



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

**Replication Server® 15.7.1
SP100**

Linux, UNIX, and Microsoft Windows

DOCUMENT ID: DC01974-01-1571100-01

LAST REVISED: May 2013

Copyright © 2013 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

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 in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Oracle and/or its affiliates in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Contents

CHAPTER 1: Conventions	1
CHAPTER 2: How to Use This Document	5
CHAPTER 3: ASE-to-SAP HANA Database	
Replication Setup	7
System Architecture	7
Prerequisites	7
Planning Your Installation	8
System Requirements	8
Identify the Sybase Installation Directory	8
Accessing and Downloading Products from SMP	8
Installing Replication Components	9
Installing Replication Server	10
Editing the Interfaces File	13
Configuring Replication Components	13
Configuring Replication Server	14
Obtaining the SAP HANA Database ODBC	
Drivers	14
Configuring Replication Server for Replication to	
SAP HANA Database	14
Configuring Replication Server for Replication	
from the Primary Database	16
Testing Replication	18
Resetting the Primary ASE Database For Replication	
.....	20

- CHAPTER 4: Oracle-to-SAP HANA Database Replication Setup23**
- System Architecture23**
- Prerequisites24**
- Planning Your Installation24**
- System Requirements24
- Identify the Sybase Installation Directory25
- Replication Agent Installation Directory25
- Accessing and Downloading Products from SMP25
- Installing Replication Components26**
- Installing Replication Server26
- Installing Replication Agent29
- Verifying the Installation31
- Configuring Replication Components32**
- Configuring Replication Server32
- Obtaining the SAP HANA Database ODBC Drivers32
- Configuring Replication Server for Replication to SAP HANA Database32
- Configuring Replication Server for Replication from the Primary Database34
- Configuring Replication Agent35
- Primary Oracle Database Configuration for Replication35
- Creating the Replication Agent Instance38
- Verifying the Replication Agent Instance Installation44
- Initializing the Replication Agent Instance46
- Creating a Replication Server Connection to the Primary Database48
- Testing Replication48**
- Resetting the Primary Oracle Database for Replication50

CHAPTER 5: Microsoft SQL Server-to-SAP HANA Database Replication Setup	51
System Architecture	51
Prerequisites	52
Planning Your Installation	52
System Requirements	52
Identify the Sybase Installation Directory	53
Replication Agent Installation Directory	53
Accessing and Downloading Products from SMP	53
Installing Replication Components	54
Installing Replication Server	54
Installing Replication Agent	57
Verifying the Installation	58
Configuring Replication Components	59
Configuring Replication Server	59
Obtaining the SAP HANA Database ODBC Drivers	59
Configuring Replication Server for Replication to SAP HANA Database	59
Configuring Replication Server for Replication from the Primary Database	61
Configuring Replication Agent	62
Primary Microsoft SQL Server Configuration and Verification	62
Configuring and Validating the Replication Agent Resource File	67
First-Time Initialization	70
Subsequent initialization	72
Creating a Replication Server Connection to the Primary Database	73
Verifying the Replication Agent Replicating State	74
Testing Replication	75

Resetting the Primary Microsoft SQL Server
Database for Replication77

**CHAPTER 6: DB2 UDB-to-SAP HANA Database
Replication Setup79**

- System Architecture79**
- Prerequisites79**
- Planning Your Installation80**
 - System Requirements80
 - Identify the Sybase Installation Directory80
 - Replication Agent Installation Directory81
 - Accessing and Downloading Products from SMP81
- Installing Replication Components82**
 - Installing Replication Server82
 - Installing Replication Agent85
 - Verifying the Installation87
- Configuring Replication Components88**
 - Setting DB2 UDB Environment Variables88
 - Node and Database Cataloging in an DB2 UDB Client
.....89
 - Configuring Replication Server89
 - Obtaining the SAP HANA Database ODBC
Drivers89
 - Configuring Replication Server for Replication to
SAP HANA Database89
 - Configuring Replication Server for Replication
from the Primary Database91
 - Configuring Replication Agent92
 - Primary DB2 UDB Configuration for Replication
.....92
 - Creating the Replication Agent Instance94
 - Verifying the Replication Agent Instance
Installation99
 - Initializing the Replication Agent Instance100

Creating a Replication Server Connection to the Primary Database	102
Testing Replication	102
Resetting the Primary DB2 UDB Database for Replication	104
 CHAPTER 7: Glossary	 105
 Index	 113

Contents

CHAPTER 1 Conventions

These style and syntax conventions are used in Sybase® 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 Adaptive Server® Enterprise, 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

CHAPTER 2 **How to Use This Document**

Use this document to set up and begin using a basic replication system with SAP® HANA database as the replicate.

This document contains these sections:

- "ASE-to-SAP HANA Database Replication Setup" provides instructions for setting up an ASE-to-SAP HANA database replication system.
- "Oracle-to-SAP HANA Database Replication Setup" provides instructions for setting up an Oracle-to-SAP HANA database replication system.
- "Microsoft SQL Server-to-SAP HANA Database Replication Setup" provides instructions for setting up a Microsoft SQL Server-to-SAP HANA database replication system.
- "DB2 UDB-to-SAP HANA Database Replication Setup" provides instructions for setting up a DB2 UDB-to-SAP HANA database replication system.

You may need to refer to these other Replication Server® documents, which are available on Sybase Infocenter at <http://infocenter.sybase.com>:

- *Replication Server Data Assurance Option Users Guide* provides information on administrative tasks, commands, performance and tuning, and troubleshooting for the Replication Server Data Assurance Option.
- *Replication Server Troubleshooting Guide* provides information on errors and troubleshooting procedures for Replication Server.
- *Replication Server Release Bulletin* provides the most current information about product compatibility, special installation instructions, known issues, and document updates for Replication Server.

Additional information for the SAP HANA database, including information on HANA Studio, is available at http://help.sap.com/hana_platform.

CHAPTER 2: How to Use This Document

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.

These tasks use the sample Replication Server instance SAMPLE_RS. If you have already installed Replication Server, skip to *Configuring Replication Components* on page 13.


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

System Architecture

There are multiple components required to implement an ASE-to-SAP HANA database replication system.

- A primary Adaptive Server data server
- A replicate SAP HANA database
- A Replication Server (with ExpressConnect for HANA DB)

Figure 1: Sybase Replication System for an Adaptive Server Primary and SAP HANA Database Replicate



Prerequisites

Before you begin setting up replication for Adaptive Server, review the prerequisites.

Make sure:

- You are familiar with Adaptive Server and SAP HANA database.

CHAPTER 3: ASE-to-SAP HANA Database Replication Setup

- You have a valid license for Replication Server with ExpressConnect for HANA DB.
- Adaptive Server is installed and configured to serve as the source (primary) database from which Replication Agent replicates transactions, and SAP HANA database is installed and configured as the target (replicate) database to which Replication Server with ExpressConnect for HANA DB replicates transactions.
- TCP/IP connectivity is available between all hosts.

During this task:

- Do not use Replication Server reserved words for object names and connection names. A complete list of reserved words is documented in the *Replication Server Reference Manual*.
- Because some directories, files, executable commands, and examples are shown only for Solaris, adjust accordingly for Microsoft Windows and for other UNIX platforms.

Planning Your Installation

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

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 Sybase Installation Directory

Determine the directory in which to install the replication components.

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 Sybase installation and configuration activities.

Accessing and Downloading Products from SMP

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

Prerequisites

You must have a login ID and password to download software from the SMP.

Task

Note: If you are installing Replication Server Options (RSO), this task requires you to perform several steps twice because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/>, and enter your login ID and password.
2. From the product listing, select:
 - **SAP Sybase Replication Server, Option for HANA** or
 - **SAP Sybase Replication Server, HANA Edition**
3. Select the software you need by version and platform.
4. 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.

5. 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.
6. 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. Check each partition to make sure it 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 the *Replication Server ASE-to-ASE Replication Quick Start Guide*.

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 the following:

```
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 Sybase Software Asset Management 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 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 Sybase 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 in place of SAMPLE_RS.

- Shut down Replication Server:

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

- If your primary database and Replication Server are installed on different hosts, change "localhost" in the Replication Server interfaces file to the host name for your primary database.
- 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 used to identify 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.

Configuring Replication Components

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

Configuring Replication Server

Review the procedures for configuring 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, see the *Replication Server Configuration Guide* for instructions.

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, *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_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 placeholder: You must provide a value, but it is not used..
- *muser* is the maintenance user for the replicate SAP HANA database instance that you created in a previous step.
- *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 1. Sample Values for the myprimary.rs File

Parameter	Description	Value
<code>sybinit.release_directory</code>	Valid path of the Replication Server software (\$SYBASE).	<code>/opt/software</code>
<code>rs.rs_name</code>	Name of the Replication Server.	<code>SAMPLE_RS</code>
<code>rs.rs_rs_sa_user</code>	User ID that has “sa” privileges on Replication Server.	<code>sa</code>
<code>rs.rs_rs_sa_pass</code>	Password of the “sa” user.	<code>sa_pass</code>
<code>rs.rs_ds_name</code>	Name of the data server that hosts the primary database.	<code>sunak1502i</code>
<code>rs.rs_ds_sa_user</code>	User ID that has “sa” privileges on data server.	<code>sa</code>
<code>rs.rs_ds_sa_password</code>	Password of the “sa” user for the data server.	<code>password</code>
<code>rs.rs_db_name</code>	Name of the primary database.	<code>pubs2</code>
<code>rs.rs_needs_repagent</code>	Specifies whether you plan to replicate from specified primary database.	<code>yes</code>

Parameter	Description	Value
rs.rs_db_maint_user	User ID whose work is not replicated when logged on the primary database. The user is called maintenance user.	<database>_maint Note: 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
rs.rs_ltm_rs_user	User that the Replication Agent will use to log into the Replication Server. The name must exist. This name typically comes from values that were set up during Replication Server creation time: rs.rs_ltm_rs_user.	<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	Specifies 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** command

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
48	DSI EXEC	Awaiting Command	102 (1) sunak1502i.pubs2
33	DSI	Awaiting Message	102 sunak1502i.pubs2
35	DIST	Awaiting Wakeup	102 sunak1502i.pubs2
36	SQT	Awaiting Wakeup	102:1 DIST sunak1502i.pubs2
34	SQM	Awaiting Message	102:1 sunak1502i.pubs2
32	SQM	Awaiting Message	102:0 sunak1502i.pubs2
37	REP AGENT	Awaiting Command	sunak1502i.pubs2
39	NRM	Awaiting Message	sunak1502i.pubs2
27	DSI EXEC	Awaiting Command	
101 (1)	SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD		
20	DSI	Awaiting Message	101
26	SQM	Awaiting Message	101:0
21	dSUB	Sleeping	
15	dCM	Awaiting Message	
18	dAIO	Awaiting Message	
23	dREC	Sleeping	dREC
9	dDELSEG	Awaiting Message	
49	USER	Active	sa
14	dALARM	Awaiting Wakeup	
24	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:

```
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;
```
4. Mark the `ptab1` table for replication:

```
sp_setreptable ptab1, 'true'
go
```

- 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 tablename, the owner of the table—*<tableowner>*—is the user that is signed on at the time 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 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.

- Create a replication definition against the primary Adaptive Server database. To create the `ptab1_repdef` replication definition for the `ptab1` table:

```
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 *pds.pdb* is the name of the primary connection created earlier.

- Create a subscription against the replicate connection. To create the `ptab1_sub` subscription for the `ptab1_repdef` replication definition:

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

where *rds.rdb* is the name of the replicate SAP HANA database connection created earlier.

For information on using the create subscription command in direct load materialization, see *Replication Server Reference Manual*.

- On the Adaptive Server database, insert data into the primary `ptab1` table:

```
insert into ptab1
values (3, "Tom Servo")
go
```

11. On the SAP HANA database, verify that the data replicated to the replicate PTAB1 table:

```
SELECT * FROM PTAB1;
```

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

2. Connect to the primary database:

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

where:

- *pds.pdb* is the name of the primary connection.
- *pds* is the name of the Adaptive Server data server of the primary connection.
- *pdb* is the name of 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 rssd
2>go
```

7. Reset the locator stored in Replication Server:

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

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

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*

These 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 on Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

There are multiple components 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)
- The Replication Server Options components:
 - Replication Agent for Oracle

Figure 2: Sybase Replication System for an Oracle Primary and SAP HANA Database Replicate



Prerequisites

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

Make sure:

- You are familiar with Oracle and SAP HANA database.
- You have valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- Oracle Database 10g or 11g Enterprise Edition is installed and configured to serve as the source (primary) database from which Replication Agent replicates transactions, and SAP HANA database is installed and configured as the target (replicate) database to which Replication Server with ExpressConnect for HANA DB replicates transactions.
- 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 is available between all hosts.
- Any operating system patches required for Java have been installed.
- The Replication Agent has direct access to the Oracle online and archived redo logs.

Do not use Replication Server reserved words for object names or connection names. A complete list of reserved words is in the *Replication Server Reference Manual*.

Planning Your Installation

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

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 Sybase Installation Directory

Determine the directory in which to install the replication components.

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 Sybase 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.

Accessing and Downloading Products from SMP

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

Prerequisites

You must have a login ID and password to download software from the SMP.

Task

Note: If you are installing Replication Server Options (RSO), this task requires you to perform several steps twice because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/>, and enter your login ID and password.
2. From the product listing, select:
 - **SAP Sybase Replication Server, Option for HANA** or
 - **SAP Sybase Replication Server, HANA Edition**
3. Select the software you need by version and platform.
4. 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.

5. 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.
6. 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 13

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. Check each partition to make sure it 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.

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 the following:

```
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 Sybase Software Asset Management 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 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 Sybase 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 ITEMPDIR (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. Insert the Replication Agent distribution media in the appropriate drive.
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, and 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**.

CHAPTER 4: Oracle-to-SAP HANA Database Replication Setup

If the directory you specified exists, and already contains a Replication Agent installation, you are warned that you will overwrite the older version.

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

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

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 primary 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 Sybase Product Download Center at <http://www.sybase.com/detail?id=1025266>, 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 *Sybase Software Asset Management Users Guide*. *SySAM Users Guide*.

Click **Next** until you see the Sybase Software Asset Management Notification window.

10. On the Sybase Software Asset Management 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).
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 Sybase copyright and the Replication Agent version string.

Configuring Replication Components

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

Configuring Replication Server

Review the procedures for configuring Replication Server for Oracle-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, see the *Replication Server Configuration Guide* for instructions.

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 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_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
```

CHAPTER 4: Oracle-to-SAP HANA Database Replication Setup

```
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 that you created in the previous step.
- *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.

Granting create object Permission to the rs_username User

Grant create object permission to the user ID specified by the Replication Agent **rs_username** parameter.

The **rs_username** user, which is used for communication between Replication Agent and Replication Server, must have **create object** permission before it can be used 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 -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 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.

Verifying the Current Archive Setting of the Redo Logs

Verify that redo logging, which Oracle uses to maintain a log infrastructure, is turned on. Turn redo logging on only 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

Verify that Oracle logs the values of the primary keys and unique indexes.

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. After LogMiner is installed, 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;
```

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 DATABASE 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;
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.LOBFrag$ 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;

```

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.

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.

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.

- For Oracle 10g:
`$ORACLE_HOME/jdbc/lib/ojdbc14.jar`
- For Oracle 11g Release 1:
`$ORACLE_HOME/jdbc/lib/ojdbc5.jar`
- For Oracle 11g Release 2:
`$ORACLE_HOME/jdbc/lib/ojdbc6.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).

- 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 this:

- 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
- (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 *Verification of a JDBC Client Installation* in the *Oracle Database JDBC Developer's Guide and Reference Guide*.

- 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 `timezlr*.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.

Determine which `timezone` file is used by the primary Oracle instance and make a copy of that file available on the machine hosting the Replication Agent.

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.

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 means 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 (`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:

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 assigned to properties in the resource file so they match the values required for your environment. The majority 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 must not contain single quotes, double quotes, or a semicolon.
 - The **pds_username** and **pds_password** configuration parameter values must not 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 2. 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
<code>ltm_admin_pw</code>	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
<code>pds_tns_connection</code>	Connection name found in the <code>tnsnames.ora</code> file which identifies the connection information for the primary database.	ORA102.JDOE_HOST.COM
<code>pds_tns_filename</code>	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>
<code>pds_username</code>	User ID that Replication Agent uses to access primary data server.	RA_USER
<code>pds_password</code>	Password for <code>pds_username</code> .	sybase
<code>rs_host_name</code>	Machine where Replication Server is installed.	jdoe_host1
<code>rs_port_number</code>	Port number where Replication Server is installed.	11752

Parameter	Description	Example Values
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. Note: 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 (UNIX or Linux) or %SYBASE%\REP-15_5\install\<server>.cfg (Windows).	<ul style="list-style-type: none"> Windows: cp850 UNIX: iso_1 Note: For a sample Replication Server instance, the rs_charset parameter must be set to utf8.
rs_source_ds	Valid name representing data server of Oracle primary database. This value is used to create a Replication Server connection from the primary database.	pds
rs_source_db	Valid name representing 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/Ora10g/oracore/zoneinfo/timezone.dat
start_instance	Start the instance that was created.	yes
initialize_instance	Initialize the Replication Agent instance.	yes
pdb_include_archives	Enables or disables the use of Oracle archive log files.	USE_DEFAULT

Parameter	Description	Example Values
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_adminrao_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.

Note: If, in your response file, you set **start_instance** to yes, your instance is also running. If you set **initialize_instance** to yes, your instance is also initialized.

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
```

```
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.

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
```

2. Verify the connection to Replication Agent:

- Open a command window in the %SYBASE% (Windows) or \$SYBASE (UNIX or Linux) directory of your Replication Agent installation.
- Set the environment variables by executing the SYBASE.bat file (Windows) or by sourcing the SYBASE.csh file (UNIX or Linux).
- 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
```

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*.

- Enter:

```
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*.

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 46

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:

```
1> resume
2> go
Msg 32000, Level 20, State 0:
Server 'rao', 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 -Uusername -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.
- 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.
9	DSI	Awaiting Message	101 SAMPLE_RS_ERSSD.
8	SQM	Awaiting Message	101:0 SAMPLE_RS_ERSSD.

CHAPTER 4: Oracle-to-SAP HANA Database Replication Setup

```
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
```

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 > Sybase Replication Products > Replication Server > Database Connections > Maintenance User Purpose*.

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

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

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

Record the name of the replication definition for use later.

6. 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 tablename, the owner of the table—`<tableowner>`—is the user that is signed on at the time the table is created.

7. 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
```

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. 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
```

10. Create a subscription against the replicate connection.

To create the `ptab1_sub` subscription for the `<repdefname>` replication definition:

```
create subscription ptab1_sub
for <repdefname>
with replicate at rds.rdb
without materialization
go
```

where `<repdefname>` is the replication definition that you altered in the previous step, and `rds.rdb` is the name of the replicate SAP HANA database connection created earlier.

CHAPTER 4: Oracle-to-SAP HANA Database Replication Setup

For information on using the create subscription command in direct load materialization, see *Replication Server Reference Manual*.

11. On the Oracle database, insert data into the primary `ptab1` table and commit:

```
insert into ptab1 values (4, 'Crow T. Robot');
commit;
```

12. On the SAP HANA database, verify that the data 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:

```
1> suspend
2> go
```

2. Archive all current redo log data:

```
ra_locator move_truncpt
```

3. 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 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
1> resume
2> 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.

These 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 on Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

There are multiple components 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)
- The Replication Server Options components:
 - Replication Agent for Microsoft SQL Server

Figure 3: Sybase Replication System for a Microsoft SQL Server Primary and SAP HANA Database Replicate



Prerequisites

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

Make sure:

- You are familiar with Microsoft SQL Server and SAP HANA database.
- You have valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- You have valid licenses for Replication Server and Replication Server Options.
- You have obtained the licenses for the Replication Server and Replication Server Heterogeneous Edition.
- Microsoft SQL Server 2008 is installed and configured to serve as the source (primary) database from which Replication Agent replicates transactions, and SAP HANA database is installed and configured as the target (replicate) database to which Replication Server with ExpressConnect for HANA DB replicates transactions.
- Microsoft SQL Server JDBC driver 1.2 is installed.
- TCP/IP connectivity is available 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. A complete list of reserved words is in the *Replication Server Reference Manual*.

Planning Your Installation

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

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 Sybase Installation Directory

Determine the directory in which to install the replication components.

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 Sybase 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.

Accessing and Downloading Products from SMP

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

Prerequisites

You must have a login ID and password to download software from the SMP.

Task

Note: If you are installing Replication Server Options (RSO), this task requires you to perform several steps twice because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/>, and enter your login ID and password.
2. From the product listing, select:
 - **SAP Sybase Replication Server, Option for HANA** or
 - **SAP Sybase Replication Server, HANA Edition**
3. Select the software you need by version and platform.
4. 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.

5. 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.
6. 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 both the primary and replicate databases. 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. Check each partition to make sure it 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.

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 the following:

```
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 Sybase Software Asset Management 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 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 Sybase 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. Insert the Replication Agent distribution media in the appropriate drive.
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, and 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 primary 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 Sybase Product Download Center at <http://www.sybase.com/detail?id=1025266>, 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 SysSAM license for the Replication Agent component in RSHE. Specify that license file here.

See the *Sybase Software Asset Management Users Guide*.

Click **Next** until you see the Sybase Software Asset Management Notification window.

- 10.** On the Sybase Software Asset Management 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 Sybase copyright and the Replication Agent version string.

Configuring Replication Components

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

Configuring Replication Server

Review the procedures for configuring Replication Server for Microsoft SQL Server-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, see the *Replication Server Configuration Guide* for instructions.

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.

CHAPTER 5: Microsoft SQL Server-to-SAP HANA Database Replication Setup

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_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 placeholder: You must provide a value, but it is not used..
- *muser* is the maintenance user for the replicate SAP HANA database instance that you created in the previous step.
- *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.

Granting create object Permission to the rs_username User

Grant create object permission to the user ID specified by the Replication Agent **rs_username** parameter.

The **rs_username** user, which is used for communication between Replication Agent and Replication Server, must have **create object** permission before it can be used 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

Review the procedures for configuring and verifying 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.

Replication Agent supports Microsoft SQL Server 2008, Microsoft SQL Server 2008 R2, and 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 "Microsoft SQL Server 2008 (100)." If it is not, select that compatibility from the drop-down list, 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 the user. 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.

Prerequisites

Task

1. 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.

2. Change the Microsoft SQL Server **remote admin connections** configuration option to 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. Open the ERRORLOG file in a text editor. 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 the string "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.

1. To determine the version of the library, in Windows Explorer, right-click `c:\windows\system32\fltlib.dll`, select **Properties**, and click the **Version** tab in the **Properties** dialog box.
2. 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 `winx86` or `winx64`.
 - `winx64` is for 64-bit Windows Server 2008, Windows Server 2008 R2, and Windows 7.
 - `winx86` is for 32-bit Windows Server 2008, Windows Server 2008 R2, and Windows 7.
 - `winvistax64` is for Windows 7.
2. In Windows Explorer, navigate to the `sybfilter` driver installation directory. This directory is located at `%SYBASE%\RAX-15_2\system\<platform>`, where `<platform>` is `winx86`, `winx64`, or `winvistax64`.
3. Right-click the `sybfilter.inf` file to install the `sybfilter` driver.

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\<platform>`, 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\<platform>`, 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 double-click the *sybfiltermgr.exe* file 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 as follows:

```
add serverName dbName logFilePath
```

For example, to add the log file named *pdb2_log.ldf* at C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data\ to the *dbName* database on the *serverName* data server, use this:

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

Note: If you add the log file path with the user manager, the user manager refreshes all log paths in the *sybfilter* driver automatically after adding the log path into the configuration file.

- 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
```

Note: 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 for your primary database before adding the log file path to the *sybfilter* driver, 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 Microsoft SQL Server Analysis Service.

Prerequisites

If you have installed and started the Microsoft SQL Server Analysis Service, first stop the process.

Task

1. From the Windows Control Panel, select:
Administrative Tools > Services
2. In the listed services, find the service named Microsoft 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.

Prerequisites

Task

1. Go to **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 this string:
`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 can be recorded and stored in a resource file. Using a resource file provides a means to record or retain the configuration information for a Replication Agent instance, allowing an instance to be 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 assigned to properties in the resource file so they match the values required for your environment. The majority 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 must not contain single quotes, double quotes, or a semicolon.
 - The **pds_username** and **pds_password** configuration parameter values must not 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 3. 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 Administration connection 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
rs_source_ds	Valid name representing data server of primary Microsoft SQL Server database.	NY
rs_source_db	Valid name representing primary Microsoft SQL Server database.	NYmss5

Parameter	Description	Example values
rs_charset	Character set that Replication Server is using. Note: 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. Note: 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.	yes

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:

- 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.

First-Time 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, skip to Subsequent Initialization.

See also

- *Subsequent initialization* on page 72
- *Verifying the Replication Agent Replicating State* on page 74

Configuring the Primary Data Server and a Replication Agent Instance for the First Time

Configure the primary data server and a Replication Agent instance for the first time.

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 the resource file has been validated, use the **ra_admin** utility to create the Replication Agent instance using the **-r** parameter. For example:

```
%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 your resource file set **start_instance** to **no**, 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) Issue this command to create server xlog objects:

```
server_xlog init
```

- c) Issue this command to shutdown 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 74

Creating and Initializing a Replication Agent Instance

Execute the resource file to create the Replication Agent instance.

1. After the resource file has been validated, use the **ra_admin** utility to create the Replication Agent instance using the **-r** parameter. For example:

```
%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 your resource file has **initialize_instance** set to **no**, start the Replication Agent instance.
 - a) Change to the directory containing your Replication Agent instance RUN file, for example:
 - 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
```

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 your resource file has **initialize_instance** set to **no**, 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.

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 > Sybase Replication Products > Replication Server > Database Connections > Maintenance User Purpose*.

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

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:

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

- 2. The `ra_status` command returns the state of the Replication Agent. It is good practice to 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.

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

- 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:

```
admin who
go
```

You see:

```
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) NY.NYmss5
DSI           Suspended          103 NY.NYmss5
24  DIST              Awaiting Wakeup   103 NY.NYmss5
25  SQT               Awaiting Wakeup   103:1 DIST NY.NYmss5
23  SQM               Awaiting Message  103:1 NY.NYmss5
22  SQM               Awaiting Message  103:0 NY.NYmss5
62  REP AGENT         Awaiting Command  NY.NYmss5
```

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 ptab1, mark
go

```

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

```

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

```

Record the name of the replication definition for use later.

6. 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 tablename, the owner of the table—*<tableowner>*—is the user that is signed on at the time the table is created.

7. 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

```

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. 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

```

10. Create a subscription against the replicate connection.

To create the *ptab1_sub* subscription for the *<repdefname>* replication definition:

```

create subscription ptab1_sub
for <repdefname>
with replicate at rds.rdb
without materialization
go

```

where *<repdefname>* is the replication definition that you altered in the previous step, and *rds.rdb* is the name of the replicate SAP HANA database connection created earlier.

For information on using the create subscription command in direct load materialization, see *Replication Server Reference Manual*.

11. On the Microsoft SQL Server database, insert data into the primary `ptab1` table and commit:

```
insert into ptab1
values (5, "Joel Robinson")
go
commit
go
```

12. On the SAP HANA database, verify that the data replicated to the replicate `PTAB1` table:

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

See also

- *Configuring Replication Server for Replication from the Primary Database* on page 61

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 using this command in the Replication Agent to archive all current redo log data:

```
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 the **ra_admin deinit, force** command followed by a normal **ra_admin init** command (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
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.

These 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 on Replication Server and its components, see the *Replication Server Release Bulletin*.

System Architecture

There are multiple components 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)
- The Replication Server Options components:
 - Replication Agent for UDB

Figure 4: Sybase Replication System for a DB2 UDB Primary and SAP HANA Database Replicate



Prerequisites

Review the prerequisites before you begin setting up DB2 UDB-to-SAP HANA database replication.

Make sure:

CHAPTER 6: DB2 UDB-to-SAP HANA Database Replication Setup

- You are familiar with DB2 UDB and SAP HANA database.
- You have valid licenses for Replication Server with ExpressConnect for HANA DB and Replication Server Options.
- You have obtained the licenses for the Replication Server and the Replication Server Heterogeneous Edition.
- DB2 UDB 9.5 or 9.7 is installed and configured to serve as the source (primary) database from which Replication Agent replicates transactions, and SAP HANA database is installed and configured as the target (replicate) database to which Replication Server with ExpressConnect for HANA DB replicates transactions.
- TCP/IP connectivity is available between all hosts.
- Any operating system patches required for Java have been installed.

Do not use Replication Server reserved words for object names or connection names. A complete list of reserved words is in the *Replication Server Reference Manual*.

Planning Your Installation

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

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 Sybase Installation Directory

Determine the directory in which to install the replication components.

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 Sybase 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.

Accessing and Downloading Products from SMP

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

Prerequisites

You must have a login ID and password to download software from the SMP.

Task

Note: If you are installing Replication Server Options (RSO), this task requires you to perform several steps twice because you are downloading Replication Server and other products.

1. Go to <https://websmp204.sap-ag.de/>, and enter your login ID and password.
2. From the product listing, select:
 - **SAP Sybase Replication Server, Option for HANA** or
 - **SAP Sybase Replication Server, HANA Edition**
3. Select the software you need by version and platform.
4. 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.

5. 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.

6. 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 both the primary and replicate databases. 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. Check each partition to make sure it 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.

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 the following:

```
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 Sybase Software Asset Management 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 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 Sybase 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. Insert the Replication Agent distribution media in the appropriate drive.
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, and 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 the directory you specified exists, and already contains a Replication Agent installation, you are warned that you will overwrite the older version.

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

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

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 primary 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 Sybase Product Download Center at <http://www.sybase.com/detail?id=1025266>, 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 *Sybase Software Asset Management Users Guide*/*SySAM Users Guide*.

Click **Next** until you see the Sybase Software Asset Management Notification window.

10. On the Sybase Software Asset Management 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).
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 Sybase copyright and the Replication Agent version string.

Configuring Replication Components

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

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 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 section concerns the configuration of a 64-bit DB2 UDB server.

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 all other platforms, 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, copy either the `db2cshrc` or `db2profile` script:

- For C shell:

```
cp db2cshrc mydb2cshrc
```

- For Bourne or Korn shell:

```
cp db2profile mydb2profile
```

3. Edit your copy of the file. 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 example, add these lines:

- For C shell, in `mydb2cshrc`:

```
# 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`:

```
# 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 an 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 on node and database cataloging in a DB2 UDB client.

Configuring Replication Server

Review the procedures for configuring Replication Server for DB2 UDB-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, see the *Replication Server Configuration Guide* for instructions.

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 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 that you created in a previous step.
- *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.

Granting create object Permission to the rs_username User

Grant create object permission to the user ID specified by the Replication Agent **rs_username** parameter.

The **rs_username** user, which is used for communication between Replication Agent and Replication Server, must have **create object** permission before it can be used 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 -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, DB2 UDB requires you to 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 that the DB2 UDB connectivity **autocommit** parameter be turned on (autocommit=1).

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. Edit the **autocommit** parameter setting in the `db2cli.ini` file in `<HOME>\sqlllib\cfg`.

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 this procedure to create a maintenance user.

1. Create a new operating system user named *ra_user* using commands appropriate for your operating system.

For example, this command can be used to create a user named *ra_user* on the UNIX and Linux operating systems:

```
useradd -gusers -Gmgmt -s/bin/shell -psybase  
-d/home/ra_user -m ra_user
```

where *password* is the password corresponding to the *ra_user* user name.

2. Start the DB2 UDB command-line processor (CLP):

```
%>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 (CLP):

```
%>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: You can determine what the DB2 UDB page size is with 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, as described in

"Creating a DB2 UDB User and Granting Permissions." 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 assigned to properties in the resource file so they match the values required for your environment. The majority 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 must not contain single quotes, double quotes, or a semicolon.
 - The `pds_username` and `pds_password` configuration parameter values must not 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 4. Resource File Parameters

Parameter	Description	Example values
<code>instance name</code>	Any valid name.	myra

Parameter	Description	Example values
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
itm_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. Note: 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	Data source name or database alias for the primary database	DBALIAS
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.	NY
rs_source_db	Valid name representing the primary DB2 UDB database.	NYudb9

Parameter	Description	Example values
rs_charset	Character set that Replication Server is using. Note: 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	Identifies 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.	yes
initialize_instance	Initialize the Replication Agent instance.	yes

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 in step 3 in the procedure, *Creating a Replication Server Connection to the Primary Database*.

- Set the DB2 UDB environment variables by sourcing `db2cshrc`, `db2profile`, or your customized versions of the scripts as described in *Setting DB2 UDB Environment Variables*.
- 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
```

```
%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.

- 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, in your response file, you set **start_instance** to **yes**, your instance is also running. If you set **initialize_instance** to **yes**, your instance is also initialized.

- Creation results are returned as either:

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

See the *Replication Agent Administration Guide*.

- If, in your resource file, you set **start_instance** to **yes**, skip this step, and continue with the next section. If your resource file did not configure **start_instance** to **false**, 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
```

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 100

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.

Note: If, in your resource file, you set **initialize_instance** to yes, skip step 1 and continue to step 2.

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.rdbDSI
EXEC      Suspended      103(1) NY.NYudb9
DSI       Suspended      103 NY.NYudb9
24  DIST        Awaiting Wakeup  103 NY.NYudb9
25  SQT         Awaiting Wakeup  103:1 DIST NY.NYudb9
23  SQM         Awaiting Message 103:1 NY.NYudb9
```

```
22 SQM      Awaiting Message 103:0 NY.NYudb9
62 REP AGENT Awaiting Command NY.NYudb9
```

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.

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. In Replication Agent, create a replication definition against the primary DB2 UDB database:

```
1> rs_create_repdef ptabl
2> go
Table/Procedure Name  RepDef Name      Status
-----
<DB2TableOwner>.PTAB1 "<repdefname>" Created
(1 row affected)
```

Record the name of the replication definition for use later.

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

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

case sensitive

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 tablename, the owner of the table—*<tableowner>*—is the user that is signed on at the time the table is created.

7. 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
```

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

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

10. Create a subscription against the replicate connection.

To create the *ptabl_sub* subscription for the *<repdefname>* replication definition:

```
create subscription ptabl_sub
for <repdefname>
with replicate at rds.rdb
without materialization
go
```

where *<repdefname>* is the replication definition that you altered in the previous step, and *rds.rdb* is the name of the replicate SAP HANA database connection created earlier.

For information on using the create subscription command in direct load materialization, see *Replication Server Reference Manual*.

11. On the DB2 UDB database, insert data into the primary PTAB1 table and commit:

```
INSERT INTO PTAB1
(IDNO, NAME) VALUES
(6, 'Dr. Forrester');
COMMIT;
```

12. On the SAP HANA database, verify that the data replicated to the replicate PTAB1 table:

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

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:

```
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
```


CHAPTER 7 **Glossary**

This glossary describes Replication Server Options terms.

- **Adaptive Server®** – the brand name for Sybase relational database management system (RDBMS) software products.
 - Adaptive Server® Enterprise manages multiple, large relational databases for high-volume online transaction processing (OLTP) systems and client applications.
 - Sybase®IQ manages multiple, large relational databases with special indexing algorithms to support high-speed, high-volume business intelligence, decision support, and reporting client applications.
 - SQL Anywhere® (formerly Adaptive Server 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 database to a replicate database in a single, atomic operation. No changes to primary data are allowed until the subscription data is captured at the primary database. See also *bulk materialization* and *nonatomic materialization*.
- **BCP utility** – a bulk copy transfer utility that provides the ability to load multiple rows of data into a table in a target database. See also *bulk copy*.
- **bulk copy** – an 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. 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 Replication Server to manage the database and distribute transactions to the database. Each database in a replication system

can have only one database connection in Replication Server. See also *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** – an abbreviation for 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*.
- **ERSSD** – an abbreviation for Embedded Replication Server System Database, which manages replication system information for a Replication Server. See also *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 *failback*.
- **function** – a data server object that represents an operation or set of operations. Replication Server distributes operations to replicate databases as functions. See also *stored procedure*.
- **function string** – a string that Replication Server uses to map a function and its parameters to a data server API. Function strings allow 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 Replication Server to spool messages received from a Replication Agent. See also *outbound queue* and *stable queue*.
- **interfaces file** – a file containing information that Sybase Open Client and Open Server™ applications need to establish connections to other Open Client and Open Server applications. See also *Open Client* and *Open Server*.
- **isql** – an Interactive SQL client application that can connect and communicate with any Sybase Open Server application, including Adaptive Server, Replication Agent, and Replication Server. See also *Open Client* and *Open Server*.
- **Java** – an object-oriented programming language developed by Sun Microsystems. A platform-independent, “write once, run anywhere” programming language.
- **Java VM** – the Java Virtual Machine. The Java VM (or JVM) is the part of the Java Runtime Environment (JRE) that is responsible for interpreting Java byte codes. See also *Java* and *JRE*.
- **JDBC** – an abbreviation for Java Database Connectivity. JDBC is the standard communication protocol for connectivity between Java clients and data servers. See also *data server* and *Java*.
- **JRE** – an abbreviation for Java Runtime Environment. The JRE consists of the Java Virtual Machine (Java VM or 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** – an abbreviation for “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** – an abbreviation for 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 *Log Transfer Manager*.

- **Log Transfer Interface** – an internal component of Replication Agent that interacts with Replication Server to forward transactions for distribution to Replication Server. See also *Log Reader* and *Log Transfer Manager*.
- **Log Transfer Language** – the proprietary protocol used between Replication Agent and Replication Server to replicate data from the primary database to Replication Server. See also *Log Reader* and *Log Transfer Interface*.
- **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 *Log Transfer Interface*.
- **maintenance user** – a special user login name in the replicate database that Replication Server uses to apply replicated transactions to the database. See also *replicate database* and *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*.
- **Multi-Path Replication™** – Replication Server feature that improves performance by enabling parallel paths of data from the source database to the target database. These multiple paths process data independently and are applicable when sets of data can be processed in parallel without transactional consistency requirements between them.
- **nonatomic materialization** – a materialization method that copies subscription data without a lock on the primary database. Changes to primary data are allowed during data transfer, which may cause temporary inconsistencies between the primary and replicate databases. Contrast with *atomic materialization*. See also *bulk materialization*.
- **ODBC** – an abbreviation for Open Database Connectivity, an industry-standard communication protocol for clients connecting to data servers. See also *client*, *data server*, and *JDBC*.
- **Open Client** – a Sybase product that provides customer applications, third-party products, and other Sybase products with the interfaces needed to communicate with Open Server applications. See also *Open Server*.
- **Open Client application** – An application that uses Sybase Open Client libraries to implement Open Client communication protocols. See also *Open Client* and *Open Server*.
- **Open Server** – a Sybase product that provides the tools and interfaces required to create a custom server. See also *Open Client*.
- **Open Server application** – a server application that uses Sybase Open Server libraries to implement Open Server communication protocols. See also *Open Client* and *Open Server*.
- **outbound queue** – a stable queue managed by 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** – in a replication system, a state in which all updates have been propagated to their destinations. Some Replication Agent and 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** – an abbreviation for 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** – an abbreviation for Replication Command Language, the command language used to manage Replication Server. See also *Replication Server*.
- **RDBMS** – an abbreviation for 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 a Replication Server for distribution to a replicate database. See also *primary database* and *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 Replication Server, includes information about the columns to be replicated and the location of the primary table or stored procedure. See also *Replication Server* and *subscription*.
- **Replication Server** – a Sybase software product that provides the infrastructure for a transaction replication system. See also *Replication Agent*.
- **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 Replication Server to a replicate Replication Server. Routes carry data-changing commands (including those for RSSDs) and replicated functions (database procedures) between separate Replication Servers. See also *Replication Server*.
- **RSSD** – an abbreviation for Replication Server System Database, which manages replication system information for a Replication Server. See also *Replication Server*.
- **SQL** – an abbreviation for 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 Replication Server. Messages written into the stable queue remain there until they can be delivered to the appropriate process or replicate database. Replication Server provides a stable queue for both incoming messages (the inbound queue) and outgoing messages (the outbound queue). See also *database connection*, *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 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 *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 *Replication Server*.
- **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** – an abbreviation for “wide area network,” a system of local-area networks (LANs) connected together with data communication lines. Contrast with *LAN*.

Index

A

accessing and downloading from SPDC 8, 25, 53, 81

ASE

- configuring Replication Server for replication from the primary database 16
- installing replication components 9
- installing Replication Server 10
- system requirements 7, 8
- testing replication 18

ASE-to-SAP HANA database replication
configuring Replication Server 14

C

cataloging a node and a database 89

configuring

- Primary Oracle database for replication 35
- Replication Agent for DB2 UDB 92
- Replication Agent for Microsoft SQL Server 62
- Replication Agent for Oracle 35
- replication components 13, 32
- replication components for Microsoft SQL Server 59
- Replication Server for ASE-to-SAP HANA database replication 14
- Replication Server for DB2 UDB replication 89
- Replication Server for Microsoft SQL Server 59
- Replication Server for Oracle-to-SAP HANA database replication 32

configuring ExpressConnect for HANA DB 14, 32, 59, 89

configuring Replication Agent

- configuring and validating the resource file 67
- configuring the primary DB2 UDB database for replication 92
- creating a Microsoft SQL Server user and granting permissions 62
- disabling Microsoft SQL Server replication 62
- enabling remote DAC 63

first-time initialization 70

obtaining Microsoft SQL Server DAC port number 63

setting the Autocommit parameter 93

subsequent initialization 72

verifying compatibility level of Microsoft SQL Server 62

verifying current archive setting of transaction log 92

verifying the Replication Agent replicating state 74

configuring Replication Server

- for replication from the primary database 14, 32, 59, 89
- for replication to SAP HANA database 14, 32, 59, 89

conventions

- style 1
- syntax 1

creating

- Oracle user 36
- Replication Agent instance 38
- Replication Agent instance in DB2 UDB 94
- Replication Server connection to the replicate database 14, 32, 59, 89
- Replication Server from primary database 48, 73, 102

D

DB2 UDB

- adding a temporary tablespace to the primary database 94
- configuring Replication Server 89
- configuring Replication Server for replication from the primary database 91
- creating a user and granting permissions 93
- creating the Replication Agent instance 94
- initializing the Replication Agent instance 100
- installing replication components 82
- resetting the primary database for replication 104
- setting environment variables 88
- system requirements 79
- testing DB2 UDB replication 102

Index

- testing replication 102
- verifying the Replication Agent instance 99

default directory 11, 27, 83

directories

- default 11, 27, 83

E

errors

- while starting installer 10, 26, 55, 82

existing installation directories 11, 27, 83

ExpressConnect for HANA DB

- configuring 14, 32, 59, 89

I

initializing

- Replication Agent instance 46
- Replication Agent instance in DB2 UDB 100

installation directories

- existing 11, 27, 83
- new 11, 27, 83

installation guidelines

- for Replication Agent 29, 57, 85
- for Replication Server 10, 26, 82
- Replication Server for Microsoft SQL Server 54

installation program

- Replication Agent 29, 57, 85

installer

- errors in starting 10, 26, 55, 82

installing 66

interfaces file

- editing 13

L

LogMiner

- configuration 36

M

Microsoft SQL Server

- configuring and verifying the primary database 62
- configuring Replication Server 59
- configuring Replication Server for replication from the primary database 61
- installing replication components 54

- installing Replication Server 54
- resetting the primary database for replication 77
- system requirements 52
- testing replication 75
- verifying the Replication Agent instance 71

Microsoft SQL Server JDBC driver 66

N

new installation directories 11, 27, 83

O

Oracle

- configuring Replication Server for replication from the primary database 34
- installing replication components 26
- installing Replication Server 26, 82
- system requirements 24
- testing replication 48
- verifying the Replication Agent instance 44

Oracle-to-SAP HANA database replication

- configuring Replication Server 32

P

planning your installation

- replication components 8, 24, 52, 80

prerequisites 53, 80

- ASE 8
- Oracle 25

prerequisites for installing replication components 52, 80

R

replication

- components 7, 23, 51, 79
- overview 7, 23, 51, 79

Replication Agent

- installing 29, 57, 85

Replication Agent for DB2 UDB

- configuring 92

Replication Agent for Microsoft SQL Server

- configuring 62

Replication Agent for Oracle

- configuring 35

Replication Agent installation directory 25, 53, 81

- replication components
 - configuration 59, 88
- Replication Server, adding
 - primary database 16
- resetting primary ASE database for replication 20
- resetting primary Oracle database for replication
 - 50

S

- setting up replication for DB2 UDB 79
- setting up replication from ASE to SAP HANA
 - database 7
- setting up replication from Oracle to SAP HANA
 - database 23, 51
- setup prerequisites 7, 24, 52, 79
- system requirements 24
 - ASE 7
 - DB2 UDB 79
 - Microsoft SQL Server 52
 - Oracle 24

T

- testing
 - ASE replication 18
 - Microsoft SQL Server replication 75
 - Oracle replication 48

V

- verifying
 - current archive setting of the redo logs 35
 - installation 31, 58, 87
 - primary Oracle configuration for replication
 - 38
 - Replication Agent instance in DB2 UDB 99
 - Replication Agent instance in Microsoft SQL
 - Server 71
 - Replication Agent instance in Oracle 44
 - supplemental logging of primary key data 35

