to utf8.

Click **Next**.

12. When the final window appears, indicating a successful installation, click **Done** to exit the installer.

Installing Replication Agent

Install Replication Agent using the GUI wizard.

Note: If there is not enough disk space in your default temporary directory, set the IATEMPDIR (Linux or UNIX) or TMP (Windows) environment variable to a directory that has enough space. Include the full path to this directory.

1. Log in to the Replication Agent host machine using an operating system user account with authority to start, stop, and administer the Replication Agent instance (for example, the “sybase” user).
2. Close all nonessential applications, and minimize any open windows.
3. Download and extract the Replication Agent installation image from the SAP Service Marketplace (SMP).
4. Start the installation program:

- On Microsoft Windows platforms, the installation program should start automatically. If it does not, start the installation program from Microsoft Windows Explorer by double-clicking the **setup.exe** file.

- On UNIX platforms, enter:

```
./setup.bin
```

5. Click **Next**.

Note: You can click **Cancel** to stop the installation at any point.

6. Select the geographic location, agree to the license agreement, then click **Next**.

7. Specify an installation directory. The default installation directory is:

- Existing `%SYBASE%` or `c:\sybase` on Microsoft Windows platforms
- Existing `$SYBASE` or `/opt/sybase` on UNIX platforms

Click **Next** to accept the default installation directory, or:

- Click **Choose** to select an installation directory in the file browser. Click **Next**.
- Enter a directory name. Click **Next**.

If you continue and the older products were not installed with the current version of Replication Agent, the installation program overwrites any common files.

If the directory name does not exist, click **Yes** when prompted to create it:

```
The directory does not exist.  
Do you want to create it?
```

If the directory you specified exists, you see:

```
Warning: You have chosen to install into an existing  
directory. If you proceed with this installation,  
any older versions of the products you choose to  
install that are detected in this directory will be  
replaced.
```

Note: On Microsoft Windows platforms, if you are prompted to overwrite a DLL, click **Yes** only if the version of the new DLL is newer than the one the installation program is attempting to overwrite.

8. Select the type of database the Replication Agent will connect to.

Click **Next**.

9. Select one of these options to enter the license:

- **Specify license keys** – browse to or specify the license file.
- **Use previously deployed license server** – use a previously deployed license server. Enter the host name of the machine where the license server is running and the port number if the port number you are using is not the default.
- **Continue installation without a license key** – install and use Replication Agent without a license for a grace period of 30 days. To continue using Replication Agent

after the end of the grace period, obtain a valid license from the SAP Service Marketplace (SMP), and install it.

Note: Replication Agent is licensed statically when installed with RSHE. If you are installing Replication Agent with RSHE, locate the file that contains a valid SySAM license for the Replication Agent component in RSHE. Specify that license file here.

See the *SySAM Users Guide*.

Click **Next** until you see the SySAM Notification window.

10. On the SySAM Notification window, configure your server for e-mail notification. When configuration is enabled, you receive information about license management events that require attention. Select **Yes** and either accept the default values that are supplied, or enter values for:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipients of the notification
- Message severity level of an event that triggers e-mail notification:
 - Informational
 - Warning
 - Error

If you choose not to have e-mail alerts or severity messages logged, select **No**.

Click **Next**.

11. Review the product features or components listed on the Preinstallation Summary window. Click **Install**.

The installation program installs the components in the installation directory you specified, and displays an installation progress indicator.

If errors occur during the installation, the installation program displays error messages. Exit the installation program wizard to correct the cause of the error, then restart the installation program.

If the software is installed successfully, you see a window confirming the successful installation.

12. Click **Done** to complete the installation and close the installation program.

Verifying the Installation

Verify that Replication Agent has been successfully installed.

- 1.** In a command window, change to the directory where you installed Replication Agent.
- 2.** Set the environment variables by sourcing the `SYBASE.csh` file (UNIX or Linux) or by executing the `SYBASE.bat` file (Windows).

DB2 UDB-to-SAP HANA Database Replication Setup

3. Change to \$SYBASE/RAX-15_5/bin (UNIX or Linux) or %SYBASE%\RAX-15_5\bin (Windows).

4. Obtain the Replication Agent version string:

- For Windows:

```
ra.bat -v
```

```
rao -v
```

- For UNIX or Linux:

```
./ra.sh -v
```

```
./rao.sh -v
```

Note: Make sure permission for the ra.sh file is set no broader than 700 (read/write/execute for user, no permissions for group and other).

If Replication Agent installs successfully, you see the SAP copyright and the Replication Agent version string.

Configuring Replication Components

Review the procedures for configuring each of the replication components in the replication system.

Configuration Checklist

Use the table as a checklist to navigate through the configuration when you set up your replication environment.

Table 12. Configuration Checklist

	Action	User	Topic
1	Replication Server: Set up the DB2 UDB environment variables.		<i>Setting DB2 UDB Environment Variables</i>
2	If Replication Agent for DB2UDB is installed on a different host than the primary database, it must use DB2 UDB client libraries.		<i>Node and Database Cataloging in a DB2 UDB Client</i>
3	Create a Replication Server connection to the primary database.	Maintenance user	<i>Creating a Replication Server Connection to the Primary Database</i>

	Action	User	Topic
4	Replication Server: Set up ExpressConnect for HANA DB for connections between SAP HANA database and Replication Server.		<i>Obtaining the SAP HANA Database ODBC Drivers</i>
5	Create a Replication Server connection to SAP HANA database.	Maintenance user	<i>Configuring Replication Server for Replication to SAP HANA Database on page 105</i>
6	Configure Replication Server for replication from the primary database.		<i>Configuring Replication Server for Replication from the Primary Database</i>
7	Configure the primary DB2 UDB database for replication.		<i>Primary DB2 UDB Configuration for Replication</i>
8	Prepare and create a Replication Agent instance for replication of each DB2 UDB database.		<i>Creating the Replication Agent Instance</i>
9	Verify that the Replication Agent instance has been successfully installed.		<i>Verifying the Replication Agent Instance Installation</i>
10	Initialize the Replication Agent instance.		<i>Initializing the Replication Agent Instance</i>
11	Test that replication works as intended.		<i>Testing Replication</i>

See also

- *User Management* on page 137

Setting DB2 UDB Environment Variables

For UNIX and Linux, the DB2 UDB installation provides two scripts for setting up the DB2 UDB environment variables: `db2cshrc` for C shell and `db2profile` for Bourne or Korn shell.

These scripts set the library path environment variable based on the bit size of the installed server or client.

Note: On Windows, the installation automatically sets all necessary environment variables. If Replication Agent is installed on Solaris, AIX or HP Itanium, you must configure a 64-bit DB2 UDB server or client instance. Information in this topic pertains to configuring a 64-bit DB2 UDB server.

DB2 UDB-to-SAP HANA Database Replication Setup

For UNIX platforms, the 32-bit and 64-bit versions of the driver and API libraries are located in `$HOME/sqlllib/lib32` and `$HOME/sqlllib/lib64`, respectively, where `$HOME` is the home directory of the DB2 UDB instance owner. If Replication Agent is installed on AIX, Solaris, HP Itanium, or Linux platforms, the library path environment variable must point to the 64-bit libraries. For Windows, the library path environment variable must point to the 32-bit libraries.

To correctly set the DB2 UDB environment variables for Replication Agent, you may need to customize either `db2cshrc` or `db2profile`. However, make any changes to saved copies of the original scripts because the scripts may be changed when you apply a FixPak to the DB2 UDB server or client.

1. Navigate to the DB2 UDB server or client `sqlllib` directory, which is located in the database instance owner user home directory. For example:

```
cd /db2home/db2inst1/sqlllib
```

2. Based on the type of UNIX shell you are using, create a copy of either the `db2cshrc` or `db2profile` script:

- For C shell:

```
cp db2cshrc mydb2cshrc
```

- For Bourne or Korn shell:

```
cp db2profile mydb2profile
```

3. Edit the copy of the file you just created. Add a statement at the very end of the file that sets the platform-specific library path variable to point to the correct DB2 UDB libraries.

- For C shell, in `mydb2cshrc`, add:

```
# force 32-bit libraries
setenv LD_LIBRARY_PATH $HOME/sqlllib/lib32:$LD_LIBRARY_PATH
```

where `$HOME` is the home directory of the DB2 UDB instance owner.

- For Bourne or Korn shell, in `mydb2profile`, add:

```
# force 32-bit libraries
LD_LIBRARY_PATH=$HOME/sqlllib/lib32:$LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

4. Before starting Replication Agent, always source your file, if you have customized it. Otherwise, source the original `db2cshrc` or `db2profile` script.

Node and Database Cataloging in a DB2 UDB Client

If Replication Agent for DB2 UDB is installed on a different host than the primary database, it must use DB2 UDB client libraries.

If you have installed Replication Agent for DB2 UDB on a different host than the primary database, see the *Replication Agent Primary Database Guide* for information about node and database cataloging in a DB2 UDB client.

Configuring Replication Server

Configure Replication Server for DB2 UDB-to-SAP HANA database replication.

Creating a Replication Server Connection to the Primary Database

Create a Replication Server connection to the primary database.

1. In Replication Server, create a connection to the primary DB2 UDB database:

```
create connection to pds.pdb
using profile rs_rs_to_udb_ra;standard
set username muser
set password mpwd
with log transfer on, dsi_suspended
go
```

where:

- *pds* is the value of the **rs_source_ds** parameter specified in Replication Agent.
- *pdb* is the value of **rs_source_db** specified in Replication Agent.
- *muser* is the maintenance user for the primary DB2 UDB database.
- *mpwd* is the maintenance user password.

For information about the maintenance user, see the *Replication Agent Primary Database Guide*.

2. Use **admin show_connections, 'primary'** to display the primary connection you created.

Obtaining the SAP HANA Database ODBC Drivers

Set up ExpressConnect for HANA DB for connections between SAP HANA database and Replication Server.

You must add the SAP HANA database ODBC driver to the ExpressConnect for HANA DB installation. Download these libraries from the SAP Service Marketplace and install them after you have installed Replication Server. See the *Replication Server Release Bulletin > Special Installation Instructions > Installing ODBC Libraries for ExpressConnect for HANA DB*.

Configuring Replication Server for Replication to SAP HANA Database

Create a Replication Server connection to SAP HANA database.

If you are not using the sample Replication Server instance, enter your Replication Server instance name instead of SAMPLE_RS.

1. If you are connecting to SAP HANA database with a standard connection and not using SAP Secure User Store, add an entry to your Replication Server interfaces file identifying the replicate SAP HANA database, and then stop and start Replication Server:

```
[dataservername]
master tcp ether hostname port
query tcp ether hostname port
```

DB2 UDB-to-SAP HANA Database Replication Setup

where *hostname* and *port* are the host and port number of the SAP HANA database, and *dataservername* is a label used to identify the host and port number. For information about stopping and starting Replication Server, see the *Replication Server Administration Guide Volume 1 > Manage a Replication System*.

Note: Each SAP HANA database installs with a unique instance number, and the port number is the instance number prefixed with 3 and suffixed with 15:

```
3in15
```

where *in* is the two-digit instance number. For example, the port number for an SAP HANA database with instance number 1 is 30115.

If you are using SAP Secure User Store, create a user store of encrypted credentials:

```
hdbuserstore set rds myhost:xxxxx my_securestore_user  
my_securestore_pwd
```

where:

- *rds* is the key for the secure store entry.
- *myhost.xxxxx* is the connection environment host name and port number.
- *my_securestore_user* and *my_securestore_pwd* are SAP Secure User Store credentials.

2. On SAP HANA database, create a maintenance user:

a) As the SAP HANA database admin user, create a user ID with a temporary password:

```
CREATE USER muser PASSWORD mpwd_temp
```

b) Log in to SAP HANA database with the temporary password, and then update the password:

```
ALTER USER muser PASSWORD mpwd_new
```

c) Log out of SAP HANA database, and log in again with the new password.

3. Log in to SAP HANA database:

```
./hdbsql -u user -p password -i id
```

where *user* is the SAP HANA database user, *password* is the user password, and *id* is the instance number.

4. On SAP HANA database, grant these authorities to the maintenance user:

```
GRANT CREATE ANY, DELETE, DROP, EXECUTE,  
INDEX, SELECT, UPDATE ON SCHEMA myschema TO muser
```

where *myschema* is your SAP HANA database schema, and *muser* is the maintenance user you just created.

5. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

6. Create a connection to the replicate SAP HANA database instance using ExpressConnect for HANA DB.

For a standard connection:

```
create connection to rds.rdb
using profile rs_udb_to_hanadb;ech
set username muser
set password mpwd
go
```

For SAP Secure User Store:

```
create connection to rds.rdb
using profile rs_udb_to_hanadb;ech
set username auser
set password apwd
set dsi_connector_sec_mech to "hdbuserstore"
go
```

where:

- *rds* is the replicate SAP HANA database. For a standard connection, this must match the data server name in the interfaces file entry. See the *Adaptive Server Enterprise Configuration Guide for UNIX > Set Up Communications Across the Network > Contents of the interfaces File* and the *Adaptive Server Enterprise Configuration Guide for Windows > Network Communications Using sql.ini*. For an SAP Secure User Store connection, this must match what you used as the key to create a user store of encrypted credentials with the **hdbuserstore** utility.
- *rdb* is placeholder; you must provide a value, but it is not used.
- *muser* is the maintenance user for the replicate SAP HANA database instance you have already created.
- *mpwd* is the replicate SAP HANA database maintenance user password.
- *auser* and *apwd* are unused values supplied only to satisfy the syntax of the **create connection** command.

Note: For an SAP Secure User Store connection, the same user who ran the **hdbuserstore** utility must run the **create connection** command.

If you have trouble connecting to the replicate SAP HANA database instance using ExpressConnect for HANA DB, see the *Replication Server Troubleshooting Guide > Common Error Messages > Connector Error Messages*.

7. Use **admin show_connections, 'replicate'** to display the replicate connection you created.

Configuring Replication Server for Replication from the Primary Database

Configure Replication Server for replication from the primary database by granting create object permission to the **rs_username** user.

The **rs_username** user, which is used for communication between Replication Agent and Replication Server, must have **create object** permission before you can use it to create a replication definition. You must grant this permission manually from Replication Server.

DB2 UDB-to-SAP HANA Database Replication Setup

1. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

```
isql -Username -Ppassword -SSAMPLE_RS
```

2. Enter:

```
grant create object to SAMPLE_RS_ra  
go
```

where *SAMPLE_RS_ra* is the value for the **rs_username** parameter you entered in the resource file.

Configuring Replication Agent

Review the procedures for configuring Replication Agent for DB2 UDB.

Primary DB2 UDB Configuration for Replication

To configure the primary DB2 UDB database for replication, you may need to connect to the primary database using the primary database instance owner user.

Verifying the Current Archive Setting of the Transaction Log

Set the primary DB2 UDB database transaction logging to archive logging, instead of circular logging for Replication Agent.

1. Determine the LOGARCHMETH1 setting.

a. Connect to the database by entering this command at the CLP prompt:

```
CONNECT TO dbalias USER db2_user USING db2_user_ps
```

where *dbalias* is the cataloged alias of the primary database, *db2_user* is the primary database user, and *db2_user_ps* is the password.

b. Determine the LOGARCHMETH1 setting:

```
GET DB CFG FOR dbalias
```

2. If the results do not show that LOGARCHMETH1 is set to LOGRETAIN or to the path name of the directory to which logs are archived, set it:

- To use the default archive location:

```
UPDATE DATABASE CONFIGURATION USING logarchmeth1 LOGRETAIN
```

- To use a specific archive location:

```
UPDATE DATABASE CONFIGURATION USING logarchmeth1 DISK:path
```

where *path* is the full path name of the directory where the archive logs are to be stored.

If you change the setting of the DB2 UDB **logarchmeth1** parameter, you must back up the database:

```
BACKUP DATABASE dbname TO "backup_dir"
```

where *dbname* is your database name, and *backup_dir* is the target directory for the backup.

Setting the DB2 UDB Connectivity Autocommit Parameter

Replication Agent requires the DB2 UDB connectivity **autocommit** parameter to be on.

The **autocommit** parameter is specified in the DB2 UDB call-level interface (CLI) configuration file for the primary database. If the **autocommit** parameter is not turned on, a deadlock problem may occur.

1. Make a copy of the sample `db2cli.ini` file located in `<HOME>\db2_software\cfg`, where `<HOME>` is the home directory of the DB2 UDB instance owner (UNIX or Linux) or the DB2 UDB client or server installation directory (Windows). Place this copy in `<HOME>\sqlllib\cfg`.
2. In the `db2cli.ini` file in `<HOME>\sqlllib\cfg`, set **autocommit=1**.

Creating a DB2 UDB User and Granting Permissions

Create a DB2 UDB user and grant the necessary permissions.

Replication Agent requires a DB2 UDB login that has permission to access data and create new objects in the primary database. The DB2 UDB login must have SYSADM or DBADM authority to access the primary database transaction log.

Note: You can also use these steps to create a maintenance user.

1. Create a new operating system user named *ra_user* using commands appropriate for your operating system.

For example, to create a user named *ra_user* on the UNIX and Linux operating systems, use:

```
useradd -gusers -Gmgmt -s/bin/shell -psybase
-d/home/ra_user -m ra_user
```

where *sybase* is the password corresponding to the *ra_user* user name.

2. Start the DB2 UDB command-line processor:

```
%>bash
%>source /db2home/db2inst1/sqlllib/db2profile
%>db2
```

3. Connect to the primary DB2 UDB database:

```
db2=>connect to pdb user db2_admin_user using db2_admin_password
```

where *db2_admin_user* and *db2_admin_password* are the administrative user ID and password for the primary database.

4. Grant all necessary authorities to *ra_user*:

```
db2=>grant DBADM on database to user ra_user
```

Adding a Temporary Tablespace to the Primary Database

Add a temporary tablespace to the primary database.

1. Start the DB2 UDB command-line processor:

```
%>bash
%>source /db2home/db2inst1/sqllib/db2profile
%>db2
```

2. Connect to the primary DB2 UDB database:

```
db2=>connect to pdb user db2_admin_user using db2_admin_password
```

where *db2_admin_user* and *db2_admin_password* are the administrative user ID and password for the primary database.

3. Create a buffer pool:

```
db2=>create bufferpool pdb_buffer_name size 1000 pagesize
same_pagesize_as_pdb
```

where *pdb_buffer_name* is the buffer name, and *same_pagesize_as_pdb* is the page size at the primary database.

4. Create a temporary tablespace:

```
db2=>create user temporary tablespace pdb_temp_space_name
pagesize same_pagesize_as_pdb managed by automatic storage
bufferpool
pdb_buffer_name
```

where *pdb_temp_space_name* is the tablespace name, *same_pagesize_as_pdb* is the page size at the primary database, and *pdb_buffer_name* is the buffer name.

Note: Determine the DB2 UDB page size using the **LIST TABLESPACES SHOW DETAIL** command.

For example, to create a temporary tablespace named *deep13* with a 16KB page size and buffer pool named *tom_servo*, enter:

```
create user temporary tablespace deep13 pagesize 16K managed
by automatic storage bufferpool tom_servo
```

Creating the Replication Agent Instance

Prepare and create a Replication Agent instance for replication of each DB2 UDB database.

A single installation of the Replication Agent can support replication from multiple databases. However, one Replication Agent instance is needed for each DB2 UDB database that is to be replicated.

Note: To replicate from multiple databases, using one Replication Agent instance for each database, you must create a separate DB2 UDB user for each primary database. You cannot use the same user for multiple Replication Agents.

1. Locate the resource file template.

The majority of configuration values required to create and initialize a Replication Agent can be recorded and stored in a resource file. Using a resource file provides a way to record or retain the configuration information for a Replication Agent instance, allowing an instance to be removed and re-created.

The resource file template for a DB2 UDB instance is `$$SYBASE/RAX-15_5/init/ibmudb.rs` (UNIX or Linux) or `%SYBASE%\RAX-15_5\init\ibmudb.rs` (Windows).

The resource file template for a DB2 UDB instance is `$$SYBASE/RAX-15_2/init/ibmudb.rs` (UNIX or Linux) or `%SYBASE%\RAX-15_2\init\ibmudb.rs` (Windows).

2. Create an instance resource file:

- Copy the resource file template to another file that contains the configuration information for a new instance, for example, `myra.rs`.
- Change the values that are assigned to properties in the resource file to match the values required for your environment. Most of these values are host, port, user, and password information for the DB2 UDB, Replication Server, and RSSD connections.
 - Set the value of `rs_source_ds` to the value you chose for pds in previous procedures.
 - Set the value of `rs_source_db` to the value you chose for pdb in previous procedures.
- When you create a Replication Agent instance with a resource file:
 - The `asa_password` configuration parameter value cannot contain single quotes, double quotes, or a semicolon.
 - The `pds_username` and `pds_password` configuration parameter values cannot contain single or double quotes if the `create_pds_username` parameter is set to yes.

The resource file is self-documenting, with each parameter described. See the *Replication Agent Administration Guide*.

Note: This table lists parameters in the order they appear in the resource file, not alphabetically.

Table 13. Resource File Parameters

Parameter	Description	Example Values
<code>instance name</code>	Any valid name.	myra
<code>admin_port</code>	Port number that Replication Agent uses.	9030 (if in use, select a different port number)
<code>ltm_admin_user</code>	User name for administering the Replication Agent instance.	sa

DB2 UDB-to-SAP HANA Database Replication Setup

Parameter	Description	Example Values
ltm_admin_pw	Password for the user administering the Replication Agent instance.	The default password policy is that the password must be between 6 and 255 characters. For example: sa_pass
pds_host_name	Machine (host) where DB2 UDB is installed.	NY
pds_port_number	Port number for DB2 UDB. Make sure you know which port your DB2 UDB instance is using. You can find this information in the <code>/etc/services</code> file.	50007
pds_database_name	Name of the primary database.	NYudb9
pds_username	User ID that Replication Agent uses to access primary data server.	ra_user
pds_password	Password for pds_username .	sybase
pds_datasource_name	Datasource name or database alias for the primary database. By default, the alias may be the same name as the database name.	NYudb9
rs_host_name	Machine where Replication Server is installed.	jdoe_host1
rs_port_number	Port number where Replication Server is installed.	11752
rs_username	Replication Server user with connect source and create object permissions.	SAMPLE_RS_ra
rs_password	Password for rs_username .	SAMPLE_RS_ra_ps
rs_source_ds	Valid name representing data server of the primary DB2 UDB database. Note: If you use direct load materialization, add an entry with this name in the interfaces file.	pds
rs_source_db	Valid name representing the primary DB2 UDB database.	pdb

Parameter	Description	Example Values
rs_charset	Character set that Replication Server is using. The value defined for the rs_charset configuration parameter must match the RS_charset value in the Replication Server configuration file, <code>\$SYBASE/REP-15_5 / install/<server>.cfg</code> (UNIX or Linux) or <code>%SYBASE%\REP-15_5 \install\<server>.cfg</code> (Windows).	<ul style="list-style-type: none"> Windows: cp850 UNIX: iso_1
rssd_host_name	Machine where RSSD resides.	jdoe_host1
rssd_port_number	Port number where RSSD resides.	11751
rssd_database_name	Database name for RSSD.	SAMPLE_RS_ERSSD
rssd_username	Valid user for RSSD.	SAMPLE_RS_RSSD_maint
rssd_password	Password for rssd_username .	SAMPLE_RS_RSSD_maint_ps
pdb_archive_remove	Enables or disables the removal of archived DB2 UDB transaction log files from the path specified by pdb_archive_path .	USE_DEFAULT
pdb_archive_path	Directory path where Replication Agent expects to find archived DB2 UDB transaction log files.	A valid directory path on the machine hosting Replication Agent that points to a location where DB2 UDB puts the archived transaction log files
start_instance	Start the instance that was created.	USE_DEFAULT
initialize_instance	Initialize the Replication Agent instance.	USE_DEFAULT

Warning! The **rs_source_ds** and the **rs_source_db** values must match the **pds.pdb** values of your Replication Server primary connection name.

- Set the DB2 UDB environment variables by sourcing `db2cshrc`, `db2profile`, or your customized versions of the scripts.
- Set the environment variables by sourcing the `SYBASE.csh` file (UNIX or Linux) or by executing the `SYBASE.bat` file (Windows).
- Create and execute the new instance resource file:
 - Validate the settings in the resource file using the **-vr** parameter.

- On Windows:

```
%SYBASE%\RAX-15_5\bin\ra_admin.bat -vr myra.rs
```

DB2 UDB-to-SAP HANA Database Replication Setup

```
%SYBASE%\RAX-15_2\bin\ra_admin.bat -vr myra.rs
```

- On UNIX or Linux:

```
$SYBASE/RAX-15_5/bin/ra_admin.sh -vr myra.rs
```

```
$SYBASE/RAX-15_2/bin/ra_admin.sh -vr myra.rs
```

where *myra.rs* is the path and name of the resource file.

Validation results are returned as either:

- Response-file processing completed, or,
- Response-file processing completed with errors

If any validation fails, the **ra_admin** utility returns an error message and information about the failure. You can repeat the validation process as many times as necessary until it executes without error. No entities are changed or created.

- b) After the resource file has been validated, allow the **ra_admin** utility to create and configure the Replication Agent instance, using the **-r** parameter.

- On Windows:

```
%SYBASE%\RAX-15_5\bin\ra_admin.bat -r myra.rs
```

```
%SYBASE%\RAX-15_2\bin\ra_admin.bat -r myra.rs
```

- On UNIX or Linux:

```
$SYBASE/RAX-15_5/bin/ra_admin.sh -r myra.rs
```

```
$SYBASE/RAX-15_2/bin/ra_admin.sh -r myra.rs
```

where *myra.rs* is the path and name of the resource file.

Note: If **start_instance** is **yes** in your response file, your instance is also running. If **initialize_instance** is **yes**, your instance is also initialized.

- c) Creation results are returned as either:

- Response-file processing completed
- or,
- Response-file processing completed with errors

See the *Replication Agent Administration Guide*.

6. If **start_instance** is **yes** in your response file, skip this step, and continue with the next section. If **start_instance** is **false** in your response file, change to the instance directory and run Replication Agent in the background.

- On Windows:

```
cd %SYBASE%\RAX-15_5\myra
```

- On UNIX or Linux:

```
cd $SYBASE/RAX-15_5/myra
```

Execute the `RUN_myra` file in the background.

Your Replication Agent for DB2 UDB is now running.

Verifying the Replication Agent Instance Installation

Verify that the Replication Agent instance has been successfully installed.

1. Use **dsedit** to update the Replication Server `sql.ini` (Windows) or `interfaces` (UNIX or Linux) file to include an entry for the Replication Agent location.

Note: You can use any Tabular Data Stream™ (TDS) client utility (**isql**, **isqlApp**, or **SQLAdvantage**) that you prefer.

To use direct load materialization, also add the server name part of the connection name to the `interfaces` file. This entry should point to Replication Agent. For example:

```
pds
master tcp ether hostname 9030
query tcp ether hostname 9030
```

Note: The server name is the same value as the **rs_source_ds** value in the resource file used to create a Replication Agent instance.

2. Verify the connection to Replication Agent:
 - a) Open a command window in the `%SYBASE%` (Windows) or `$SYBASE` (UNIX or Linux) directory of your Replication Agent installation.
 - b) Set the environment variables by executing the `SYBASE.bat` file (Windows) or by sourcing the `SYBASE.csh` file (UNIX or Linux).
 - c) Log in to Replication Agent:

```
isql -Usa -Psa_pass -Smyra
```

Note: These verification steps are optional because they were performed when you verified the resource file.

3. Verify the Replication Agent connection to Replication Server: Verify the Replication Agent connection to Replication Server by entering:

```
test_connection RS
go
```

This result is returned:

```
Type Connection
-----
RS succeeded
(1 row affected)
```

If the result indicates a failure, either the server is not responding or the connection properties (*host*, *port*, *user*, or *password*) are incorrect. Verify the host and port configuration values, and manually log in to the Replication Server as the configured user to determine which property is incorrectly defined.

See the *Replication Agent Administration Guide*.

- a) Enter:

```
test_connection RS
go
```

b) This result is returned:

```
Type Connection
-----
RS succeeded
(1 row affected)
```

c) If the result indicates a failure, either the server is not responding or the connection properties (*host*, *port*, *user*, or *password*) are incorrect. Verify the host and port configuration values, and manually log in to the Replication Server as the configured user to determine which property is incorrectly defined.

See the *Replication Agent Administration Guide*.

4. Verify the Replication Agent connection to the primary database:

a) Enter:

```
test_connection PDS
go
```

You see:

```
Type Connection
-----
PDS succeeded
(1 row affected)
```

b) If the result indicates a failure:

- The server is not responding, or
- One of the connection properties is incorrect.

Manually log in to the primary DB2 UDB database as the configured user to find which property is incorrectly defined.

See *Testing Network Connectivity* in the *Replication Agent Administration Guide*.

See also

- *Initializing the Replication Agent Instance* on page 116

Initializing the Replication Agent Instance

Initialize the Replication Agent instance.

The **ra_admin init** command verifies that the primary DB2 UDB database is correctly configured to provide archived logging and that the DB2 UDB user ID used by the Replication Agent has the necessary permissions. The **ra_admin init** command also creates objects in the primary database to support replication.

1. Initialize the Replication Agent instance:

```
ra_admin init
go
```

You see a message indicating that the procedure was successful.

2. Enter:

```
resume
go
```

If the Replication Agent successfully transfers to a replicating state, you see:

```
State          Action
-----
REPLICATING   Ready to replicate data.
```

3. The **ra_status** command returns the state of the Replication Agent. It is good practice to verify that the Replication Agent remains in replication state, even after the resume command executes successfully.

To detect an error that occurred after replication start-up, execute:

```
ra_status
go
```

If the Replication Agent is in replicating state, **ra_status** returns:

```
State          Action
-----
REPLICATING   Ready to replicate data.
```

4. Validate that both primary and replicate connections are active:

```
isql -Usa -Psa_pass -SSAMPLE_RS
admin who
go
```

Note: Be aware that:

- The DSI connection for the primary database connection is usually suspended because you are not replicating data back to the primary database.
 - The Replication Agent connection, if established for the replicate database connection, is usually down, because you are not replicating data from the replicate database.
-

Do not proceed until **admin who** returns thread status similar to:

```
Spid Name          State          Info
-----
13  DSI EXEC   Awaiting Command 101(1) SAMPLE_RS_ERSSD.
SAMPLE_RS_ERSSD
9   DSI       Awaiting Message 101 SAMPLE_RS_ERSSD.
SAMPLE_RS_ERSSD
8   SQM       Awaiting Message 101:0 SAMPLE_RS_ERSSD.
SAMPLE_RS_ERSSD
54  DSI EXEC   Awaiting Command 102(1) DCService.rdb
53  DSI       Awaiting Message 102 DCService.rdb
17  SQM       Awaiting Message 102:0 DCService.rdb
   DSI EXEC   Suspended        103(1) pds.pdb
   DSI       Suspended        103 pds.pdb
24  DIST     Awaiting Wakeup 103 pds.pdb
25  SQT      Awaiting Wakeup 103:1 DIST pds.pdb
23  SQM      Awaiting Message 103:1 pds.pdb
```

```
22    SQM           Awaiting Message 103:0 pds.pdb
62    REP AGENT    Awaiting Command pds.pdb
```

Testing Replication

When you finish setting up the replication system, test that replication works as intended.

1. Connect to the primary DB2 UDB instance as a regular user, rather than as the maintenance user. Make sure the regular user also exists in the replicate database.

a) To connect to the database with CLP, use:

```
CONNECT TO dbalias USER db2_user USING db2_user_ps
```

where *dbalias* is the cataloged alias of the primary database, *db2_user* is the primary database user, and *db2_user_ps* is the password.

2. At the primary DB2 UDB database, create a table named PTAB1 to replicate:

```
CREATE TABLE PTAB1
(IDNO INTEGER NOT NULL,
NAME VARCHAR(20),
PRIMARY KEY(IDNO))
```

3. Grant permissions to any new or existing object to be replicated in the primary database:

```
grant all on PTAB1 to public
```

4. Connect to Replication Agent through **isql**, and mark the PTAB1 table for replication:

```
pdb_setreptable PTAB1, mark
go
```

5. Add rows to the ptab1 table:

```
insert into ptab1 values (1, 'Burt Gringlesby');
insert into ptab1 values (2, 'Douglas Wong');
insert into ptab1 values (3, 'Tom Servo');
insert into ptab1 values (4, 'Innes del Castillo');
insert into ptab1 values (5, 'Akiko Yokomoto');
commit;
```

6. In Replication Agent, create a replication definition against the primary IBM DB2 UDB database:

```
rs_create_repdef ptab1
go
Table/Procedure Name      RepDef Name      Status
-----
<DB2UDBTableOwner>.PTAB1 "<repdefname>" Created
(1 row affected)
```

Record the name of the replication definition for use later.

7. At the replicate SAP HANA database instance, create a table named PTAB1:

```
CREATE TABLE <tableowner>.PTAB1
(IDNO INT PRIMARY KEY,
NAME VARCHAR(20));
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

If an owner is not specified with the table name, the owner of the table—*<tableowner>*—is the user who is logged in when the table is created.

8. Grant permissions to any new or existing object to be replicated in the replicate database so that the Replication Server maintenance user can update this table:

```
grant all privileges on <tableowner>.PTAB1 to public
```

9. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

10. Alter the replication definition that you created earlier to point to the correct replicate table:

```
alter replication definition <repdefname>
with replicate table named <tableowner>.ptab1
go
```

Note: SAP HANA database is not case-sensitive and uses uppercase characters if lowercase characters are provided.

11. Create a subscription against the replicate connection to materialize the PTAB1 table on the SAP HANA database.

Create the *ptab1_sub* subscription with **direct_load** option for the *<repdefname>* replication definition. Replication Server logs in to Replication Agent for IBM DB2 UDB, and Replication Agent issues a **SELECT** command to select data from the IBM DB2 UDB database, enter:

```
create subscription ptab1_sub
for <repdefname>
with replicate at rds.rdb
without holdlock
direct_load
user puser password ppwd
go
```

where:

- *rds.rdb* is the replicate SAP HANA database connection created earlier.
- *puser* is the user who connects to Replication Agent for IBM DB2 UDB.
- *ppwd* is the password of the *puser*. You must enter a password.

Note: The interfaces file used by the replicate Replication Server must contain an entry for the Replication Agent for IBM DB2 UDB that is used to select data from the IBM DB2 UDB database.

To create the subscriptions at the replicate Replication Server using the no materialization method, use the **create subscription** command with the **without materialization** clause.

For information about using the **create subscription** command in **direct_load** materialization, see *Replication Server Reference Manual*.

12. Check the subscription status to ensure there are no errors:

```
check subscription ptabl_sub
for <repdefname>
with replicate at rds.rdb
go
```

If there are errors, see *Troubleshooting Subscription Problems*.

13. Connect to the IBM DB2 UDB database as a non-maintenance user with the ability to insert rows into the table:

```
insert into ptabl values (10, 'Michel DeFrance');
insert into ptabl values (11, 'Dirk Stringer');
commit;
```

14. On the Replication Server, check the status of the subscription created with the **direct_load** option by using the **check subscription** command. When the subscription is in the VALID state, materialization is complete:

```
check subscription ptabl_sub
for <repdefname>
with replicate at rds.rdb
go
Subscription ptabl_sub is VALID at the replicate.
```

15. After the status is VALID at the primary and replicate database, on the replicate database, verify that the data is replicated to the replicate PTABL table:

```
SELECT * FROM <tableowner>.PTABL
```

Resetting the Primary DB2 UDB Database for Replication

In a test environment, you may occasionally need to reset the replication environment.

1. Log in to the Replication Agent and issue:

```
isql -Usa -Psa_pass -Smyra
suspend
go
```

2. Reset the locator stored in the Replication Agent:

```
ra_locator zero
go
```

3. Reset the locator stored in Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
connect
go
rs_zeroltm pds, pdb
go
exit
```

You see:

```
Locator has been reset to zero.  
(return status = 0)
```

4. Resume replication with the Replication Agent **resume** command:

```
isql -Usa -Psa_pass -Smyra  
1> resume  
2> go
```


SAP Business Suite-to-SAP HANA Database Replication Setup

Implement a sample replication system from SAP Business Suite application running on Oracle, Microsoft SQL Server, or IBM DB2 UDB into SAP HANA database.

Note: If your primary and replicate databases are not already participating in replication using SAP® Replication Server®, use SAP Replication Management Agent Designer to setup the replication environment. See *Replication Management Agent Configuration and Users Guide*.

Use the tasks in this document to set up a proof-of-concept or test replication system.

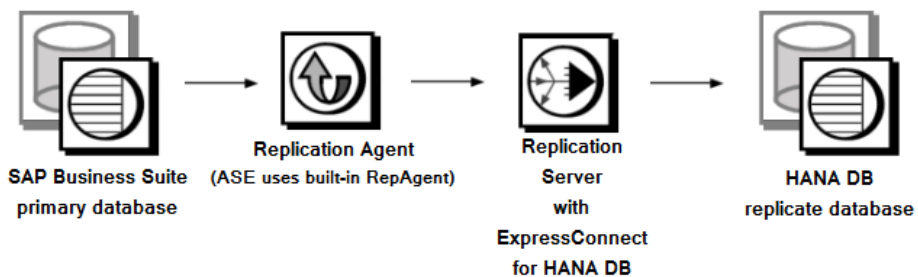
The tasks use the sample Replication Server instance SAMPLE_RS. If you have already installed Replication Server, skip to the configuration procedures. The procedures to install and configure the replication components are the same as for other supported databases, except for the procedure to configure Replication Server.

For the most current information about Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

Multiple components are required to implement an SAP Business Suite application data-to-SAP HANA database replication system: a primary data server, a replicate SAP HANA database, a Replication Server with ExpressConnect for SAP HANA Database, and Replication Agent for Oracle, Microsoft SQL Server, or IBM DB2 UDB.

Figure 5: SAP Business Suite-to-SAP HANA Database Replication Setup



Prerequisites

Review the prerequisites before you begin setting up SAP Business Suite-to-SAP HANA database replication.

Make sure:

- You are familiar with the SAP Business Suite primary database and SAP HANA database.
- You have valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- SAP HANA version is 1.00.62 or above.
- SAP Kernel is 720, 721, 740, 741.

Note: If your SAP Business Suite application is running on Microsoft SQL Server, make sure that you use the SAP Kernel described in the SAP Note 1972365 – Retry on error 4711: <https://css.wdf.sap.corp/sap/support/notes/1972365>.

- The library, `libSQLDBCHDB.dll` (Windows) or `libSQLDBCHDB.so` (UNIX), Revision 1.00.72.00 or above, is downloaded and installed.
- The library, `libodbchdb.dll` (Windows) or `libodbchdb.so` (UNIX), Revision 1.00.72.00 or above, is downloaded and installed.
- TCP/IP connectivity is available between all hosts.
- Any operating system patches required for Java have been installed.

Planning Your Installation

To verify the system requirements, identify the installation directories, and obtain the installation software, refer to the section specific to the primary database for your SAP Business Suite application.

See:

- *Planning Your Installation for ASE-to-SAP HANA Database Replication*
- *Planning Your Installation for Oracle-to-SAP HANA Database Replication*
- *Planning Your Installation for Microsoft SQL Server-to-SAP HANA Database Replication*
- *Planning Your Installation for IBM DB2 UDB-to-SAP HANA Database Replication*

Downloading Products from SMP

Download replication components from the SAP Service Marketplace (SMP).

Prerequisites

An SMP login ID and password.

Task

1. Go to <https://websmp204.sap-ag.de/patches>, and enter your login ID and password.
2. Select **SAP Software Download Center > Installations and Upgrades > A-Z Index** and click **R** for SAP Replication Server products.
3. From the product listing, select **Sybase Replication Server, HANA Edition**.
4. Select the software you need by version and platform, and follow the instructions.

Installing Replication Components

To review the procedures for installing the replication components, see the section specific to the primary database of your SAP Business Suite application.

See:

- *Installing Replication Components for ASE-to-SAP HANA Database Replication*
- *Installing Replication Components for Oracle-to-SAP HANA Database Replication*
- *Installing Replication Components for Microsoft SQL Server-to-SAP HANA Database Replication*
- *Installing Replication Components for IBM DB2 UDB-to-SAP HANA Database Replication*

Configuring Replication Components

Review the procedures for configuring each of the replication components in the replication system.

Configuring Replication Server

Configure Replication Server to perform replication from SAP Business Suite applications running on supported databases into SAP HANA database.

Creating a Replication Server Connection to the Primary Database

Create a Replication Server connection to the primary Oracle, Microsoft SQL Server, or IBM DB2 UDB database. If your primary database is Adaptive Server, you do not need to create a Replication Server connection.

1. In Replication Server, create a connection to the primary database:

```
create connection to pds.pdb
using profile connection_profile;standard
set username muser
set password mpwd
with log transfer on, dsi_suspended
go
```

where:

- *pds* is the value of the **rs_source_ds** parameter specified in Replication Agent.
- *pdb* is the value of **rs_source_db** specified in Replication Agent.
- *connection_profile* is **rs_rs_to_oracle_ra** for Oracle, **rs_rs_to_msss_ra** for Microsoft SQL Server, and **rs_rs_to_udb_ra** for IBM DB2 UDB.
- *muser* is the maintenance user for the primary Oracle database. This user must already exist and have the necessary select permissions in the primary database. See the *Replication Server Options Primary Database Guide*.
- *mpwd* is the maintenance user password.

For information about the maintenance user, see the *Replication Server Heterogeneous Replication Guide > SAP Replication Products > Replication Server > Database Connections > Maintenance User Purpose*.

2. Use **admin show_connections, 'primary'** to display the primary connection you created.

Obtaining the SAP HANA Database ODBC Drivers

Set up ExpressConnect for HANA DB for connections between SAP HANA database and Replication Server.

You must add SAP HANA database ODBC driver to the ExpressConnect for SAP HANA Database installation. If you are replicating cluster tables, add these libraries to the ExpressConnect for SAP HANA Database installation:

- The library, **libSQLDBCHDB.dll** (Windows) or **libSQLDBCHDB.so** (UNIX), Revision 1.00.72.00 or above.
- The library, **libodbcHDB.dll** (Windows) or **libodbcHDB.so** (UNIX), Revision 1.00.72.00 or above.

You can download these libraries from the SAP Service Marketplace and install them after you have installed Replication Server.

See the *SAP Replication Server Release Bulletin > Special Installation Instructions > Installing ODBC Libraries for ExpressConnect for SAP HANA Database*.

Configuring Replication Server for Replication to SAP HANA Database

Configure Replication Server for replication from SAP Business Suite running on supported databases to SAP HANA database.

If you are not using the sample Replication Server instance, enter your Replication Server instance name in place of SAMPLE_RS.

1. Add an entry to your Replication Server interfaces file identifying the replicate SAP HANA database, and then stop and start Replication Server:

```
[dataservername]
master tcp ether hostname port
query tcp ether hostname port
```

where *hostname* and *port* are the host and port number of the SAP HANA database, and *dataservername* is a label used to identify the host and port number. For information about stopping and starting Replication Server, see the *Replication Server Administration Guide Volume 1 > Manage a Replication System*.

Note: Each SAP HANA database installs with a unique instance number, and the port number is the instance number prefixed with 3 and suffixed with 15:

```
3in15
```

where *in* is the two-digit instance number. For example, the port number for an SAP HANA database with instance number 1 is 30115.

2. On SAP HANA database, create a maintenance user:
 - a) As the SAP HANA database admin user, create a user ID with a temporary password:

```
CREATE USER muser PASSWORD mpwd_temp
```

- b) Log in to SAP HANA database with the temporary password, and then update the password:

```
ALTER USER muser PASSWORD mpwd_new
```

- c) Log out of SAP HANA database, and log in again with the new password.

3. Log in to SAP HANA database:

```
./hdbsql -u user -p password -i id
```

where *user* is the SAP HANA database user, *password* is the user password, and *id* is the instance number.

4. On SAP HANA database, grant these authorities to the maintenance user:

```
GRANT CREATE ANY, DELETE, DROP, EXECUTE,
INDEX, SELECT, UPDATE ON SCHEMA myschema TO muser
```

SAP Business Suite-to-SAP HANA Database Replication Setup

where *myschema* is your SAP HANA database schema, and *muser* is the maintenance user you just created.

5. Log in to Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

If you have not set up the sample Replication Server instance, enter your Replication Server instance name instead of `SAMPLE_RS`.

6. Create a connection to the replicate SAP HANA database instance using ExpressConnect for SAP HANA Database.

For a standard connection:

```
create connection to rds.rdb
using profile connection_profile;ech
set username muser
set password mpwd
go
```

where:

- *rds* is the replicate SAP HANA database. For a standard connection, this must match the data server name in the interfaces file entry.
- *rdb* is a placeholder; you must provide a value, but it is not used.
- *connection_profile* is:

Primary Database	Connection Profile Value
Oracle	rs_oracle_to_hanadb
Microsoft SQL Server	rs_msss_to_hanadb
IBM DB2 UDB	rs_udb_to_hanadb

- *muser* is the maintenance user for the replicate SAP HANA database instance you have already created.
- *mpwd* is the replicate SAP HANA database maintenance user password.
- *auser* and *apwd* are unused values supplied only to satisfy the syntax of the **create connection** command.

If you have trouble connecting to the replicate SAP HANA database instance using ExpressConnect for SAP HANA Database, see the *Replication Server Troubleshooting Guide > Common Error Messages > Connector Error Messages*.

7. Use **admin show_connections, 'replicate'** to display the replicate connection you created.

8. Enable the replication of trailing zeros in varbinary values:

```
configure replication server set varbinary_strip_trailing_zeros
to 'off'
```

9. Set the value of \$SAPSYSTEMNAME environment variable for the SAP HANA database <sid>adm user to the **sapsystemname** configuration parameter:

```
alter connection to rds.rdb
set sapsystemname to 'sid_adm_value'
```

where:

- *rds.rdb* is the connection name for the SAP HANA database.
- *sid_adm_value* is the name of the SID (SAP SYSTEM NAME).

For example, the SYSTEM NAME is *D01*. You must enter:

```
alter connection to rds.rdb set sapsystemname to 'D01'
```

10. Map primary schema to replicate schema:

```
alter connection to rds.rdb
set schemamap with decluster from pds.pdb.primary_schema to
HANA_schema
```

where:

- *rds.rdb* is the connection name for the SAP HANA database.
- *pds.pdb* is the is the connection name used in the **create connection** command for primary database.
- *primary_schema* is the schema name of the primary database.
- *HANA_schema* is the schema name of the SAP HANA database.

11. Connect to Replication Agent through **isql**, and mark these SAP Data Dictionary (DD) tables—DD16S, DD03L, and DD02L—for replication:

```
pdb_setreptable primary_schema.tablename, mark, owner
go
```

where, *primary_schema.tablename* is the SAP table name.

Note: Make sure that you do not mark the cluster tables CLU4 and VER_CLUSTR for replication. For more information, see SAP Note 902817 – Inconsistencies in cluster tables (log vs phys field name): <https://css.wdf.sap.corp/sap/support/notes/902817>

12. In Replication Agent, create a replication definition for each SAP DD table:

```
rs_create_repdef primary_schema.tablename
go
```

where, *primary_schema.tablename* is the name of the SAP DD table.

13. Create a subscription against the replicate connection to materialize the SAP DD tables on the SAP HANA database.

```
create subscription sub_name
for table_repdef
with replicate at rds.rdb
without holdlock
direct_load init replicate table
with create_or_truncate
user username password pass
go
```

where:

- *sub_name* is the name of the subscription.
- *table_repdef* is the table replication definition the subscription *sub_name* is for.

SAP Business Suite-to-SAP HANA Database Replication Setup

- *rds.rdb* is the replicate SAP HANA database connection created earlier.
- *username* is the user who connects to Replication Agent for your primary database.
- *pass* is the password of the *username*. You must enter a password.

Note: You can also use other **create subscription** options such as **create**, **recreate**, or **truncate**.

14. (Only for SAP Business Suite running on Oracle database) Set **sap_trim_len** to **on** to trim the first two bytes for certain SAP columns such as RAW or LRAW whose DDIC type length is greater than 255 bytes:

```
alter connection to data_server.database
set sap_trim_len to 'on'
```

15. (Only for SAP Business Suite running on Oracle database or Microsoft SQL Server) Set **reblock_ddntf** to **off** to allow Replication Server to disable the FIELDS column length adjustment in the DDNTF table:

```
alter connection to rds.rdb
set reblock_ddntf to 'off'
```

16. Connect to Replication Agent through **isql**, and mark these SAP tables—DDNTT, DDNTF, TADIR, SVERS, DD09L, DD12L, and DD17S:

```
pdb_setreptable primary_schema.tablename, mark, owner
go
```

where, *primary_schema.tablename* is the SAP table name.

17. In Replication Agent, create a replication definition for each SAP table that you marked for replication:

```
rs_create_repdef primary_schema.tablename
go
```

where, *primary_schema.tablename* is the name of the SAP table.

18. Create a subscription against the replicate connection to materialize the SAP tables on the SAP HANA database.

```
create subscription sub_name
for table_repdef
with replicate at rds.rdb
without holdlock
direct_load init replicate table
with create_or_truncate
user username password pass
go
```

where:

- *sub_name* is the name of the subscription.
- *table_repdef* is the table replication definition the subscription *sub_name* is for.
- *rds.rdb* is the replicate SAP HANA database connection created earlier.
- *username* is the user who connects to Replication Agent for your primary database.
- *pass* is the password of the *username*. You must enter a password.

Note: You can also use other **create subscription** options such as **create**, **recreate**, or **truncate**.

19. Create table replication definition and subscription with materialization for all other SAP tables.

Configuring Replication Server for Replication from the Primary Database

To configure Replication Server for replication from the primary database, see the section specific to the primary database for your SAP Business Suite application.

See:

- *Configuring Replication Server for Replication from ASE*
- *Configuring Replication Server for Replication from Oracle*
- *Configuring Replication Server for Replication from Microsoft SQL Server*
- *Configuring Replication Server for Replication from IBM DB2 UDB*

Configuring Replication Agent

To review the procedures for configuring Replication Agent, see the section specific to the primary database for your SAP Business Suite application.

See:

- *Configuring Replication Agent for Oracle*
- *Configuring Replication Agent for Microsoft SQL Server*
- *Configuring Replication Agent for IBM DB2 UDB*
- *Configuring Replication Server for Replication from IBM DB2 UDB*

SAP HANA-to-SAP HANA Database Replication Setup

Setup replication from SAP HANA database to SAP HANA database using the Replication Agent™ for SAP HANA®.

Replication Agent for SAP HANA is a lightweight server that performs initial data extraction and load, and also replicates data changes on the primary SAP HANA server to the replicate SAP HANA server. See *Replication Agent for SAP HANA Configuration and Users Guide*.

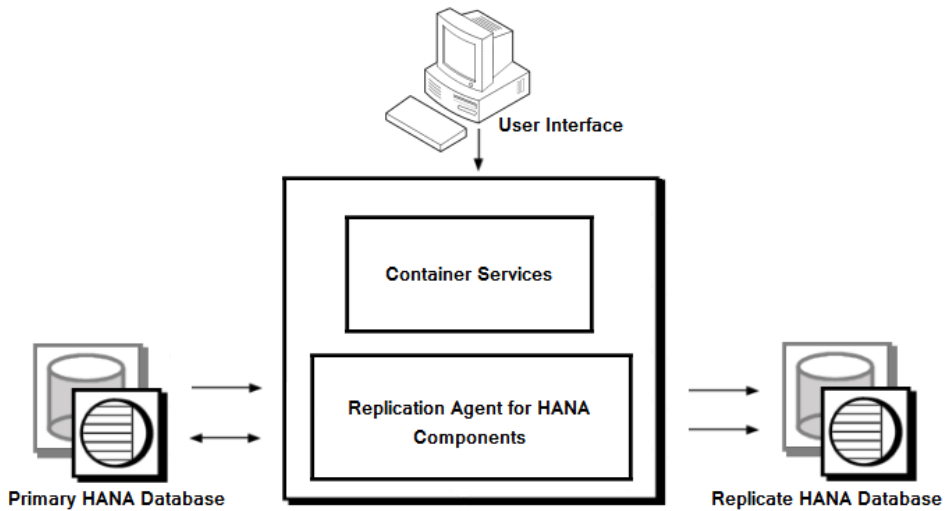
For the most current information about Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

Replication Agent for SAP HANA is an application that runs inside a Replication Management Agent (RMA) container.

The RMA container is a server that provides many services, such as a management application programming interface (API), a logging service, and a repository for metadata. The container may run applications, such as Replication Agent for SAP HANA, that consume these services, or may simply be used as an agent for communication with other containers. You can connect multiple containers to form a peer-to-peer grid. When you begin to replicate data from SAP HANA database, you must first start a container that contains the Replication Agent for SAP HANA components.

Figure 6: SAP HANA-to-SAP HANA Database Replication Setup



These RMA components are used by Replication Agent for SAP HANA:

- **Command API** – is an API that executes commands for the Replication Agent for SAP HANA. For example, **Start Replication** returns results that depend on the command executed.
- **Builder** – builds shadow tables and other control tables. For example, the trigger queue replicates data at the primary database, builds user tables, and sends change data capture (CDC) tables to the replicate database.
- **Scanner** – reads data from the shadow tables on the source database.
- **Applier** – writes data to the CDC tables using one of the CDC modes and the replicate database user-defined tables when a non-CDC mode is enabled.

You can use the Replication Agent in the SAP® Control Center graphical user interface (GUI) or from the command line interface (CLI) to issue replication commands. The CLI commands are processed by the Tabular Data Stream™ (TDS) communication protocol in the container.

Support for Multitenancy

Multitenancy allows system administrators to deploy a single instance of the software that runs on a server, serving multiple client organizations known as tenants.

Replication Agent for SAP HANA allows you to deploy multiple replication models to a single container as long as each model has a distinct schema. You can configure, manage, deploy, and undeploy each replication model independently.

Replication Agent for SAP HANA does not support deployment of multiple replication models that have the same schema. For example, you cannot deploy replication models where both read from or write to have the same database schema.

Installing Replication Agent for SAP HANA

Use the Replication Server media to install Replication Agent for SAP HANA.

For more information about installing Replication Agent for SAP HANA, deploying a replication model, and starting replication, see *Replication Agent for SAP HANA Configuration and Users Guide*.

User Management

Replication Server, Replication Agent, and database users.

Table 14. Replication Server Users

User/Login	Server	Client	Usage
rs_user	Replication Server	Replication Server administrators and users	Executes commands on the replication server. The commands executed depend on the security level of the user and include commands such as creating replication definitions, subscriptions, and creating other users.
Replication Server login username rs_username	Replication Server	Replication Server users and Replication Agent	Replication Agent uses the rs_user user to send commands to the Replication Server. For example, replication definition commands and LTL commands. This user must have the permissions to send LTL, create, and drop replication definitions.
Primary database maintenance user	Primary database	Replicate Server	Executes DML at the primary database in the case of bi-directional replication. Use the maintenance user if you do not want the operations at the primary database to be sent to the replicate database. This user cannot be the same as the pds_username user.
Replicate database maintenance user	Replicate database	Replication Server	Executes DDL, DML, and stored procedures at the replicate database. The transactions applied by the maintenance user are not replicated out of the replicate database. This user cannot be the same as the pds_username user.

User/Login	Server	Client	Usage
RSSD primary user rssd_username	Replication Server	Replication Agent	<p>Replication Agent uses this user to login to the Replication Server RSSD for replication definition information.</p> <p>Replication Server provides a gateway for access to the RSSD. Replication Server users other than the Replication Agent can access the RSSD using the rs_username login. This user cannot be the same as the pds_username user.</p>
RSSD maintenance user	Replicate RSSD	Replicate, Replication Server	Applies modifications to replicate RSSDs to coordinate with route information in the primary RSSD.

Table 15. Replication Agent Users

User/Login	Server	Client	Usage
pds_username	Primary database	Replication Agent	<p>Replication Agent uses this user login name to access primary database objects and to create, remove, or manage the transaction log objects in the primary database. The login name cannot be the same as the Replication Agent user for RSSD or the maintenance user of the primary database.</p> <p>For Replication Agents for Oracle, IBM DB2 UDB, and Microsoft SQL Server, the login name selects the data if you use direct load materialization.</p> <p>For Replication Agent for IBM DB2 UDB, this user must have sysdba permission to read DB2 logs and these permissions:</p> <ul style="list-style-type: none"> • SELECT permissions for reading the log. • SELECT permissions on the required system and user tables. <p>pds_username cannot be the same as the primary database maintenance user or the RSSD primary user.</p>

User/Login	Server	Client	Usage
Replication Agent administrator ra_username	Replication Agent	Replication Agent user and administrator	Administers the Replication Agent instance.

Table 16. Database Users

User/Login	Server	Client	Usage
Data server login	Primary database	Primary database user	Data server login names are used with a client application to connect to a data server. This is the login used by the application or user whose activities in the primary database (for example, DML, DDL, and stored procedures) are to be replicated.
Primary database administrative user	Primary database	Primary database user and Replication Agent	Creates other users in the primary database and performs configuration activities such as granting permissions to database objects. The Replication Agent for IBM DB2 UDB requires this user to read the DB2 log. The Replication Agent for Microsoft SQL Server required this user to add system stored procedures.
ASM username asm_username	Primary database	Replication Agent	(Oracle only) Replication Agent uses this user to access the Oracle ASM server.
ddl_username	Replicate database	Replication Server	The ddl_username user has the permissions to perform tasks on behalf of other users. This allows consistent object ownership across primary and replicate database.

Best Practices and Troubleshooting

Best practices for using and troubleshooting common problems that arise while setting up a basic replication system with SAP HANA database as the replicate.

Example Scenarios: Marking Tables for Replication

Scenarios for marking tables for replication in Oracle, IBM DB2 UDB, or Microsoft SQL Server.

- Scenario 1: If the owner and table names are the same at the primary and replicate database, mark the table in the primary database using **pdb_setreptable** with the **owner** mode set to on. You can use **rs_create_repdef** to create the replication definition.

For example:

```

pdb_setreptable T1, mark, owner
go

rs_create_repdef T1
go

```

- Scenario 2: If the primary and replicate tables have the same name but different owners, mark the table in the primary database using **pdb_setreptable** with the **owner** mode set to off. Create a replication definition using **rs_create_repdef** and then alter the replication definition to point to the replicate table.

For example:

```

pdb_setreptable T1, mark
go

rs_create_repdef T1
go

```

In the Replication Server run:

```

isql -Usa -Psa_pass -SSAMPLE_RS
alter replication definition <repdefname>
with replicate table named <replicate tableowner>.<replicate
tablename>
go

```

- Scenario 3: If you are replicating a primary table to a table of the same name, and the owner of the replicate table is the same as the maintenance user of the replicate connection, mark the table in the primary database using **pdb_setreptable** with the **owner** mode set to off. You can use **rs_create_repdef** to create the replication definition.

For example:

```

pdb_setreptable T1, mark
go

```

```
rs_create_repdef T1  
go
```

- Scenario 4: If there are two primary tables with the same schema, each with its own owner and if you are replicating them to the same replicated table of the same table name, mark the table in the primary database using **pdb_setreptable** with the **owner** mode set to off. You can use **rs_create_repdef** to create the replication definition.

For example:

```
pdb_setreptable T1, mark  
go
```

```
rs_create_repdef T1  
go
```

For more information about marking tables for replication with Replication Agent for Oracle, IBM DB2 UDB, or Microsoft SQL Server, see *Replication Agent Administration Guide > Setup and Configuration > Primary Database Object Marking*. For information about creating replication definitions, see *Replication Server Reference Manual*.

Troubleshooting Subscription Problems

Troubleshoot problems that arise if a subscription is not in a valid state, or if the **check subscription** command returns an error message.

1. Check whether there are any error messages in the Replication Server error log file and rectify the problem. If there are no error messages, go to step 2.

By default, Replication Server error log files are in the `$SYBASE/$SYBASE_REP/install` directory. The Replication Server error log file name has this format:

```
rs_name.log
```

where *rs_name* is the Replication Server name. The `SAMPLE_RS` log file name is `SAMPLE_RS.log`, in the `$SYBASE/$SYBASE_REP/samp_repserver` directory.

2. Check whether there are any error messages in the Replication Agent error log file and rectify the problem.
3. Restart the **direct_load** process:

- a. Log in to the Replication Server:

```
isql -Usa -Psa_pass -SSAMPLE_RS
```

- b. Drop the **direct_load** subscription:

```
drop subscription ptab1_sub  
for <repdefname>  
with replicate at rds.rdb  
without purge  
go
```

- c. Ensure that the subscription is dropped:

```
check subscription ptab1_sub  
for <repdefname>
```

```
with replicate at rds.rdb
go
ptabl_sub doesn't exist.
```

- d. Truncate the table on the SAP HANA database to ensure that all the rows loaded from the primary database to the replicate HANA table can be rematerialized.
- e. Re-create the subscription with the **direct_load** option:

```
create subscription ptabl_sub
for <repdefname>
with replicate at rds.rdb
without holdlock
direct_load
user <puser> password <ppwd>
go
```

The subscriptions created with the **direct_load** option materialize subscription data regardless of whether a table is marked and enabled for continuous replication.

Note: If you suspend the SAP HANA connection or drop a subscription during subscription materialization, the **direct_load** materialization process stops.

Glossary

This glossary describes SAP® Replication Server® Options terms.

- **SAP Adaptive Server** – the brand name for SAP relational database management system (RDBMS) software products.
 - SAP® Adaptive Server® Enterprise manages multiple, large relational databases for high-volume online transaction processing (OLTP) systems and client applications.
 - SAP® IQ manages multiple, large relational databases with special indexing algorithms to support high-speed, high-volume business intelligence, decision support, and reporting client applications.
 - SAP® SQL Anywhere® manages relational databases with a small DBMS footprint, which is ideal for embedded applications and mobile device applications.

See also *DBMS* and *RDBMS*.

- **atomic materialization** – a materialization method that copies subscription data from a primary to a replicate database through the network in a single atomic operation, using a select operation with a holdlock. No changes to primary data are allowed until data transfer is complete. See also *bulk materialization* and *nonatomic materialization*.
- **BCP utility** – a feature that improves SAP Replication Server performance when replicating large batches of insert statements on the same table in SAP Adaptive Server Enterprise supported versions. SAP Replication Server implements bulk copy-in in Data Server Interface (DSI), the SAP Replication Server module responsible for sending transactions to replicate databases, using the SAP® Open Client™/SAP® Open Server™ Bulk-Library. See also *bulk copy*.
- **bulk copy** – an SAP® Open Client™ interface for the high-speed transfer of data between a database table and program variables. Bulk copying provides an alternative to using SQL **insert** and **select** commands to transfer data.
- **bulk materialization** – a materialization method whereby subscription data in a replicate database is initialized outside of the replication system. You can use bulk materialization for subscriptions to table replication definitions or function replication definitions. For example, data may be transferred from a primary database using media such as magnetic tape, diskette, CDROM, or optical storage disk. Bulk materialization involves a series of commands, starting with define subscription. See also *atomic materialization* and *nonatomic materialization*.
- **client** – in client/server systems, the part of the system that sends requests to servers and processes the results of those requests. See also *client application*.
- **client application** – software that is responsible for the user interface, including menus, data entry screens, and report formats. See also *client*.
- **commit** – an instruction to the DBMS to make permanent the changes requested in a transaction. See also *transaction*. Contrast with *rollback*.

- **database** – a collection of data with a specific structure (or schema) for accepting, storing, and providing data for users. See also *data server*, *DBMS*, and *RDBMS*.
- **database connection** – a connection that allows SAP Replication Server or Replication Agent to manage the database and send or receive transactions from or to a database. Each database in a replication system can have only one database connection in SAP Replication Server and one primary database connection in Replication Agent. See also *SAP Replication Server* and *route*.
- **data client** – a client application that provides access to data by connecting to a data server. See also *client*, *client application*, and *data server*.
- **data distribution** – a method of locating (or placing) discrete parts of a single set of data in multiple systems or at multiple sites. Data distribution is distinct from data replication, although a data replication system can be used to implement or support data distribution. Contrast with *data replication*.
- **data replication** – the process of copying primary data to remote locations and synchronizing the copied data with the primary data. Data replication is different from data distribution. Replicated data is a stored copy of data at one or more remote sites throughout a system, and it is not necessarily distributed data. Contrast with *data distribution*. See also *transaction replication*.
- **data server** – a server that provides the functionality necessary to maintain the physical representation of a table in a database. Data servers are usually database servers, but they can also be any data repository with the interface and functionality a data client requires. See also *client*, *client application*, and *data client*.
- **datatype** – a keyword that identifies the characteristics of stored information on a computer. Some common datatypes are: *char*, *int*, *smallint*, *date*, *time*, *numeric*, and *float*. Different data servers support different datatypes.
- **DBMS** – database management system

a computer-based system for defining, creating, manipulating, controlling, managing, and using databases. The DBMS can include the user interface for using the database, or it can be a standalone data server system. Compare with *RDBMS*.
- **Direct Load Materialization** – an automatic materialization method to consume data between different primary and replicate databases. Data is loaded directly from a primary table into a replicate table, no materialization queue is used. See also *materialization*, *atomic materialization*, *bulk materialization*, and *nonatomic materialization*.
- **ERSSD** – Embedded Replication Server System Database

The SAP SQL Anywhere database that stores Replication Server system tables. You can choose whether to store SAP Replication Server system tables on the ERSSD or the SAP Adaptive Server Enterprise RSSD. See also *SAP Replication Server*.
- **failback** – a procedure that restores the normal user and client access to a primary database, after a failover procedure switches access from the primary database to a replicate database. See also *failover*.
- **failover** – a procedure that switches user and client access from a primary database to a replicate database, particularly in the event of a failure that interrupts operations at the

primary database, or access to the primary database. Failover is an important fault-tolerance feature for systems that require high availability. See also *fallback*.

- **function** – a data server object that represents an operation or set of operations. SAP Replication Server distributes operations to replicate databases as functions. See also *stored procedure*.
- **function string** – a string that SAP Replication Server uses to map a function and its parameters to a data server API. Function strings allow SAP Replication Server to support heterogeneous replication, in which the primary and replicate databases are different types, with different SQL extensions and different command features. See also *function*.
- **gateway** – connectivity software that allows two or more computer systems with different network architectures to communicate.
- **inbound queue** – a stable queue managed by SAP Replication Server to spool messages received from a Replication Agent. See also *outbound queue* and *stable queue*.
- **interfaces file** – a file containing entries that define network access information for server programs in an SAP client/server architecture. Server programs may include SAP Adaptive Servers, gateways, Replication Servers, and Replication Agents. The interfaces file entries enable clients and servers to connect to each other in a network. See also *SAP Open Client* and *SAP Open Server*.
- **isql** – an Interactive SQL client application that can connect and communicate with any SAP® Open Server™ application, including SAP Adaptive Server Enterprise, Replication Agent, and SAP Replication Server. See also *SAP Open Client* and *SAP Open Server*.
- **Java** – an object-oriented programming language developed by Oracle (formerly Sun Microsystems). A platform-independent, “write once, run anywhere” programming language.
- **JVM** – Java Virtual Machine

the JVM is the part of the Java Runtime Environment (JRE) that is responsible for interpreting Java byte codes. See also *Java* and *JRE*.

- **JDBC** – Java Database Connectivity

JDBC is the standard communication protocol for connectivity between Java clients and data servers. See also *data server* and *Java*.

- **JRE** – Java Runtime Environment

the JRE consists of the Java Virtual Machine (JVM), the Java Core Classes, and supporting files. The JRE must be installed on a machine to run Java applications, such as Replication Agent. See also *Java VM*.

- **LAN** – local area network

a computer network located on the user premises and covering a limited geographical area (usually a single site). Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary can be subject to some form of regulation. Contrast with *WAN*.

- **latency** – in transaction replication, the time it takes to replicate a transaction from a primary database to a replicate database. Specifically, latency is the time elapsed between

committing an original transaction in the primary database and committing the replicated transaction in the replicate database.

In disk replication, latency is the time elapsed between a disk write operation that changes a block or page on a primary device and the disk write operation that changes the replicated block or page on a replicate device.

See also *transaction replication*.

- **LOB** – large object

a large collection of data stored as a single entity in a database.

- **Log Reader** – an internal component of Replication Agent that interacts with the primary database to capture transactions for replication. See also *Log Transfer Interface* and *LTM*.

- **LTI** – Log Transfer Interface

an internal component of Replication Agent that interacts with SAP Replication Server to forward transactions for distribution to Replication Server. See also *Log Reader* and *LTM*.

- **LTL** – Log Transfer Language

a subset of the Replication Command Language (RCL). A Replication Agent such as RepAgent uses LTL commands to submit to SAP Replication Server the information it retrieves from primary database transaction logs. See also *Log Reader* and *LTI*.

- **LTM** – Log Transfer Manager

an internal component of Replication Agent that interacts with the other Replication Agent internal components to control and coordinate Replication Agent operations. See also *Log Reader* and *LTI*.

- **maintenance user** – a special user login name in the replicate database that SAP Replication Server uses to apply replicated transactions to the database. See also *replicate database* and *SAP Replication Server*.

- **materialization** – the process of copying the data from a primary database to a replicate database, initializing the replicate database so that the replication system can begin replicating transactions. See also *atomic materialization*, *bulk materialization*, and *nonatomic materialization*.

- **MPR** – Multi-Path Replication™

SAP Replication Server feature that improves performance by enabling parallel paths of data from the source database to the target database. You can configure multi-path replication in warm standby and multisite availability (MSA) environments. These multiple paths process data independently of each other and are applicable when sets of data can be processed in parallel without transactions consistency requirements between them while still maintaining data consistency within a path, but not adhering to the commit order across different paths.

- **nonatomic materialization** – a materialization method that copies subscription data from a primary to a replicate database through the network in a single operation, without a

holdlock. Changes to the primary table are allowed during data transfer, which may cause temporary inconsistencies between replicate and primary databases. Data is applied in increments of ten rows per transaction, which ensures that the replicate database transaction log does not fill. Nonatomic materialization is an optional method for the **create subscription** command. Contrast with *atomic materialization*. See also *bulk materialization*.

- **ODBC** – Open Database Connectivity
an industry-standard communication protocol for clients connecting to data servers. See also *client*, *data server*, and *JDBC*.
- **SAP Open Client** – an SAP product that provides customer applications, third-party products, and other SAP products with the interfaces needed to communicate with SAP Open Server applications. See also *SAP Open Server*.
- **Open Client application** – An application that uses SAP Open Client libraries to implement Open Client communication protocols. See also *SAP Open Client* and *SAP Open Server*.
- **SAP Open Server** – an SAP product that provides the tools and interfaces required to create a custom server. See also *SAP Open Client*.
- **Open Server application** – a server application that uses SAP Open Server libraries to implement Open Server communication protocols. See also *SAP Open Client* and *SAP Open Server*.
- **outbound queue** – a stable queue managed by SAP Replication Server to spool messages to a replicate database. See also *inbound queue*, *replicate database*, and *stable queue*.
- **primary data** – the data source used for replication. Primary data is stored and managed by the primary database. See also *primary database*.
- **primary database** – the database that contains the data to be replicated to another database (the replicate database) through a replication system. The primary database is the source of replicated data in a replication system. Sometimes called the active database. Contrast with *replicate database*. See also *primary data*.
- **primary key** – a column or set of columns that uniquely identifies each row in a table.
- **primary site** – the location or facility at which primary data servers and primary databases are deployed to support normal business operations. Sometimes called the active site or main site. See also *primary database* and *replicate site*.
- **primary table** – a table used as a source for replication. Primary tables are defined in the primary database schema. See also *primary data* and *primary database*.
- **primary transaction** – a transaction that is committed in the primary database and recorded in the primary database transaction log. See also *primary database*, *replicated transaction*, and *transaction log*.
- **quiesce** – to cause a system to go into a state in which further data changes are not allowed. See also *quiescent*.
- **quiescent** – a state in which log scanning has stopped and all scanned records have been propagated to their destinations in a replication system. Some Replication Agent and SAP Replication Server commands require that you first quiesce the replication system.

In a database, a state in which all data updates are suspended so that transactions cannot change any data, and the data and log devices are stable.

This term is interchangeable with quiesced and in quiesce. See also *quiesce*.

- **RASD** – Replication Agent System Database

Information in the RASD is used by the primary database to recognize database structure or schema objects in the transaction log.

- **RCL** – Replication Command Language

the command language used to manage SAP Replication Server. See also *SAP Replication Server*.

- **RDBMS** – relational database management system

an application that manages and controls relational databases. Compare with *DBMS*. See also *relational database*.

- **relational database** – a collection of data in which data is viewed as being stored in tables, which consist of columns (data items) and rows (units of information). Relational databases can be accessed by SQL requests. Compare with *database*. See also *SQL*.

- **replicate data** – A set of data that is replicated from a primary database to a replicate database by a replication system. See also *primary database*, *replication system*, and *replicate database*.

- **replicate database** – a database that contains data replicated from another database (the primary database) through a replication system. The replicate database is the database that receives replicated data in a replication system. Contrast with *primary database*. See also *replicate data*, *replicated transaction*, and *replication system*.

- **replicated transaction** – a primary transaction that is replicated from a primary database to a replicate database by a transaction replication system. See also *primary database*, *primary transaction*, *replicate database*, and *transaction replication*.

- **replicate site** – the location or facility at which replicate data servers and replicate databases are deployed to support normal business operations during scheduled downtime at the primary site. Contrast with *primary site*. See also *replicate database*.

- **Replication Agent** – an application that reads a primary database transaction log to acquire information about data-changing transactions in the primary database, processes the log information, and then sends it to an SAP Replication Server for distribution to a replicate database. See also *primary database* and *SAP Replication Server*.

- **replication definition** – a description of a table or stored procedure in a primary database, for which subscriptions can be created. The replication definition, maintained by SAP Replication Server, includes information about the columns to be replicated and the location of the primary table or stored procedure. See also *SAP Replication Server* and *subscription*.

- **replication system** – a data processing system that replicates data from one location to another. Data can be replicated between separate systems at a single site, or from one or more local systems to one or more remote systems. See also *transaction replication*.

- **rollback** – an instruction to a database to back out of the changes requested in a unit of work (called a transaction). Contrast with *commit*. See also *transaction*.
- **route** – A one-way message stream from a primary SAP Replication Server to a replicate SAP Replication Server. Routes carry data-changing commands (including those for RSSDs) and replicated functions (database procedures) between separate SAP Replication Servers. See also *SAP Replication Server*.
- **RSSD** – Replication Server System Database
the SAP Adaptive Server Enterprise (SAP ASE) database containing an SAP Replication Server system tables. The user can choose whether to store Replication Server system tables on SAP ASE or embedded in an SAP SQL Anywhere database hosted by SAP Replication Server. See also *SAP Replication Server*.
- **SAP Replication Server** – an SAP software product that provides the infrastructure for a transaction replication system. See also *Replication Agent*.
- **SQL** – Structured Query Language
a nonprocedural programming language used to process data in a relational database. ANSI SQL is an industry standard. See also *transaction*.
- **stable queue** – a disk device-based, store-and-forward queue managed by SAP Replication Server. Messages written into the stable queue remain there until they can be delivered to the appropriate process or replicate database. SAP Replication Server provides a stable queue for both incoming messages (the inbound queue) and outgoing messages (the outbound queue). See also *database connection*, *SAP Replication Server*, and *route*.
- **stored procedure** – a data server object that represents an operation or set of operations. This term is often used interchangeably with *function*.
- **subscription** – a request for SAP Replication Server to maintain a replicated copy of a table, or a set of rows from a table, in a replicate database at a specified location. See also *replicate database*, *replication definition*, and *SAP Replication Server*.
- **table** – in a relational DBMS, a two-dimensional array of data or a named data object that contains a specific number of unordered rows composed of a group of columns that are specific for the table. See also *database*.
- **transaction** – a unit of work in a database that can include zero, one, or many operations (including **insert**, **update**, and **delete** operations), and that is either applied or rejected as a whole. Each SQL statement that modifies data can be treated as a separate transaction, if the database is so configured. See also *SQL*.
- **transactional consistency** – A condition in which all transactions in the primary database are applied in the replicate database, and in the same order that they were applied in the primary database.
- **transaction log** – generally, the log of transactions that affect the data managed by a data server. Replication Agent reads the transaction log to identify and acquire the transactions to be replicated from the primary database. See also *Replication Agent*, *primary database*, and *SAP Replication Server*.

Glossary

- **transaction replication** – a data replication method that copies data-changing operations from a primary database to a replicate database. See also *data replication*.
- **UDB** – IBM DB2 Universal Database (formerly IBM DB2 for Linux, UNIX, and Windows).
- **WAN** – wide-area network

a system of local-area networks (LANs) connected together with data communication lines. Contrast with *LAN*.

Index

A

- accounts, creating sybase user 10
- ASE
 - configuring Replication Server for replication from the primary database 20
 - installing replication components 14
 - installing Replication Server 14
 - system requirements 11, 12
 - testing replication 23
- ASE-to-SAP HANA database replication
 - configuring Replication Server 18

B

- before you begin 5

C

- cataloging a node and a database 104
- configuring
 - Primary Oracle database for replication 43
 - Replication Agent for DB2 UDB 108
 - Replication Agent for Microsoft SQL Server 76
 - Replication Agent for Oracle 43
 - replication components 18, 38
 - replication components for Microsoft SQL Server 72
 - Replication Server for ASE-to-SAP HANA database replication 18
 - Replication Server for DB2 UDB replication 105
 - Replication Server for Microsoft SQL Server 73
 - Replication Server for Oracle-to-SAP HANA database replication 40
- configuring ExpressConnect for HANA DB 18, 40, 73, 105, 126
- configuring Replication Agent
 - configuring and validating the resource file 81
 - configuring the primary DB2 UDB database for replication 108
 - creating a Microsoft SQL Server user and granting permissions 77

- disabling Microsoft SQL Server replication 77
 - enabling remote DAC 77
 - first-time initialization 84
 - obtaining Microsoft SQL Server DAC port number 78
 - setting the Autocommit parameter 109
 - subsequent initialization 86
 - verifying compatibility level of Microsoft SQL Server 76
 - verifying current archive setting of transaction log 108
 - verifying the Replication Agent replicating state 87
- configuring Replication Server
 - for replication from the primary database 19, 41, 74, 105
 - for replication to SAP HANA database 19, 41, 74, 105
- conventions
 - style 1
 - syntax 1
- creating
 - Oracle user 47
 - Replication Agent instance 49
 - Replication Agent instance in DB2 UDB 110
 - Replication Server connection to the replicate database 19, 41, 74, 105
 - Replication Server from primary database 40, 73, 105
- creating sybase user accounts 10

D

- DB2 UDB
 - adding a temporary tablespace to the primary database 110
 - configuring Replication Server 105
 - configuring Replication Server for replication from the primary database 107
 - creating a user and granting permissions 109
 - creating the Replication Agent instance 110
 - initializing the Replication Agent instance 116
 - installing replication components 96

Index

- resetting the primary database for replication 120
- setting environment variables 103
- system requirements 93
- testing DB2 UDB replication 118
- testing replication 118
- verifying the Replication Agent instance 115

default directory 15, 34, 97

directories

- default 15, 34, 97

dsedit utility

- adding server entries to interfaces file 9

E

errors

- while starting installer 14, 33, 67, 97

existing installation directories 15, 34, 97

ExpressConnect for HANA DB

- configuring 18, 40, 73, 105, 126

I

initializing

- Replication Agent instance 57
- Replication Agent instance in DB2 UDB 116

installation directories

- existing 15, 34, 97
- new 15, 34, 97

installation guidelines

- for Replication Agent 36, 70, 99
- for Replication Server 14, 33, 96
- Replication Server for Microsoft SQL Server 67

installation program

- Replication Agent 36, 70, 99

installer

- errors in starting 14, 33, 67, 97

installing 80

interfaces file 9

- editing 17

L

LogMiner

- configuration 44

M

Microsoft SQL Server

- configuring and verifying the primary database 76

- configuring Replication Server 73
- configuring Replication Server for replication from the primary database 76
- installing replication components 66
- installing Replication Server 67
- resetting the primary database for replication 91
- system requirements 63
- testing replication 89
- verifying the Replication Agent instance 85

Microsoft SQL Server JDBC driver 80

N

new installation directories 15, 34, 97

O

Oracle

- configuring Replication Server for replication from the primary database 43
- installing replication components 32
- installing Replication Server 33, 96
- system requirements 30, 31
- testing replication 59
- verifying the Replication Agent instance 56

Oracle-to-SAP HANA database replication

- configuring Replication Server 40

P

planning your installation

- replication components 12, 30, 64, 94

prerequisites 65, 95

- ASE 12
- Oracle 31

prerequisites for installing replication components 65, 94

R

replication

- components 11, 29, 63, 93
- overview 11, 29, 63, 93

Replication Agent

- installing 36, 70, 99

Replication Agent for DB2 UDB

- configuring 108

Replication Agent for Microsoft SQL Server
 configuring 76
Replication Agent for Oracle
 configuring 43
Replication Agent installation directory 31, 65, 95
replication components
 configuration 72, 102
Replication Server, adding
 primary database 20
resetting primary ASE database for replication 26
resetting primary Oracle database for replication
 61

S

setting up replication for DB2 UDB 93, 133
setting up replication from ASE to SAP HANA
 database 11
setting up replication from Oracle to SAP HANA
 database 29, 63
setup prerequisites 11, 30, 63, 93
sybase user account
 creating 10
system requirements
 ASE 11
 DB2 UDB 93
 Microsoft SQL Server 63

Oracle 30

T

testing
 ASE replication 23
 Microsoft SQL Server replication 89
 Oracle replication 59

U

user account
 sybase 10

V

verifying
 current archive setting of the redo logs 44
 installation 38, 71, 101
 primary Oracle configuration for replication
 49
 Replication Agent instance in DB2 UDB 115
 Replication Agent instance in Microsoft SQL
 Server 85
 Replication Agent instance in Oracle 56
 supplemental logging of primary key data 44

