ASE-to-ASE Replication Quick Start Guide

SAP Replication Server® 15.7.1 SP200

DOCUMENT ID: DC01251-01-1571200-01
LAST REVISED: March 2014

Contents

Conventions

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

Style conventions

Key	Definition
`monospaced(fixed-width)`	• SQL and program code • Commands to be entered exactly as shown • File names • Directory names
`italic monospaced`	In SQL or program code snippets, placeholders for user-specified values (see example below).
italic	• 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 san serif	• 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 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

SAP® Replication Server® works with various components to enable replication between supported database such as, SAP® Adaptive Server® Enterprise (SAP® ASE), SAP HANA® database, SAP® IQ, Oracle, IBM DB2 UDB, and Microsoft SQL Server. SAP Replication Server uses SAP ASE for its Replication Server System Database (RSSD) or it uses SAP® SQL Anywhere® for its embedded Replication Server System Database (ERSSD).

Replication Agent™ is a generic term used to describe the Replication Agents for SAP ASE, SAP HANA database, Oracle, IBM DB2 UDB, and Microsoft SQL Server. The specific names are:

- RepAgent – Replication Agent thread for SAP ASE
- Replication Agent for Oracle
- Replication Agent for Microsoft SQL Server
- Replication Agent for UDB – for IBM DB2 on Linux, Unix, and Windows
- Replication Agent for DB2 for z/OS

Before You Begin

Learn about the prerequisites, guidelines, and assumptions for creating a simple replication environment.

Prerequisites

Become familiar with the prerequisites for setting up the replication system.

- You must be familiar with SAP ASE data servers and have an understanding of SAP replication.
- You must have a valid license for SAP Replication Server and SAP ASE.
- Use two different servers with the same named database. However, you can subsitute the replicate database hosted on a different server to a different database on the same server that hosts the primary database.
- TCP/IP connectivity must be available.

Obtaining a License

Obtain valid SySAM licenses before you begin the installation.

See *Obtaining a License* in the installation guide to learn more about SAP Replication Server product editions, and licensing options such sub-capacity licensing. For detailed instructions to configure SySAM licensing, see the *SySAM Users Guide* on the Product Documentation Web site at *http://sybooks.sybase.com*.

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 *System Requirements* in the installation guide for your platform.

Using the Examples to Set Up Your Replication Environment

Supplied examples walk you through setting up your replication environment.

The examples explain how to:

- Install SAP Replication Server
- Configure and verify the software installation
- Add the primary and replicate database to the SAP Replication Server
- Replicate data and DDL from a primary to a replicate database
- Manage warm standby applications
- Manage materialization

Follow these general instructions when using the examples to set up your replication environment:

- Some directories, files, executable commands, and examples are provided only for Solaris; adjust them accordingly for Windows and for other UNIX and Linux platforms.
- Do not use SAP Replication Server-reserved words for object names and connection names. The software also reserves all keywords and identifiers beginning with "rs_". See the *Reserved Words* in the *Reference Manual*.
- The values in the examples can be customized to meet the needs of your replication environment.

The examples make several assumptions:

- The simple replication environment uses three SAP Adaptive Server Enterprise (SAP ASE) databases and one SAP Replication Server. The primary database is the source of changes to be replicated; the replicate database is where changes are applied, and the Replication Server System Database (RSSD) holds the system tables for one SAP Replication Server. Each SAP Replication Server requires an RSSD, and in a production environment, the RSSD must be maintained like any other critical database or file.

  Note: The sample scenarios do not discuss the Embedded Replication Server System Database (ERSSD) database.
- You can have many SAP Replication Servers in your production environment. The examples discuss how to install and configure a single SAP Replication Server, which has a RSSD database in an existing SAP ASE.
- The replication environment uses the `pubs2` database, which is available in the standard SAP ASE installation. See *Installating Sample Databases* in the *SAP Adaptive Server Enterprise Installation Guide* for information on installing `pubs2` database.

Identifying the Installation Directory

The SAP **SYBASE** environment variable identifies the installation path for installing SAP Replication Server.

Select a path on a host drive to be the recipient of the installation and configuration activities. For example: `/software/sybase`

Downloading SAP Replication Server from SMP

Download SAP Replication Server from the SAP Service Marketplace (SMP).

Prerequisites

An SMP login ID and password.

Task

1. Go to *https://websmp204.sap-ag.de/,* and enter your login ID and password.
2. From the product listing, select SAP Replication Server; select the software you need by version and platform.
3. 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.
4. 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.
5. Download your installation files, and uncompress and extract all the installation images to your local drive.

Server Entries to the Interfaces Files

An interfaces file contains network definitions for each SAP Replication Server and data server in a replication system. The interfaces file is located in `$SYBASE/interfaces` on UNIX platforms, or `%SYBASE% \ini\sql.ini` on Microsoft Windows platforms.

When you install SAP Replication Server in an existing SAP Sybase installation directory, the installer adds SAP Replication Server information to the existing interfaces file. When you

install SAP Replication Server in a unique installation directory, the installer creates a new interfaces file. In this instance, you will have two interfaces files—one for your existing SAP applications and one for SAP Replication Server.

To allow the SAP Adaptive Server Enterprise (SAP ASE) and SAP Replication Server to communicate, you must:

- Add an entry for SAP Replication Server in the SAP interfaces file used by the SAP ASEs
- Add entries for the primary and replicate SAP ASEs in the SAP Replication Server interfaces file

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

For each server, you need:

- Server name
- Host name
- Port number

Adding Server Entries to the Interfaces Files

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

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

   On UNIX, enter:

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

   ```
   chaucer, 35356
   ```
7. Click **OK**.
8. Select **Server Object > Add**.
9. Enter the name of the replicate database server. Click **OK**.
10. Enter the host name and port number of the replicate database server. For example:

    ```
    johnson, 9756
    ```
11. Click **OK** twice, then select **File > Exit**.

Creating the Sybase User Account

To ensure that SAP product files and directories are created with consistent ownership and privileges, create a system administrator account.

A user, such as the system administrator— who has read, write, and execute privileges—must perform all installation and configuration tasks. For Windows, the user must have the administrator privilege.

1. To create system administrator account, choose an existing account, or create a new account and assign a user ID, group ID, and password for it.

   This account is sometimes called the “sybase” user account. See your operating system documentation for instructions on creating a new user account.

   If you have already installed other SAP software, the “sybase” user already exists.
2. Verify that you can log in to the machine using this account.

Installation and Configuration

Installation and configuration topics describe how to set up your SAP Replication Server and then configure it to suit your environment.

Mounting the Installation Media

If you are installing using the CD or DVD, mount the installation media.

The location of the **mount** command is site-specific and may differ from the instructions shown here. If you cannot mount the installation media in the appropriate drive using the path shown, check your operating system documentation or contact your system administrator.

- On HP-UX:
  Log out, log in as “root”, and issue:
  ```
  mount -F cdfs -o ro device_name /mnt/cdrom
  ```
  Then, log out as “root” and log in again as “sybase.”
- On IBM AIX:
  Log in as “sybase” and issue:
  ```
  mount -v 'cdrfs' -r device_name /mnt/cdrom
  ```
- On Solaris:
  The operating system automatically mounts the CD or DVD. Log in as “sybase.” If CD- or DVD-reading errors occur, check your operating system kernel to make sure that the ISO 9660 option is turned on. If you have previously installed a CD or DVD on your system, the # sign interferes with the installation process. Before installing your current CD or DVD, either:
  - Restart your system, or,
  - Eject the CD or DVD. Delete the *Volume Label* file in `/vol/dsk`, and reinsert the CD or DVD.
- On Linux and Linux on POWER:
  Log in as “sybase” and issue:
  ```
  # mount -t iso9660 /dev/cdrom /mnt/cdrom
  ```

Installing in GUI Mode

The installer creates a target directory and installs the selected components into that directory.

Prerequisites

Verify that the drive on which you install the SAP Replication Server has enough disk space for the components being installed, and at least 1GB of extra disk space for the installation program.

Task

1. Insert the installation media in the appropriate drive, or download and extract the SAP Replication Server installation image from the SAP Service Marketplace (SMP).
2. Launch the **setup** program.

   If the installer does not start automatically, double-click `setup.exe` or select **Start** > **Run** and enter the following, where *x* is your CD or DVD drive:

   ```
   x:\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 *directory_name*, include its full path.
3. If you downloaded the product from SMP, go to the directory where you extracted the installation image and start the installer:

   ```
   ./setup.bin
   ```

   If you are installing using the CD or DVD, mount it and start the installer.

   - On HP-UX:

     ```
     cd /cdrom
     ./setup.bin
     ```
   - On IBM AIX:

     ```
     cd /device_name
     ./setup.bin
     ```
   - On Solaris:

     ```
     cd /cdrom/Volume Label
     ./setup.bin
     ```
   - On Linux and Linux on POWER:

     ```
     cd /mnt/cdrom
     ./setup.bin
     ```

   where:

- *cdrom* and */mnt/cdrom* is the directory (mount-point) you specified when mounting the CD or DVD drive.
- *device_name* is the name of the CD or DVD device drive.
- `setup.bin` is the name of the executable file name for installing SAP 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 *tmp_dir*, include its full path.

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

   Do not use double-byte, single-, or double-quote characters in the installation path. The installer does not recognize these, and displays an error.

Option	Description
Click **Choose**.	Browse for and select the installation directory.
Enter a new directory path.	Create a new directory.
Click **Restore Default Folder**.	Use the default directory instead of the directory you have entered or selected.

   - If the directory you chose does not exist, click **Yes** to create it.
   - If it does exist, and already contains a current installation, you are warned that you will overwrite the older version. Click **Next**.

     The installer checks that the version you want to update is compatible with the version of the SAP Replication Server you are installing. If the version is incompatible, the Check Upgrade Incompatible Version dialog appears, and you see this:

     ```
     Warning: The current "SAP Replication Server" in your
     destination directory is not compatible with this version
     upgrade; some bug fixes may be unavailable
     if you proceed. See the release note for more information.
     ```

     You may see a similar message if your SAP Replication Server is an out-of-band release, such as an emergency bug fix, one-off, controlled, or instrumental release:

     ```
     Warning: The current "SAP Replication Server" in your
     destination directory is an out-of-band release; some bug fixes
     may be unavailable if you proceed.
     Verify that the bug fixes you need are in this newer version
     before proceeding with the upgrade.
     ```

     If you see such messages, click **Cancel** to stop the installation process. To override the error and continue with the installation, select **Proceed installation with incompatible version** and click **Next**.

Warning! Upgrading to an incompatible version may cause software regression. SAP recommends that you cancel the installation and obtain a compatible version of an SAP Replication Server.

If you perform the installation in silent, or unattended mode, and the versions are incompatible, the installer prompts you to re-run the installer with this argument, then quit:

```
DALLOW_UPGRADE_TO_INCOMPATIBLE_VERSION=true
```

6. Select the type of installation:

Option	Description
Typical	Installs the default components. This is recommended for most users.
Full	Installs every 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. In the SySAM License Server window:

Option	Choose
Specify license keys	Either: • Click **Browse** to select the license file. • Use **Shift+Click** or **Ctrl+Click** to select multiple license files. The license pane displays the license information. • 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. Select: • **Next** – to install the new SySAM license server. 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	Enter the: • Host name of the machine where the license server is running • Port number if the port number you are using is not the default Click **Next**.

Option	Choose
Continue installation without a license key	If you do not have licenses for any of the components, select this option, then click **Next** to proceed. The installer allows you to install and use the components without a license for a grace period of 30 days. To continue using these components after the end of the grace period, obtain valid licenses, and install these licenses using the SAP Replication Server license installer.

Use **sysadmin lmconfig** to configure and show license management-related information in the SAP Replication Server. See **sysadmin lmconfig** in the *Reference Manual*.

9. Set SySAM e-mail configuration if necessary. Click **Next**.
10. The installation summary window displays the selections you have made. Review the information, and click **Install**.
11. In the Start Sample Replication Server window, select:

Option	Description
Yes	Configure and start a sample Replication Server. The installer displays the configuration information for the sample Replication Server. Record this information. You can enter a maximum of 30 bytes in the password field. For a: • Single-byte character – enter 6 to 30 characters. • Double-byte character – enter 3 to 15 characters.
No	Manually configure a full-featured SAP 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.

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

Click **Next**.

12. Click **Done**.

Next

Verify that the installation is valid and successful.

- View the log files in the `log` directory to check for errors. A valid installation does not contain the word “ERROR.”
- Check that the date of the `si_reg.xml` file in the `Sybase_Install_Registry` directory reflects the date of the current installation.

Configuring an SAP Replication Server

Configure an SAP Replication Server using the values in the example.

Prerequisites

- Allocate a disk partition of at least 20MB for each SAP Replication Server you are installing. You can add more partitions later, if necessary.
- Ensure that the raw device or file system is available and has write permissions. If using a raw device, allocate the entire partition to the SAP Replication Server. You can add more space for the stable device later, if necessary. If you allocate only a part of the partition to SAP Replication Server, you cannot use the remainder for any other purpose.

Task

Each replication system has a single SAP Replication Server that manages the overall environment information. This SAP Replication Server is called an ID Server. A primary database is the source of replication and can only belong to one SAP Replication Server environment. A replicate database is the destination database.

This example creates a single SAP Replication Server, which has a Replication Server System Database (RSSD) in an existing SAP Adaptive Server Enterprise (SAP ASE). The script, in this example, creates the data and log devices that the RSSD database will use. In this example:

- PRS – name of the SAP Replication Server
- sunak1505i – name of the SAP ASE where the RSSD will reside

1. Go to the *$SYBASE* directory.
2. Add the primary SAP ASE, the replicate SAP ASE, the SAP Replication Server, and the server containing the RSSD database to the interfaces file.
3. Set the environment variables by sourcing the `SYBASE.csh` file.
4. Go to *$SYBASE*`/REP-15_5/init/rs`.
5. Make a copy of the `install.rs` file and rename it as `PRS.rs`.
6. Edit the parameter values in `PRS.rs` file.

Note: You need to modify the parameters described in the *Sample values for PRS.rs file* table.

Table 1. Sample Values for PRS.rs File

Parameter	Description	Values
sybinit.release_directory	Valid path where SAP Replication Server instance is created (`$SYBASE`).	`/opt/sybase`
rs.rs_idserver_name	Name of the SAP Replication Server.	PRS
rs.rs_id_server_is_rs_server	Specifies whether the SAP Replication Server created is the ID server or not.	yes
rs.rs_idserver_user	Because this is the ID server, user name that other SAP Replication Servers will use to connect to this SAP Replication Server. If this is not the ID server, then it is the user name of the ID server identified in the parameter, rs.rs_idserver_name.	PRS_id_user **Note:** The default name is the name of the SAP Replication Server followed by _id_user. For example, PRS_id_user.
rs.rs_idserver_pass	Password of the ID server user.	PRS_id_passwd
rs.rs_name	The name of the SAP Replication Server.	PRS
rs.rs_rs_sa_pass	Password for the sa user ID.	sa_pass **Note:** You can enter a maximum of 30 bytes in the password field. For: • single-byte character – enter 6 to 30 characters. • double-byte character – enter 3 to 15 characters.
rs.rs_needs_repagent	Specifies whether the RSSD needs a RepAgent. This is needed when multiple replication severs are used to replicate data from the primary to the replicate.	no **Note:** If the primary database and the replicate database will be replicating through multiple SAP Replication Servers, enter yes.
rs.rs_rssd_sqlsrvr	The name of the SAP ASE where the RSSD will reside.	sunak1505i

Parameter	Description	Values
rs.rs_rssd_db	Specifies the name of the RSSD database.	USE_DEFAULT **Note:** The default name is the name of the SAP Replication Server followed by _RSSD. For example, PRS_RSSD.
rs.rs_create_rssd	Specifies whether the script will create the RSSD database.	yes
rs.rs_rssd_sa_login	Specifies the name of the login that has "sa" privileges.	sa
rs.rs_rssd_sa_pass	Specifies the password of the login entered in the rs.rs_rssd_sa_login parameter.	password
rs.rs_rssd_prim_user	User that the SAP Replication Server uses to log into the RSSD.	PRS_RSSD_prim **Note:** The default name is the name of the RSSD database followed by _prim. For example, PRS_RSSD_prim.
rs.rs_rssd_prim_pass	Specifies the password of the login entered in the rs.rs_rssd_prim_user parameter.	PRS_RSSD_prim_ps
rs.rs_rssd_maint_user	User that the SAP Replication Server uses to log into the RSSD when getting work from other SAP Replication Servers.	PRS_RSSD_maint **Note:** The default name is the name of the RSSD database followed by _maint. For example, PRS_RSSD_maint.
rs.rs_rssd_maint_pass	Specifies the password of the login entered in the rs.rs_rssd_maint_user parameter.	PRS_RSSD_maint_ps
rs.rs_rsdddb_size	Specifies the size (in MB) of the system database device.	40
rs.rs_rssd_log_size	Specifies the size (in MB) of the RSSD database log device.	32

Parameter	Description	Values
rs.rs_rssd_db_device_name	The name of the SAP ASE device that stores the data portion of the RSSD database.	PRS_RSSD_data
rs.rs_create_rssd_data-base_dev	Specifies whether a new SAP ASE device needs to be created for the data portion of the RSSD database.	yes
rs.rs_rssd_db_device_path	Specifies the physical path for file system and file (or raw device) for the RSSD database data device.	/opt/sybase/ PRS_RSSD_data
rs.rs_rssddb_device_size	Specifies the size (in MB) of the RSSD database device.	40 The value must be the size of the data portion of the RSSD database specified in the **rs.rs_rsdddb_size** or greater.
rs.rs_rssd_log_device_name	Specifies the logical name of RSSD database log device.	PRS_RSSD_log
rs.rs_create_rssd_log_dev	Specifies whether the device for the log will be created.	yes **Note:** If the same device is used for both data and log, then enter no for this parameter.
rs.rs_rssd_log_device_path	Specifies the physical path for file system and file (or raw device) for the RSSD database log device.	/opt/sybase/ PRS_RSSD_log
rs.rs_rssd_log_device_size	Specifies the size (in MB) of the RSSD database log device.	32 The value must be the size of the log portion of the RSSD database specified in the **rs.rs_rsdddb_log_size** or greater.
rs.rs_diskp_name	Specifies the path and raw device (or file name) for the SAP Replication Server stable device. This file or device must already exist.	/opt/sybase/ PRSpart1.dat
rs.rs_diskp_lname	Specifies the logical name of the stable device partition.	part1

Parameter	Description	Values
rs.rs_diskp_size	Specifies the size, in MB, of the stable device partition.	20 The minimum size is 20MB.
rs.rs_rs_user	User that other SAP Replication Servers can use when a route between Replication Servers is created .	PRS_rsi **Note:** The default name is the name of the Replication Server followed by _rsi. For example, PRS_rsi.
rs.rs_rs_pass	Specifies the password of the login entered in the rs.rs_rs_user parameter.	PRS_rsi_ps
rs.rs_ltm_rs_user	Default user for Replication Agents to log into the SAP Replication Server.	PRS_ra **Note:** The default name is the name of the SAP Replication Server followed by _ra. For example, PRS_ra.
rs.rs_ltm_rs_pass	Specifies the password of the login entered in the rs.rs_ltm_rs_user parameter.	PRS_ra_ps

7. Save the file.
8. Go to the `$SYBASE/REP-15_5/install` directory.
9. Create the SAP Replication Server and the RSSD:

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

Note: If the **rs_init** command does not complete, check the log file in *$SYBASE/*`$SYBASE_REP/init/log`, correct the error, and resubmit the **rs_init**. If the problem occurred because the RSSD database had no space, then run the **isql** command on the data server and delete the RSSD database (for example, PRS_RSSD). If the **rs_init** command created the devices for the RSSD database, drop the devices (**sp_dropdevice**), and resubmit the **rs_init** command.

Verifying the Configuration

Verify that the SAP Replication Server installation directory and subdirectories (*%SYBASE%* on Microsoft Windows platforms and *$SYBASE* on UNIX platforms) have been created, and that other required software has been installed.

1. Change to the directory where you installed SAP Replication Server.
2. Set the environment variables by sourcing the `SYBASE.csh` file.
3. Log in to SAP Replication Server:

   ```
   isql -Usa -Psa_pass -SPRS
   ```

   PRS – name of the SAP Replication Server.

   Note: The default system administrator user ID is “sa”. The password to the "sa" user ID is the value that was filled in the resource file for rs.rs_rs_sa_pass.

4. Enter:

   ```
   admin who
   go

   The output from the command looks similar to:
   Spid   Name        State              Info
   ----   -----       -----              -----------
   27     DSI EXEC    Awaiting Command
   101                                   (1)sunak1505i.PRS_RSSD
   20     DSI         Awaiting Message   101 sunak1505i.PRS_RSSD
   26     SQM         Awaiting Message   101:0 sunak1505i.PRS_RSSD
   21     dSUB        Sleeping
   15     dCM         Awaiting Message
   18     dAIO        Awaiting Message
   23     dREC        Sleeping           dREC
   9      dDELSEG     Awaiting Message
   28     USER        Active             sa
   14     dALARM      Awaiting Wakeup
   24     dSYSAM      Sleeping
   ```

5. Verify the SAP Replication Server version. Enter:

   ```
   admin version
   go
   ```

   Note: The listing from the command looks similar to:

   ```
   Version
   ------------------------------------------------
   Replication Server/15.7.1/P/Sun_svr4/OS 5.8/1/OPT64/Sun
   Apr 22 18:37:00 2012
   ```

6. Exit the **isql** session using:

   ```
   quit
   ```

Manage SAP Replication Server

Manage SAP Replication Servers and related components in the replication environment.

Starting an SAP Replication Server

Start the SAP Replication Server manually if it is shut down.

1. Go to the `$SYBASE` directory and set the SYBASE environment variables.

   ```
   cd /opt/sybase
   source SYBASE.csh
   ```

2. Go to the directory that contains the startup file for the SAP Replication Server.

   The typical directory for UNIX is `$SYBASE/REP*/install`. The directory for the SAMPLE_RS is `$SYBASE/REP*/samp_repserver`

3. For Windows, double-click on the `RUN*.bat` file. For UNIX, execute the `RUN` file.

Stopping an SAP Replication Server

Shut down an SAP Replication Server.

1. Go to the `$SYBASE` directory and set the SYBASE environment variables.

   ```
   cd /opt/sybase
   source SYBASE.csh
   ```

2. Log in with **isql** to connect to the SAP Replication Server.

   ```
   isql –Usa –Psa_pass –Sreplication_server
   ```

   where:

   - U – is a user id with sa privileges
   - P – is the password for the user
   - S – is the SAP Replication Server. Alternatively, one can use the hostname and port number with a colon in-between, for example: "chaucer:35356"

3. Issue the **shutdown** command.

   ```
   shutdown
   go
   ```

Viewing the SAP Replication Server Logs

View a record of the SAP Replication Server actions. Viewing the log can help you identify problems.

Note: The `SAMPLE_RS.log` file is available if you choose to install and start sample Replication Server during installation.

Use a text editor to open the `SAMPLE_RS.log` file.

The log file used at startup can be found in the -E parameter of the `RUN` file. For instance, given the following start command in the RUN_SAMPLE_RS file:

- `/opt/sybase/REP-15_5/bin/repserver -SSAMPLE_RS \`
- `-C/opt/sybase/REP-15_5/samp_repserver/SAMPLE_RS.cfg \`
- `-E/opt/sybase/REP-15_5/samp_repserver/SAMPLE_RS.log \`
- `-I/opt/sybase/interfaces`

The log file can be found at: `/opt/sybase/REP-15_5/samp_repserver/SAMPLE_RS.log`

All log files for the sample Replication Server are located in the `samp_repserver` directory. SAP Replication Server records its actions and notes failures in `SAMPLE_RS.log`

Multisite Availability

Multisite Availability (MSA) extends SAP Replication Server replication capabilities and simplifies setting up a replication system.

Key features of MSA include:

- A simple replication methodology that requires only one replication definition for the primary database and only one subscription for each subscribing database.
- A replication filtering strategy that lets you choose whether or not to replicate individual tables, transactions, functions, system stored procedures, and data definition language (DDL).
- Replication of DDL to any replicate database—including non–warm standby databases. Replication to multiple replicate sites—for standby as well as nonstandby databases.

Multisite Availability Diagram

The diagram illustrates a simple example of multisite availability replication.

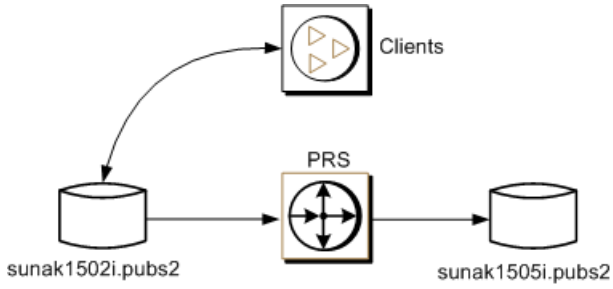

Adding the Primary Database to the SAP Replication Server

Add a primary database to the SAP Replication Server for MSA.

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

   Note: Restart the SAP ASE 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 `primary_pubs2.rs`.
4. Edit the `primary_pubs2.rs file`.

Note: You need to modify the parameters described in the sample values table.

Table 2. Sample Values for the primary_pubs2.rs File

Parameter	Description	Value
sybinit.release_directory	Valid path of the SAP Replication Server software ($SYBASE).	/opt/sybase
rs.rs_name	Name of the SAP Replication Server.	PRS
rs.rs_rs_sa_user	Specifies the user ID that has "sa" privileges on SAP Replication Server.	sa
rs.rs_rs_sa_pass	Specifies the password of the "sa" user.	sa_pass
rs.rs_ds_name	Name of the data server that hosts the primary database.	sunak1502i
rs.rs_ds_sa_user	Specifies the user ID that has "sa" privileges on data server.	sa
rs.rs_ds_sa_password	Specifies the password of the "sa" user for the data server.	password
rs.rs_db_name	Specifies the name of the primary database.	pubs2
rs.rs_needs_repagent	Specifies whether you plan to replicate from specified primary database.	yes
rs.rs_db_maint_user	Specifies the 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 cannnot be the name of an alias.
rs.rs_db_maint_password	Specifies the password for the maintenance user.	<database>_maint_ps

Parameter	Description	Value
rs.rs_ltm_rs_user	Specifies the user that the Replication Agent will use to log into the SAP Replication Server. The name must exist. This name typically comes from values that were set up during SAP Replication Server creation time: rs.rs_ltm_rs_user.	PRS_ra
rs.rs_ltm_rs_pass	Specifies the password of the rs.rs_ltm_rs_user.	PRS_ra_ps
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 SAP Replication Server by running the resource file:

   ```
   ./rs_init -r ../init/rs/primary_pubs2.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 SAP ASE 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 -S PRS
   ```

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)sunak1505i.PRS_RSSD
20       DSI        Awaiting Message   101 sunak1505i.PRS_RSSD
26       SQM        Awaiting Message   101:0 sunak1505i.PRS_RSSD
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.

Adding the Replicate Database to the SAP Replication Server

Add a replicate database to the SAP Replication Server for MSA.

1. Add the replicate SAP ASE to the `interfaces` file of the SAP Replication Server.

Note: Restart the SAP Replication Server 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 `replicate_pubs2.rs`.

4. Edit the `replicate_pubs2.rs file`.

Note: You need to modify the parameters described in the sample values table.

Table 3. Sample Values for the replicate_pubs2.rs file

Parameter	Description	Value
sybinit.release_directory	Valid path of the SAP Replication Server software ($SYBASE).	/opt/sybase
rs.rs_name	Name of the SAP Replication Server.	PRS
rs.rs_rs_sa_user	Specifies the user ID that has “sa” privileges on SAP Replication Server.	sa
rs.rs_rs_sa_pass	Specifies the password of the “sa” user.	sa_pass
rs.rs_ds_name	Name of the data server that hosts the replicate database.	sunak1505i

Parameter	Description	Value
rs.rs_ds_sa_user	Specifies the user ID that has "sa" privileges on data server.	sa
rs.rs_ds_sa_password	Specifies the password of the "sa" user for the data server.	password
rs.rs_db_name	Specifies the name of the replicate database.	pubs2
rs.rs_needs_repagent	Specifies whether you plan to replicate from specified rs.rs_db_name.	no
rs.rs_db_maint_user	Specifies the user ID who will apply the work at the replicate database. The user is called the maintenance user.	<database>_maint **Note:** If the user ID does not exist the script creates the user ID on the replicate database. The user ID cannnot be the name of an alias.
rs.rs_db_maint_password	Specifies the password for the user specified in rs.rs_db_maint_user.	<database>_maint_ps
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 SAP Replication Server by running the resource file:

   ```
   ./rs_init -r ../init/rs/replicate_pubs2.rs
   ```

   If the **rs_init** command fails, correct the issue, and re-run the **rs_init**.
8. Validate the primary connection:

   ```
   isql -Usa -Psa_pass -S<PRS>
   ```
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) sunak1505i.PRS_RSSD
20      DSI         Awaiting Message     101 sunak1505i.PRS_RSSD
26      SQM         Awaiting Message     101:0 sunak1505i.PRS_RSSD
55      DSI EXEC    Awaiting Command     103(1) sunak1505i.pubs2
54      DSI         Awaiting Message     103 sunak1505i.pubs2
53      SQM         Awaiting Message     103:0 sunak1505i.pubs2
21      dSUB        Sleeping
15      dCM         Awaiting Message
18      dAIO        Awaiting Message
23      dREC        Sleeping             dREC
9       dDELSEG     Awaiting Message
56      USER        Active               sa
14      dALARM      Awaiting Wakeup
24      dSYSAM      Sleeping
```

10. Exit the `isql` session.

Marking the Primary Database for Replication

Use database replication definition and subcription to replicate the entire primary database.

1. Log in to the primary database with system administrator privileges:

```
% isql -Usa -P -Ssunak1502i
```

2. Connect to the pubs2 database:

```
use pubs2
go
```

3. Mark the primary database for replication. For example:

```
sp_reptostandby pubs2, 'all'
go
```

4. Set the RepAgent parameter **send warm standby xacts** to true so that RepAgent sends Data Manipulation Transactions (DML) and Data Definition Language (DDL) to the replicate database. For example, at the primary data server, enter:

```
sp_config_rep_agent
pubs2,send_warm_standby_xacts,true
go

Parameter_Name Default_Value Config_Value Run_Value
-------------- --------- ---------------- ---------
send warm standby xacts false  true           true
(1 row affected)

RepAgent configuration changed for database pubs2.
The changes will take effect the next time the
RepAgent thread is started. (return status = 0)
```

5. Stop and restart the RepAgent:

```
sp_stop_rep_agent pubs2
go
```

```
sp_start_rep_agent pubs2
go
```

6. Exit the **isql** session.
7. At the Replication Server, create a database replication definition that also replicates DDL:

```
isql -Usa -Psa_pass -SPRS
create database replication definition name
with primary at pds.pdb
replicate DDL
```

where:

- *name* – is the unique identifer for this replication definition.
- *pds* – is the name of the primary ASE data server.
- *pdb* – is the name of the primary database.

For example:

```
create database replication definition pubs2_repdef
with primary at sunak1502i.pubs2
replicate DDL
go

Database replication definition pubs2_repdef for
sunak1502i.pubs2 is created.
```

8. Create a database subscription for the replicate database. This example creates a database subscription that uses no materialization method and will replicate the truncate table command:

```
create subscription sub_name
for database replication definition name
with primary at pds.pdb
with replicate at rds.pdb
without materialization
subscribe to truncate table
```

where:

- *sub_name* – is the unique identifer for this subscription.
- *name* – is the unique identifer for the replication definition.
- *pds* – is the name of the primary ASE data server.
- *pdb* – is the name of the primary database.
- *rds* – is the name of the replicate ASE data server.

for example:

```
create subscription pubs2_sub
for database replication definition pubs2_repdef
with primary at sunak1502i.pubs2
with replicate at sunak1505i.pubs2
without materialization
subscribe to truncate table
go
```

Note: For the subscription to be successfully created or dropped, the connection to the replicate database must be available.

9. Check the subscription status at the primary and replicate data servers:

```
check subscription pubs2_sub
for database replication definition pubs2_repdef
with primary at sunak1502i.pubs2
with replicate at sunak1505i.pubs2
go
```

The status shows:

```
Subscription pubs2_sub is VALID at the replicate.
Subscription pubs2_sub is VALID at the primary.
```

The database is now ready for replication

Note: If the subscription results in errors, you can drop the subscription using the **drop subscription** command. For example:

```
drop subscription pubs2_sub
for database replication definition pubs2_repdef
with primary at sunak1502i.pubs2
with replicate at sunak1505i.pubs2
without purge
go
```

The replicate connection must be available before you can drop a subscription.

Viewing Information About Database Replication Definitions

Use **rs_helpdbrep** to view information about a specific database replication definition or all database replication definitions for a database or a data server.

For example, to view information about the **rep_1B** database replication definition, enter:

```
rs_helpdbrep rep_1B, PDS, pdb
```

See ***rs_helpdbrep*** in the *Reference Manual* for syntax and usage information.

Viewing Information About Database Subscriptions

Use **rs_helpdbsub** to view information about a specific database subscription or all database replication definitions for a database or a data server.

For example, to view information about the **sub_2B** database subscription, enter:

```
rs_helpdbsub sub_2B, dsA, db
```

See *Replication Server Reference Manual* > *RSSD Stored Procedures* > ***rs_helpdbsub*** *for syntax and usage information*.

Replicating DDL

Replicate DDL.

1. Log in to the primary database using a user ID that exists on both the primary and replicate data server and that has permission to create a table on both the primary and replicate databases.

   Note: Do not use the same maintenance user ID that you used to set up the primary connection.

2. Create a new table:

   ```
   % isql -Usa -P -Ssunak1502i
   use pubs2
   go
   create table t1 (a char(10), b integer, c text)
   go
   ```

3. Create unique indexes to ensure data integrity:

   ```
   create unique clustered index t1_idx1 on t1 (a,b)
   go
   ```

4. Log in to the replicate database:

   ```
   % isql -Usa -P -Ssunak1505i
   use pubs2
   go
   ```

5. Verify that the table and index exists in the replicate database. If the table and index do not exist, follow the instructions in steps 6 and 7 otherwise, go to step 8.

6. Check the SAP Replication Server log file at: *$SYBASE*/REP- 15_5/install/PRS.log. Correct the errors and restart the connection to the replicate database:

   ```
   resume connection to rds.rdb
   go
   ```

   where:

   - *rds* – is the name of the data server that hosts the replicate database.
   - *rdb* – is the name of the replicate database.

   Note: Make sure that the user making the corrections is not same the maintenance user that was used to set up the primary connection.

7. If you want SAP Replication Server to skip any current transactions when trying to resume connection to the replicate database, use:

   ```
   resume connection to rds.rdb
   skip transaction
   go
   ```

See ***resume connection*** in the *Reference Manual* for other available **resume connection** options.

Note: If you see this message:

```
"Message from server: Message: 2762, State 3, Severity 16 - 'The
'CREATE TABLE' command is not allowed within a multi-statement
transaction in the 'pubs2 database."
```

Ensure that the Replication Agent has been configured to send warm standby transactions and that the RepAgent has been stopped and restarted since the last time when the RepAgent configuration parameter was changed.

8. Grant **insert**, **update**, and **delete** permission for the new table to the replicate database maintenance user at the replicate database:

```
grant all on t1 to pubs2_maint
go
```

Replicating DML

Replicate DML.

1. Log in to the primary database using a user ID that has permission to **insert**, **update**, **delete** and **truncate** a table. See *Manage the Maintenance User* in the *Administration Guide Volme 1* for information on how to grant permission to the maintenance user ID.

   Note: Do not use the same maintenance user ID that you have used to set up the primary connection.

2. In the primary database, insert a row to t1:

```
insert into t1 values('a',1,'this is the first row')
go
```

3. Check whether the row exists in the replicate database:

```
select * from t1
go
```

   If the row does not exist, follow instructions in steps 4 and 5 otherwise, go to step 6.

4. Check the SAP Replication Server log file at: *$SYBASE*/REP- 15_5/install/PRS.log. Correct the errors and restart the connection to the replicate database:

```
resume connection to rds.rdb
go
```

5. If you want SAP Replication Server to skip any current transaction when trying to resume connection to the replicate database, use:

```
resume connection to rds.rdb
skip transaction
go
```

See ***resume connection*** in the *Reference Manual* for other available options for the **resume connection** command.

6. Log in to the primary database and update the row:

```
update t1 set c = 'this is an update' where b = 1
go
```

7. Log in to the replicate database and verify that the row was updated:

```
select * from t1
go
```

8. Log in to the primary database and enter:

```
truncate table t1
go
```

9. Log in to the replicate database and enter:

```
select count (*) from t1
go
```

   The number of rows at the replicate table, t1, should now be zero.

Warm Standby Applications

A warm standby application is a pair of SAP Adaptive Server Enterprise (SAP ASE) databases, one of which is a backup copy of the other.

Client applications update the active database; SAP Replication Server maintains the standby database as a copy of the active database. If the active database fails, or if you need to perform maintenance on the active database or on the data server, a switch to the standby database allows client applications to resume work with little interruption. Warm standby applications topics describe how to set up and configure a warm standby application between two SAP ASE databases—the active database and the standby database.

See *Manage Warm Standby Applications* in the *Administration Guide Volume 2* for detailed information about warm standby applications, how it works in SAP Replication Server, and other related topics.

Most of the times in SAP Replication Server, databases are defined as "primary" and "replicate." However, in discussing warm standby applications databases are also defined as "active" and "standby."

Warm Standby Application Diagrams

Learn how a warm standby works. The diagrams illustrate the normal operation, and switching of active and standby databases of an example warm standby application.

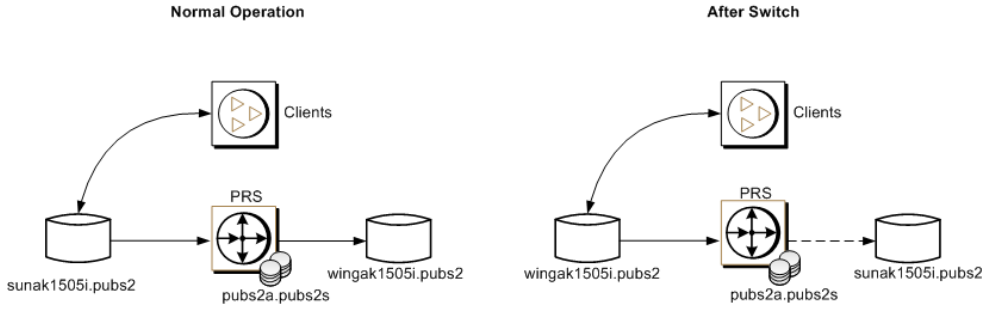

Creating a Logical Connection

Create a logical connection, which establishes one symbolic name for both the active and standby databases.

In a warm standby application, the active database and the standby database appear in the replication system as a connection from the SAP Replication Server to a single logical database. A logical connection is created to establish one symbolic name for both the active and standby databases. The names do not have to exist in the `interfaces` file nor in the replication system.

1. Log in to the SAP Replication Server:

   ```
   isql -Usa -Psa_pass -SPRS
   ```

2. Create the logical connection for the warm standby:

   ```
   create logical connection to pubs2a.pubs2s
   go
   ```

3. Enable SQL statement replication:

   ```
   alter logical connection to pubs2a.pubs2s
   set ws_sqldml_replication to "on"
   go
   ```

   Note: SQL statement replication is only available in SAP ASE 15.0.3 and later.

4. Exit the **isql** session.

Adding the Active Database to the SAP Replication Server

Add the active database.

1. Add entries for SAP Replication Server and SAP ASE that host the primary or active database to the corresponding `interfaces` files.

   Note: Restart the SAP ASE or SAP Replication Server if the `interfaces` file has changed.

2. Go to *$SYBASE*`/REP-15_5/init/rs`.
3. Make a copy of the `setupdb.rs` file and rename it as `active_pubs2.rs`.
4. Edit the `active_pubs2.rs` file.

Table 4. Sample Values for the active_pubs2.rs File

Parameter	Description	Value
sybinit.release_directory	Specifies the valid path of the SAP Replication Server software ($SYBASE).	/opt/sybase
rs.rs_name	Specifies the name of the SAP Replication Server	PRS
rs.rs_rs_sa_user	Specifies the user ID that has "sa" privileges on SAP Replication Server.	sa
rs.rs_rs_sa_pass	Specifies the password of the "sa" user.	sa_pass
rs.rs_ds_name	The name of the data server that hosts the primary database.	sunak1505i
rs.rs_ds_sa_user	Specifies the user ID that has "sa" privileges on data server.	sa
rs.rs_ds_sa_password	Specifies the password of the "sa" user for the data server.	password
rs.rs_db_name	Specifies the name of the primary database	pubs2
rs.rs_needs_repagent	Specifies whether you plan to replicate from specified primary database.	yes
rs.rs_db_maint_user	Specifies the user ID that performs **insert**, **update**, **delete**, and **truncate** table commands to the replicate database. The user ID must have permissions to perform DML commands at the replicate database. The user is called maintenance user. **Note:** Do not use a user ID that has been aliased to another user ID.	USE_DEFAULT The default value is *database-name*_maint.
rs.rs_db_maint_password	Specifies the password for the maintenance user.	<database>_maint_ps

Parameter	Description	Value
rs.rs_ltm_rs_user	The user that the Replication Agent will use to log into the SAP Replication Server. The name must exist. This name typically comes from values that were set up during RS creation time: rs.rs_ltm_rs_user.	PRS_ra
rs.rs_ltm_rs_pass	Specifies the password of the login entered in the rs.rs_ltm_rs_user parameter.	PRS_ra_ps
rs.rs_db_physical_for_logical	Specifies whether this is a warm standby database.	yes
rs.rs_db_active_or_standby	Specifies whether the configuration is for active or standby database.	active
rs.rs_db_logical_ds_name	Specifies the data server portion of the logical connection name.	pubs2a
rs.rs_db_logical_db_name	Specifies the database portion of the logical connection name.	pubs2s

5. Save the file.
6. Go to *$SYBASE*/REP-15_5/install.
7. Create the connection from the active database to the SAP Replication Server by running the resource file:

   ```
   ./rs_init -r ../init/rs/active_pubs2.rs
   ```

   If the **rs_init** command fails, correct the issue, then disable the RepAgent as follows:

   a) Log in to the primary SAP ASE using an "sa" user role and access the primary database
   b) Disable the RepAgent thread in the primary database:

   ```
   isql -Usa -P -Ssunak1505i
   use pubs2
   go
   sp_config_rep_agent pubs2,'disable'
   go
   ```

   Re-run **rs_init** command
8. Validate the primary connection:

   ```
   isql -Usa -Psa_pass -SPRS
   ```
9. Enter:

```
admin who
go
The output generated from admin who is similar to
Spid        Name       State             Info
------      -------    ------------      -----------------
36          DIST       Awaiting Wakeup   102 pubs2a.pubs2s
37          SQT        Awaiting Wakeup   102:1 DIST pubs2a.pubs2s
31          SQM        Awaiting Message  102:1 pubs2a.pubs2s
30          SQM        Awaiting Message  102:0 pubs2a.pubs2s
27          DSI EXEC   Awaiting Command  101(1)sunak1505i.PRS_RSSD
20          DSI        Awaiting Message  101 sunak1505i.PRS_RSSD
26          SQM        Awaiting Message  101:0 sunak1505i.PRS_RSSD
49          DSI EXEC   Awaiting Command  103(1)sunak1505i.pubs2
35          DSI        Awaiting Message  103 sunak1505i.pubs2
38          REP AGENT  Awaiting Command  sunak1505i.pubs2
39          NRM        AwaitingMessage   sunak1505i.pubs2
21          dSUB       Sleeping
15          dCM        Awaiting Message
18          dAIO       Awaiting Message
23          dREC       Sleeping dREC
9           dDELSEG    Awaiting Message
29          USER       Active sa
14          dALARM     Awaiting Wakeup
24          dSYSAM     Sleeping
```

Note: The RepAgent for the active database must be available.

10. Validate the status of the active database connection:

```
admin logical_status
go
```

The output generated from **admin logical_status** is similar to:

Logical Connection Name	Active Connection Name	Active Conn State	Standby Connection Name	Standby Conn State
[102] pubs2a.pubs2s	[103] sunak1505i.pubs2	Active/	None	None

Controller RS	Operation in Progress	State of Operation in Progress	Spid
[16777317] PRS	None	None	

Marking the Active Database for Replication

Use **sp_reptostandby** or **sp_setreptable** to replicate tables in the active database.

The database is enabled for both SQL statement replication (available in Adaptive Server 15.0.3 and later) and non-SQL statement replication.

You can enable replication for tables in the active database in either of these ways:

- Use **sp_reptostandby** to mark the database for replication, enabling replication of data and supported schema changes or,
- Use **sp_setreptable** to mark individual tables for replication of data changes.

In this example, we have used **sp_reptostandby** to mark the database for replication.

1. Log in to the Adaptive Server as the system administrator:

```
% isql -Usa -P -Ssunak1505i
use pubs2
go
```

2. Mark the database tables for replication for both DDL and DML commands and procedures:

```
sp_reptostandby pubs2,'all'
go

The replication mode for database 'pubs2' has been
set to 'ALL'.
(return status = 0)
```

3. Mark the database to send SQL statements for **update**, **delete**, **insert**, and **select** into commands, if the SQL statement threshold has been met:

```
sp_setrepdbmode pubs2,'UDIS','on'
go

The replication mode for database 'pubs2' is 'udis'.
(return status = 0)
```

4. Set the database threshold for SQL statement replication to 10.

```
sp_setrepdbmode pubs2,'threshold','10'
go

The replication threshold for 'pubs2' is '10'.
(return status = 0)
```

Note: You can also set the threshold per table:

```
sp_setrepdefmode t1,'threshold','10'
```

5. Exit the **isql** session.

Adding the Standby Database to the SAP Replication Server

Add the standby database. The example procedure for configuring warm standby uses a dump of the active database to load the standby database and starts replication after the setup is complete.

1. Add entries for SAP Replication Server and SAP ASE that host the replicate database to the corresponding `interfaces` files.

   Note: Restart the SAP ASE and SAP Replication Server if the `interfaces` file has changed.

2. Add the maintenance user login name for the standby connection in the standby data server.
3. Go to *`$SYBASE`*`/REP-15_5/init/rs`.
4. Make a copy of the `setupdb.rs` file and rename it as `standby_pubs2.rs`.
5. Edit the `standby_pubs2.rs` file.

Table 5. Sample Values for the standby_pubs2.rs File

Parameter	Description	Value
sybinit.release_directory	Specifies the valid path of the SAP Replication Server software ($SYBASE).	/opt/sybase
rs.rs_name	Specifies the name of the SAP Replication Server	PRS
rs.rs_rs_sa_user	Specifies the user ID that has “sa” privileges on SAP Replication Server.	sa
rs.rs_rs_sa_pass	Specifies the password of the “sa” user.	sa_pass
rs.rs_ds_name	The name of the data server that hosts the standby database.	wingak1505i
rs.rs_ds_sa_user	Specifies the user ID that has “sa” privileges on data server.	sa
rs.rs_ds_sa_password	Specifies the password of the “sa” user for the data server.	password
rs.rs_db_name	Specifies the name of the standby database	pubs2

Parameter	Description	Value
rs.rs_needs_repagent	Specifies whether you plan to replicate from specified standby database.	yes
rs.rs_db_maint_user	Specifies the user ID that performs **insert**, **update**, **delete**, and **truncate** table commands to the replicate database. The user ID must have permissions to perform DML commands at the replicate database. The user is called maintenance user.	The maintenance user (rs.rs_db_maint_user) that was defined for the active database. Using the same user for the active and standby database simplifies synchronizing the server user IDs after the database dump is loaded to the standby server.
rs.rs_db_maint_password	Specifies the password for the maintenance user.	The password for the user specified rs.rs_db_maint_user
rs.rs_ltm_rs_user	The user that the Replication Agent will use to log into the SAP Replication Server. The name must exist. This name typically comes from values that were set up during SAP Replication Server creation time: rs.rs_ltm_rs_user.	PRS_ra
rs.rs_ltm_rs_pass	Specifies the password of the login entered in the rs.rs_ltm_rs_user parameter.	PRS_ra_ps **Warning!** When creating the replicate connection, comment out the line for this parameter. Otherwise, the resource file will not work.
rs.rs_db_physical_for_logical	Specifies whether this is a warm standby database.	yes
rs.rs_db_active_or_standby	Specifies whether the configuration is for active or standby database.	standby
rs.rs_db_logical_ds_name	Specifies the data server portion of the logical connection name.	pubs2a
rs.rs_db_logical_db_name	Specifies the database portion of the logical connection name.	pubs2s

Parameter	Description	Value
rs.rs_db_active_ds_name	Specifies the server that hosts the standby database.	sunak1505i
rs.rs_db_active_db_name	Specifies the name of the active database.	pubs2
rs.rs_db_active_sa	Specifies the user ID that has "sa" privileges on the active database.	sa
rs.rs_db_active_sa_pw	Specifies the password of the "sa" user.	password
rs.rs_init_by_dump	Specifies that the standby database is initialized using a dump of the active database.	yes
rs.rs_db_use_dmp_marker	Specifies that "dump marker" option is used to notify replication when to begin forwarding transactions to the standby database.	yes

6. Save the file.
7. Go to *$SYBASE*/$SYBASE_REP/install.
8. Create the connection from the Replication Server to the standby database by running the resource file:

   ```
   ./rs_init -r ../init/rs/standby_pubs2.rs
   ```

   If the **rs_init** command fails, 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:

   ```
   isql -Usa -P -Swingak1505i
   use pubs2
   go
   sp_config_rep_agent pubs2,'disable'
   go
   ```

   Re-run **rs_init** command
9. Validate the replicate or standby connection:

   ```
   isql -Usa -P<sa_pass> -S<PRS>
   ```

10. Enter:

    ```
    admin who
    go
    The output generated from admin who is similar to
    Spid    Name       State               Info
    ```

```
------  -------     ------------        -----------------
36      DIST        Awaiting Wakeup     102 pubs2a.pubs2s
37      SQT         Awaiting Wakeup     102:1 DIST pubs2a.pubs2s
31      SQM         Awaiting Message    102:1 pubs2a.pubs2s
30      SQM         Awaiting Message    102:0 pubs2a.pubs2s
27      DSI EXEC    Awaiting Command    101(1)sunak1505i.PRS_RSSD
20      DSI         Awaiting Message    101 sunak1505i.PRS_RSSD
26      SQM         Awaiting Message    101:0 sunak1505i.PRS_RSSD
49      DSI EXEC    Awaiting Command    103(1)sunak1505i.pubs2
35      DSI         Awaiting Message    103 sunak1505i.pubs2
38      REP AGENT   Awaiting Command    sunak1505i.pubs2
39      NRM         Awaiting Message    sunask1505i.pubs2
        DSI EXEC    Suspended           104(1)wingak1505i.pubs2
        DSI         Suspended           104 wingak1505i.pubs2
21      dSUB        Sleeping
15      dCM         Awaiting Message
18      dAIO        Awaiting Message
23      dREC        Sleeping dREC
9       dDELSEG     Awaiting Message
29      USER        Awaiting Command    sa
55      USER        Active              sa
14      dALARM      Awaiting Wakeup
24      dSYSAM      Sleeping
```

Note: The RepAgent for the active database must be available.

11. Validate the status of the active database connection:

```
admin logical_status
go
```

The output generated from **admin logical_status** is similar to:

Logical Connection Name	Active Connection Name	Active Conn State	Standby Connection Name	Standby Conn State
[102] pubs2a.pubs2s	[103] sunak1505i.pubs2	Active/	[104] wingak1505i.pubs2	Suspended / Waiting for Enable Marker

Controller RS	Operation in Progress	State of Operation in Progress	Spid
[16777317] PRS	None	None	

12. Exit the **isql** session

Initializing the Standby Database

Use SAP Adaptive Server Enterprise (SAP ASE) commands and utilities to initialize the standby database.

This example uses the "dump marker" option to initially materialize the standby database. See *Manage Database Connections* in the *Administration Guide Volume 1* for information on how to grant permissions to the maintenance user ID for all the tables in the database.

1. Log in to the active data server and dump the active database:

   ```
   dump database pubs2 to '/backup/data/sybase1550/ASE-
   15_5/pubs2.dmp'
   ```

   Note: Make sure that the Backup Server is running. Dumps and loads are performed through Backup Server.

2. Exit from the active database.
3. Load the standby database with the dump from the active database:

   ```
   load database pubs2 from
   '/backup/data/sybase1550/ASE-15_5/pubs2.dmp'
   ```

4. After completing the load operations, bring the standby database online:

   ```
   online database pubs2
   go
   Started estimating recovery log boundaries for
   database 'pubs2'.
   Database 'pubs2', checkpoint=(1564, 65),
   first=(1564, 65), last=(1565, 17).
   Completed estimating recovery log boundaries for
   database 'pubs2'.
   Started ANALYSIS pass for database 'pubs2'.
   Completed ANALYSIS pass for database 'pubs2'.
   Recovery of database 'pubs2' will undo incomplete
   nested top actions.
   Database 'pubs2' is now online
   ```

5. Check the "suid" for the maintenance user defined for the standby database at the server level:

   ```
   use master
   go

   select suid,name from syslogins
   where name ='pubs2_maint'
   go

   suid     name
   ------- -----------
   3        pubs2_maint
   ```

6. Check the "suid" for the maintenance user defined in the standby database:

```
use pubs2
go

select suid,name from sysusers
where name = 'pubs2_maint'
go

suid     name
------- -----------
8        pubs2_maint
```

Note: If the suid for the maintenance user does not exist, add it with the **sp_addlogin** command.

7. Change the "suid" in the `sysusers` table in the standby database to match the "suid" in the SAP ASE server that hosts the standby database:

```
sp_configure "allow updates to system tables",1
go
```

Parameter	Default	Memory Used	Config Value	Run Value	Unit	Type
allow updates to system tables	0	0	1	1	switch	dynamic

```
Configuration option changed. ASE need not be rebooted
since the option is dynamic.
Changing the value of 'allow updates to system tables'
does not increase the amount of memory Adaptive Server
uses.
(return status = 0)
```

```
update sysusers set suid = 3 where name = "pubs2_maint"
go
(1 row affected)
```

Note: If there is a difference in the server user IDs (suids) assigned to the users at the active database versus the standby database, modify the sysusers table in the newly loaded database to match both the logins.

If the maintenance user does not exist on the database, add the user with the **sp_adduser** command and skip step 7.

8. Exit the **isql** session.
9. Log in to SAP Replication Server and resume the connection to the standby database:

```
resume connection to wingak1505i.pubs2
go

Connection to 'wingak1505i.pubs2' is resumed
```

Note: Validate that the connection is not suspended or down. If the connection is down, check the SAP Replication Server log for errors and correct the errors, and then resume the connection.

10. Check the warm standby status:

```
admin logical_status
go
```

The output generated from **admin logical_status** is similar to:

Logical Connection Name	Active Connection Name	Active Conn State	Standby Connection Name	Standby Conn State
[102] pubs2a.pubs2s	[103] su-nak1505i.pubs2	Active/	[104] wing-ak1505i.pubs2	Active`

Controller RS	Operation in Progress	State of Op-eration in Progress	Spid
[16777317] PRS	None	None	

11. Release the secondary truncation point of the standby database server:

```
isql -Usa -P -Swingak1505i
use pubs2
go

dbcc settrunc ('ltm','ignore')
go
```

Switching the Active and Standby Databases

Switch from the active to the standby database if the active database will be unavailable for a long time.

In general, do not switch the active and standby databases when the active data server experiences a transient failure from which the SAP ASE recovers upon restarting with no need for additional recovery steps. You may want to switch if the active database will be unavailable for a long period of time.

1. Ensure that the RepAgent is shut down at the active database. If the RepAgent is still active, issue:

```
isql -Usa -P -Ssunak1505i
use pubs2
go
sp_stop_rep_agent pubs2
go
The Replication Agent thread for database 'pubs2' is
```

```
being stopped.
(return status = 0)
```

2. At the SAP Replication Server, enter:

```
isql -Usa -Psa_pass -SPRS
switch active for pubs2a.pubs2s to wingak1505i.pubs2
go
Switch active to wingak1505i.pubs2 for logical
connection to pubs2a.pubs2s is in progress
```

3. To monitor the progress of a switch, use:

```
admin logical_status
go
```

Logical Connection Name	Active Connection Name	Active Conn State	Standby Connection Name	Standby Conn State
[102] pubs2a.pubs2s	[104] wing-ak1505i.pubs2	Active/	[104] su-nak1505i.pubs2	Suspended/Wait-ing for Enable Marker

Controller RS	Operation in Progress	State of Op-eration in Progress	Spid
[16777317] PRS	None	None	

When the switch is complete, you can see the following in the SAP Replication Server log:

```
I. 2009/10/28 22:43:18. SQM starting: 102:1
pubs2a.pubs2s
I. 2009/10/28 22:43:18. Resetting Replication Agent
starting log position for wingak1505i.pubs2


I. 2009/10/28 22:43:19. DIST for 'pubs2a.pubs2s' is
Starting
I. 2009/10/28 22:43:19. Resuming LogTransfer for
wingak1505i.pubs2
I. 2009/10/28 22:43:19. Switch completed :
pubs2a.pubs2s
I. 2009/10/28 22:43:19. The DSI thread for database
'wingak1505i.pubs2' is started.
```

4. When the active database switch is complete, restart RepAgent for the new active database:

```
isql -Usa -P -Swingak1505i
sp_start_rep_agent pubs2
go
Replication Agent thread is started for database
'pubs2'.
(return status = 0)
```

A successful start-up writes messages to the SAP Replication Server log:

```
I. 2009/10/28 22:52:25. Replication Agent for
wingak1505i.pubs2 connected in passthru mode.
I. 2009/10/28 22:52:25. Setting system upgrade
locater for version 1100 to 00000000000000
00000000000000000000000000000000000000000000
000000000000000000 for
database wingak1505i.pubs2.
I. 2009/10/28 22:52:26. Distributor for
'pubs2a.pubs2s' received and proceessed enable
marker.
```

5. Resume the standby connection to replicate the data to the standby database, if the new standby database does not need to be resyncronized with the new active database, go to step 7:

```
resume connection to sunak1505i.pubs2
go
Connection to 'sunak1505i.pubs2' is resumed.
```

6. To verify whether the warm standby is operational or not, use:

```
admin logical_status
go
```

Logical Connection Name	Active Connection Name	Active Conn State	Standby Connection Name	Standby Conn State
[102] pubs2a.pubs2s	[104] wing-ak1505i.pubs2	Active/	[104] su-nak1505i.pubs2	Active/

Controller RS	Operation in Progress	State of Op-eration in Progress	Spid
[16777317] PRS	None	None	

7. If the old active database needs to be resynchronized with the new active database, first drop the standby connection from the SAP Replication Server:

```
drop connection to sunak1505i.pubs2
go
Connection to 'sunak1505i.pubs2' is dropped.
```

8. Rebuild the standby side of the warm standby connection using a dump of the new active database and synchronize with a dump marker by following the steps previously described in these topics:
   - Adding the standby database to the SAP Replication Server
   - Initializing the standby database

See also

- *Adding the Standby Database to the SAP Replication Server* on page 43

- *Initializing the Standby Database* on page 47

Replicating DDL in a Warm Standby Environment

Replicate DDL in a warm standby environment.

1. Log in to the active database using a user ID that exists on both the active and standby data server, and that has permission to create a table.

   Note: Do not use the maintenance user ID with the same password that you have defined when the active connection was set up.

2. Create a new table:

   ```
   % isql -Usa -P -Ssunak1505i
   use pubs2
   go
   create table t1 (a integer, b char(10), c datetime))
   go
   ```

3. Create unique indexes for better performance:

   ```
   create unique clustered index t1_idx1 on t1 (a,b)
   go
   ```

4. Log in to the standby database:

   ```
   % isql -Usa -P -Swingak1505i
   use pubs2
   go
   ```

5. Verify that the table and index exists in the standby database. If the table and index do not exist, follow the instructions in steps 6 and 7 otherwise, go to step 8.
6. Check the SAP Replication Server log file in: *$SYBASE*/REP-15_5/install/PRS.log. Correct the errors and resume the connection to the standby database:

   ```
   resume connection to rds.rdb
   go
   ```

   where:
   - *rds* – is the name of the data server that hosts the standby database.
   - *rdb* – is the name of the standby database.

   Note: If there are errors in the SAP Replication Server log, make sure that the user making the corrections is not the same maintenance user used to set up the active connection.

7. If you want the SAP Replication Server to skip any current transaction when trying to resume connection to the standby database, use:

   ```
   resume connection to rds.rdb
   skip transaction
   go
   ```

   See the***resume connection*** in the *Reference Manual* for other available options for the **resume connection** command.

8. Grant **insert**, **update**, and **delete** permission for the new table to the replicate database maintenance user at the standby database:

```
grant all on t1 to pubs2_maint
go
```

Replicating DML in a Warm Standby Environment

Replicate DML in a warm standby.

1. Log in to the active database using a user ID that exists on both the active and standby data server, and that has permission to create a table.
2. In the active database, insert a row to t1:

```
insert into t1 values (1,'first row',getdate())
go
```

3. Check whether the row exists in the standby database:

```
select * from t1
go
```

   If the row does not exist, follow instructions in steps 4 and 5 otherwise, go to step 6.

4. Check the SAP Replication Server log file at: *$SYBASE*/REP-15_2/install/PRS.log. Correct the errors and restart the connection to the replicate database:

```
resume connection to rds.rdb
go
```

5. If you want SAP Replication Server to skip any current transaction when trying to resume connection to the standby database, use:

```
resume connection to rds.rdb
skip transaction
go
```

   See ***resume connection*** in the *Reference Manual* for other available options for this command.

6. Log in to the active database and update the row:

```
update t1 set b = 'changed row' where a = 1
go
```

7. Log in to the standby database and verify whether the row exists:

```
select * from t1
go
```

8. Log in to the active database and enter:

```
truncate table t1
go
```

9. Log in to the standby database and enter:

```
select count (*) from t1
go
```

Tracing the SAP Replication Server Transactions to Target Databases

Trace transactions that SAP Replication Server sends to all replicate databases.

Because the trace is not specific to a connection and can get quite large, SAP recommends that you do not keep the trace function turned on.

1. Use **isql** to log in to SAP Replication Server.
2. Turn the trace on. Output from the trace function is written to the SAP Replication Server log.

   ```
   trace "on",dsi,dsi_buf_dump
   go
   ```

3. Use **isql** to insert one row in the active database:

   ```
   insert into t1 values (1,'first row',getdate())
   go
   (1 row affected)
   ```

4. Update the row in the active database:

   ```
   update t1 set c = getdate()
   go
   (1 row affected)
   ```

5. View the trace output in the SAP Replication Server log file, *$SYBASE*/REP-15_5/install/repservername.log:

   ```
   T. 2009/10/28 22:09:08. (138): Command(s) to
   'wingak1505i.pubs2':
   T. 2009/10/28 22:09:08.
   (138):'begin transaction
   [0a] update dbo.t1 set
   c='20091028 22:09:07:703'
   where a=1 and b='first row'
   and c='20091028
   22:05:53:843' '
   ```

Note: To turn off the tracing function in SAP Replication Server, use:

```
trace "off",dsi,dsi_buf_dump
go
```

Replication Definitions for Improving Performance

A replication definition describes the source table to SAP Replication Server, specifying the columns you want to copy. It may also describe attributes of the destination table. Destination tables that match the specified characteristics can subscribe to the replication definition.

SAP Replication Server does not require replication definitions to maintain a standby database, although using replication definitions can improve performance when you are replicating into a standby database. In addition, **create replication definitions** on tables that contain approximate numeric datatypes so that the where clause used to construct the SQL statement applied to the standby database does not include these columns. The range and storage precision of approximate numeric datatypes (real, float) is machine-dependent, and may result in the wrong rows being modified, or the correct row not being found.

When you specify that you want to use a replication definition for replicating into a standby database:

- SAP Replication Server optimizes updates and deletes by using the primary key defined in the replication definition to generate the **where** clause.
- You can specify whether SAP Replication Server uses the replication definition's replicate minimal columns setting for replicating into the standby database. This setting indicates whether updates replace the values for all columns or only the columns with changed values.

See *Manage Replicated Tables* in the *Admininstration Guide Volume 1* for detailed information on replication definitions.

Creating a Replication Definition

Create a replication definition to describe a replicated object. You can create replication definitions for databases, functions, or tables.

1. Use **isql** to log in to SAP Replication Server.
2. Create a replication definition for the warm standby for table "t1":

```
create replication definition t1_ws_repdef
with primary at pubs2a.pubs2s
with all tables named t1
(a integer, b char(10), c datetime)
primary key (a)
send standby replication definition columns
replicate SQLDML
go
Replication definition 't1_ws_repdef' is created.
```

Note: To replicate using SQL statement replication, which is available only in SAP ASE version 15.0.3 and later, use the replicate SQLDML clause.

3. Update the row in table “t1”:

```
update t1 set c = getdate()
go
```

4. Compare the trace output with the output that was generated before the replication definition was created:

```
T. 2009/10/28 22:10:43. (138): Command(s) to
'wingak1505i.pubs2':
T. 2009/10/28 22:10:43. (138): 'begin transaction [0a]
update dbo.t1 set a=1, b='first row', c='20091028
22:10:42:383' where a=1 '
```

   The **where** clause in the update statement, now contains only column “a” because the replication definition specified that column “a”, uniquely identifies the row.

See also

- *Tracing the SAP Replication Server Transactions to Target Databases* on page 54

Using SQL Statement Replication for Warm Standby

SQL statement replication complements log-based replication and addresses performance degradation caused by batch jobs.

Prerequisites

You can perform SQL statement replication only with SAP ASE 15.0.3 and later.

Task

1. In step 4 of the task, "Marking the active database for replication," the threshold for the database is set to 10. Therefore, SQL statement replication is used only when there are more than 10 rows. Insert 10 rows in to the table “t1”:

```
insert into t1 values (2,'first row',getdate())
insert into t1 values (3,'first row',getdate())
insert into t1 values (4,'first row',getdate())
insert into t1 values (5,'first row',getdate())
insert into t1 values (6,'first row',getdate())
insert into t1 values (7,'first row',getdate())
insert into t1 values (8,'first row',getdate())
insert into t1 values (9,'first row',getdate())
insert into t1 values (10,'first row',getdate())
insert into t1 values (11,'first row',getdate ())
```

2. Update a number of rows less than or equal to the threshold value.

```
update t1 set b = 'no SQL' where a < 3
go
(2 rows affected)
T. 2009/10/28 22:18:55. (138): Command(s) to
'wingak1505i.pubs2':
T. 2009/10/28 22:18:55. (138):
```

```
'begin transaction [0a]
update dbo.t1 set a=1, b='no SQL',
c='20091028 22:10:42:383'
where a=1 [0a] update dbo.t1
set a=2, b='no SQL', c='20091028
22:12:24:093' where a=2 '
```

The trace statements show the individual SQL updates to each row.

3. Update all the rows:

```
update t1 set b = 'yes SQL'
go
(11rows affected)
```

The trace output shows the SQL statement, not the individual statements for each row.

```
T. 2009/10/28 22:23:35. (138): Command(s) to
'wingak1505i.pubs2':
T. 2009/10/28 22:23:35. (138): 'begin transaction
[0a] update dbo.t1 set b = 'yes SQL' '
```

Note: To turn off the tracing function in SAP Replication Server, use:

```
trace "off",dsi,dsi_buf_dump
go
```

See also

- *Marking the Active Database for Replication* on page 42

Materialization and Resynchronization

Materialization is copying data specified by a subscription from a primary or source database or table, to a replicate or target database or table. Resynchronization makes all of the data in the primary database or table and the replicate database or table identical. You can resynchronize a database, a table, or individual rows. The materialization process causes resynchronization.

The materialization method depends upon the amount of data to be transmitted, the portion of the data that resides at the replicate site, the time available for the process, and whether the primary database or table can have any activity against it.

The provided scenarios make several assumptions and considerations:

- Marking tables with text or image columns can take a long time. In SAP Adaptive Server Enterprise (SAP ASE)15.0 ESD #1 and later, the **use_index** option can significantly speed up marking these columns. When text in a table is marked for replication using one method, those text pages are not remarked if a different method is added. For example, if a database is marked for replication, marking a table containing text columns in that database does not require all the text pages to be re-marked for replication.
- The database materialization and synchronization options discussed in the examples use the SAP ASE database dump and load method. See the SAP ASE manuals for any requirements for using these routines.
- When you dump a database, the secondary truncation point and the RepAgent for the database are included in the dump. Users and permissions are copied; but logins and roles are not.
- Adding and dropping a subscription requires the Data Server Interface (DSI) to be active or awaiting command to the replicate database.
- These scenarios assume that the connections from the SAP Replication Server to the database exists and that replication definitions are already defined.

Note: The scenarios provided do not work with synchronous replication.

See *Manage Subscriptions* in the *Administration Guide Volume 1* for complete descriptions of various materializations methods.

Scenario 1

The primary and the replicate database can be refreshed from a different database source while existing replication definitions and subscriptions continue to be used.

If the source database has never participated in replication, temporarily add the database to an SAP Replication Server so that it has all the tables and stored procedures needed for replication, before making a copy of it.

This scenario uses a third database (for example, a production database) to populate the source and target database environment (for example, a test database). You would use this scenario when you want to refresh a test system from a copy of a production database.

Before You Begin

In the current primary, validate that the objects that have been marked for replication. If the database has `text` and `image` columns and these columns will be replicated, decide whether to mark the database for replication prior to the copy or after the primary database has been established.

1. Check whether the database has been marked for replication:

```
use pri
go
sp_reptostandby pri
go
The replication status for database 'pri' is 'ALL'.
The replication mode for database 'pri' is 'off'.
(return status = 0)
```

Note: In this scenario, the database has been marked for database replication and has not been marked for SQL statement replication.

2. If the database has not been marked for replication, check to see whether the tables and stored procedures are marked for replication:

   For tables, execute:

```
use pri
go
sp_setreptable
go
Name       Type
------- ----------------
t1         user table
t2         user table
(2 rows affected)
(return status = 0)
```

   For stored procedures, execute:

```
use pri
go
sp_setrepproc
go
Name       Type
------- ----------------
rs_marker  stored procedure
(1 rows affected)
(return status = 0)
```

Note: Any table or stored procedure beginning with "rs_" is created by SAP Replication Server. Therefore, in this scenario, t1 and t2 are application tables, which were marked for

replication and **rs_marker** is a stored procedure created by adding the database to the replication system.

3. Get the current generation number of the primary database.

```
use pri
go
dbcc gettrunc
go
secondary trunc page secondary trunc state dbrepstat
------------------- ---------------------- -------
2669                              1            173
generation id database id database name ltl version
------------- ------------ ------------- ---------
0                 7            pri           720
```

4. At the RSSD, obtain the maintenance users for the connections:

```
use PRS2_RSSD
go
rs_helpuser
go

User Name      Permission(s) Name
------------ ------------------------------
PRS2_id_user    no grants
sa              sa
PRS2_ra         connect source
PRS2_rsi        connect source
PRS2_RSSD_prim  source, primary subscr

Maintenance Users
User name       Destination DS.DB
------------ -------------------------------------
PRS2_RSSD_maint  sunak1505x.PRS2_RSSD
pri_maint        sunak1505x.pri
pri_maint        sunak1505x.rep
(return status = 0)
```

In this example, “pri_maint” is the maintenance user for both the connection to the primary database, sunak1505x.pri. and replicate database, sunak1505x.rep.

Begin the Resynchronization

You can stop all user activity to the source database; however, it is not required.

When a new database is brought online, SAP ASE creates a database that is transactionally consistent at the time of the dump.

1. Stop all user activity of the primary database including the existing RepAgent:

```
sp_stop_rep_agent pri
go
The Replication Agent thread for database 'pri' is
being stopped.
(return status = 0)
```

2. Suspend the connection to the primary and replicate database:

```
isql –Usa –Psa_pass –SPRS2
suspend connection to sunak1505x.pri
```

```
go
Connection to 'sunak1505x.pri' is suspended
suspend connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is suspended.
```

3. Dump the database of the source database:

```
dump database diffprim to
'/c11014900/sybase1520x/diffprim.dmp'
go
```

4. Load the source database dump to the primary database:

```
load database pri from
'/c11014900/sybase1520x/diffprim.dmp'
go
```

5. Bring the database online:

```
online database pri
go
```

6. Add the maintenance user of the primary connection to the primary database and grant the appropriate privileges. If the maintenance user already exists on this server, synchronize the "suid" of the maintenance user and any other user that will be logging in to the primary database.

```
use pri
go
sp_adduser pri_maint
go
New user added.
(return status = 0)
grant execute on rs_get_lastcommit to pri_maint
go
grant all on rs_lastcommit to pri_maint
go
```

Note: If the maintenance user is new to this server, add the maintenance user login to this server, and grant the replication role:

```
sp_addlogin <maintenance user>, <maintenance userpassword>,
<database>
go
sp_role "grant", replication_role, <maintenance
user>
go
```

If the password to the maintenance user is unknown, set a new password in SAP ASE. In the SAP Replication Server, change the password of the maintenance user to match with what was assigned in the SAP ASE:

```
alter connection to <dataserver>.<database>
set password to <new password>
go
```

---

7. If the source database dump came from a primary database that had a Replication Agent, remove the secondary truncation point and the existing Replication Agent from the new primary database:

```
use pri
go
dbcc settrunc (ltm,ignore)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- -------------------- --------
2667                        0                     166
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                7               pri         720
```

Note: A secondary truncation state of 0 indicates that the secondary truncation point is inactive.

```
sp_config_rep_agent pri,'disable'
go
Replication Agent disabled for database 'pri'. The
secondary truncation point
in the database is no longer active.
(return status = 0)
```

8. Enable the Replication Agent at the primary database:

```
sp_config_rep_agent pri,
'enable','PRS2','PRS2_ra','PRS2_ra_ps'
go
Replication Agent enabled for database 'pri'. The
Replication Agent thread needs
to be started using sp_start_rep_agent.
(return status = 0)
```

Note: Any SAP Replication Server user with “connect source” privilege can be used. If the password of the SAP Replication Server is unknown, it can be reset using the alter user SAP Replication Server command.

9. Modify the settings in the RSSD and the primary database to have the Replication Agent start at the end of the transaction log:

```
use PRS2_RSSD
go
rs_zeroltm sunak1505x,pri
go
Locater has been reset to zero.
(return status = 0)
isql –Usa –P –Ssunak1505x
use pri
go
dbcc settrunc (ltm,valid)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- -------------------- --------
2670                          1                 167
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                7          pri            720
```

10. Increase the generation number of the new primary database by 1 in case the log pages of the new database are numerically less than the log pages of the previous copy of the database:

```
dbcc settrunc (ltm,gen_id,1)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- --------------------- ---------
2670                            1               167
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                7          pri             720
```

Note: The previous value of the generation id was obtained in step 3 of "Before you begin" section . If `text` or `image` columns need to be marked for replication, mark the tables and columns for replication here.

11. Load the replicate database with the source database dump:

```
load database rep from
'/c11014900/sybase1520x/diffprim.dmp'
go
```

12. Bring the replicate database online:

```
online database rep
go
```

13. Add the maintenance user of the replicate connection, to the replicate database and grant all appropriate privileges using the information gathered prior to the resynch.

```
use rep
go
sp_adduser pri_maint
go
New user added.
(return status = 0)
grant execute on rs_get_lastcommit to pri_maint
go
grant all on rs_lastcommit to pri_maint
go
```

Note: If the maintenance user is new to this server, add the maintenance user login to this server, and grant the replication role:

```
sp_addlogin <maintenance user>, <maintenance userpassword>,
<database>
go
sp_role "grant", replication_role, <maintenance
user>
go
```

If the password to the maintenance user is unknown, set a new password in SAP ASE. In the SAP Replication Server, change the password of the maintenance user to match with what was assigned in the SAP ASE:

```
alter connection to <dataserver>.<database>
set password to <new password>
go
```

14. If the source database had a RepAgent, release the secondary truncation point and remove the RepAgent:

```
use rep
go
dbcc settrunc (ltm,ignore)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- --------------------- --------
2667                        0                   166
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                7             pri          720
```

Note: A secondary truncation state of 0 indicates that the secondary truncation point is inactive.

```
sp_config_rep_agent rep,'disable'
go
Replication Agent disabled for database 'rep'. The
secondary truncation point
in the database is no longer active.
(return status = 0)
```

15. Truncate the rs_lastcommit table at the replicate database. Enter:

```
truncate table rs_lastcommit
go
```

16. Grant appropriate permissions to the replicate tables so that the maintenance user can apply commands to the replicate database:

```
grant all on t1 to pri_maint
go
grant all on t2 to pri_maint
go
```

17. Purge the inbound queue of the primary connection and the outbound queue of the replicate connection on the SAP Replication Server to eliminate any in-process work from the old replication setup:

```
isql -Usa -Psa_pass -SPRS2
sysadmin hibernate_on
go
The Replication Server has now entered hibernation
mode.
sysadmin sqm_purge_queue, 104,1
go
sysadmin sqm_purge_queue, 106,0
go
sysadmin hibernate_off
go
The Replication Server has now finished hibernation
mode.
```

18. Resume connection to the primary and the replicate database:

```
resume connection to sunak1505x.pri
go
Connection to 'sunak1505x.pri' is resumed.
resume connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is resumed.
```

19. Start the Replication Agent of the primary database:

```
sp_start_rep_agent pri
go
Replication Agent thread is started for database
'pri'.
(return status = 0)
```

20. On the SAP Replication Server, ensure that both the RepAgent and the DSI threads were successfully started

```
admin who
go
Spid     Name          State               Info
----     --------      ------------        ------------------
139      DSI EXEC      Awaiting Command    104(1) sunak1505x.pri
135      DSI           Awaiting Message    104 sunak1505x.pri
129      DIST          Awaiting Wakeup     104 sunak1505x.pri
130      SQT          Awaiting Wakeup     104:1 DIST sunak1505x.pri
123      SQM           Awaiting Message    104:1 sunak1505x.pri
56       SQM           Awaiting Message    104:0 sunak1505x.pri
142      REP AGENT     Awaiting Command    sunak1505x.pri
143      NRM           Awaiting Command    sunak1505x.pri
145      DSI EXEC      Awaiting Command    106(1) sunak1505x.rep
144      DSI           Awaiting Message    106 sunak1505x.rep
124      SQM           Awaiting Message    106:0 sunak1505x.rep
```

21. Validate that replication is available and working.
22. Allow users on the primary database.

Scenario 2

You can materialize a replicate database from a primary database where all user activity has stopped at the primary database.

Use this materialization method when replication is broken between the primary database and the replicate database for a significant period of time and the queues are filling up, which needs to be purged. Refresh the replicate database with a copy from the current primary database. This can be used with either table replication or database replication. This scenario assumes that user activity is stopped at the primary while the database dump is taken.

Before You Begin

- At the RSSD, obtain the maintenance users for the connections:

```
use PRS2_RSSD
go
```

```
rs_helpuser
go

User Name      Permission(s) Name
------------ ------------------------------
PRS2_id_user    no grants
sa              sa
PRS2_ra         connect source
PRS2_rsi        connect source
PRS2_RSSD_prim  source, primary subscr

Maintenance Users
User name       Destination DS.DB
------------ -------------------------------------
PRS2_RSSD_maint  sunak1505x.PRS2_RSSD
pri_maint        sunak1505x.pri
pri_maint        sunak1505x.rep
(return status = 0)
```

In this example, “pri_maint” is the maintenance user for both the connection to the primary database, sunak1505x.pri. and replicate database, sunak1505x.rep.

Begin the Resynchronization

1. If it is not already suspended, suspend the connection to the replicate database:

   ```
   isql –Usa –P<sa_pass> –SPRS2
   suspend connection to sunak1505x.rep
   go
   Connection to 'sunak1505x.rep' is suspended.
   ```

2. In the primary database, release the secondary truncation point—any log records not previously replicated are already reflected in the data.

   ```
   isql –Usa –P –SSunak1505x
   use pri
   go
   sp_stop_rep_agent pri
   go
   The Replication Agent thread for database 'pri' is
   being stopped.
   (return status = 0)
   dbcc settrunc (ltm,ignore)
   go
   secondary trunc page secondary trunc state dbrepstat
   -------------------- --------------------- -------
   2669                    0                     172
   generation id database id database name ltl version
   ------------- ------------ ------------- ----------
   0                7               pri         720
   ```

   Note: When secondary truncation state is 0, the secondary truncation point is inactive on the primary database.

3. Dump the primary database:

```
dump database pri to '/c11014900/sybase1520x/ASE-
15_0/bin/pri.dmp'
go
```

4. Load the replicate database:

```
load database rep from '/c11014900/sybase1520x/ASE-
15_0/bin/pri.dmp'
go
```

5. Purge queues of any existing work from the primary database (inbound queue) to the replicate database (outbound queue):

```
isql -Usa -Psa_pass -SPRS2
sysadmin hibernate_on
go
The Replication Server has now entered hibernation
mode.
sysadmin sqm_purge_queue, 104,1
go
sysadmin sqm_purge_queue, 106,0
go
sysadmin hibernate_off
go
The Replication Server has now finished hibernation
mode.
```

6. Bring the replicate database online:

```
online database rep
go
```

7. To the replicate database, add the maintenance user of the replicate connection, and grant the appropriate privileges.

```
use rep
go
sp_adduser pri_maint
go
New user added.
(return status = 0)
grant execute on rs_get_lastcommit to pri_maint
go
grant all on rs_lastcommit to pri_maint
go
```

Note: If the maintenance user is new to this server, add the maintenance user login to this server, and grant the replication role:

```
sp_addlogin <maintenance user>, <maintenance userpassword>,
<database>
go
sp_role "grant", replication_role, <maintenance
user>
go
```

If the password to the maintenance user is unknown, set a new password in Adaptive Server. In the Replication Server, change the password of the maintenance user to match with what was assigned in the Adaptive Server:

```
alter connection to <dataserver>.<database>
set password to <new password>
go
```

8. Resume the connection to the replicate database:

```
resume connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is resumed.
```

9. On the Replication Server, confirm that the connection to the replicate database has successfully resumed:

```
admin who
go
Spid    Name        State               Info
-----   -------     ------------        -------------------
103     DSI EXEC    Awaiting Command    104(1) sunak1505x.pri
100     DSI         Awaiting Message    104 sunak1505x.pri
95      DIST        Awaiting Wakeup     104 sunak1505x.pri
96      SQT         Awaiting Wakeup     104:1 DISTsunak1505x.pri
92      SQM         Awaiting Message    104:1 sunak1505x.pri
56      SQM         Awaiting Message    104:0 sunak1505x.pri
        REP AGENT   Down                sunak1505x.pri
        NRM         Down                sunak1505x.pri
107     DSI EXEC    Awaiting Command    106(1)sunak1505x.rep
106     DSI         Awaiting Message    106 sunak1505x.rep
91      SQM         Awaiting Message    106:0 sunak1505x.rep
```

Note: If the connection does not successfully resume, look at the Replication Server log to see what errors occurred, correct the errors, and resume the connection.

10. Add the replication definition and subscription to the Replication Server, if needed. In this example, we are using a database replication definition and a database subscription. Enter:

```
create database replication definition pri_db_repdef
with primary at sunak1505x.pri
replicate DDL
go
Database replication definition 'pri_db_repdef'
for sunak1505x.pri is created.
create subscription rep_db_repdef
for database replication definition pri_db_repdef
with primary at sunak1505x.pri
with replicate at sunak1505x.rep
without materialization
subscribe to truncate table
go

Subscription 'rep_db_repdef' is in the process of
being created
```

11. Reset the start point on the primary database to the end of the primary database log.
    a. Zero the `rs_locater` entry for this connection in the RSSD:

```
use PRS2_RSSD
go
rs_zeroltm sunak1505x,pri
go
```

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

b. Enable the secondary truncation point in the primary database:

```
isql -Usa -P -SSunak1505x
use pri
go
dbcc settrunc (ltm,valid)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- -------------------- ---------
2669                              1               173
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                 7                 pri     720
```

12. Start the Replication Agent on the primary database:

```
use pri
go
sp_start_rep_agent pri
go
Replication Agent thread is started for database
'pri'.
(return status = 0)
```

13. Validate that the Replication Server connections are ready to replicate:

```
admin who
go
Spid     Name      State               Info
----     --------  ------------        -------------
103      DSI EXEC  Awaiting Command     104(1) sunak1505x.pri
100      DSI       Awaiting Message     104 sunak1505x.pri
95       DIST      Awaiting Wakeup      104 sunak1505x.pri
96       SQT       Awaiting Wakeup      104:1 DIST sunak1505x.pri
92       SQM       Awaiting Message     104:1 sunak1505x.pri
56       SQM       Awaiting Message     104:0 sunak1505x.pri
108      REP AGENT Awaiting Command     sunak1505x.pri
109      NRM       Awaiting Command     sunak1505x.pri
107      DSI EXEC  Awaiting Command     106(1)sunak1505x.rep
106      DSI       Awaiting Message     106 sunak1505x.rep
91       SQM       Awaiting Message     106:0 sunak1505x.rep
```

14. Validate that data can be replicated from the primary to the replicate.:

Note: If Replication Agent goes down, look in the ASE error log for messages. If the DSI is suspended, look in the Replication Server error log or the ASE error log for messages. To start the Replication Agent on the primary database, use **sp_start_rep_agent** <*dbname*> command. To resume the DSI at the Replication Server, use the resume connection to <*dataserver*>.<*database*> command. To move past and not apply transactions in the outbound queue, use the skip transaction clause of the resume connection command.

15. Allow users on the primary database.

Scenario 3

You can materialize a replicate database from a primary database where user activity cannot be stopped and using MSA replication.

It is assumed that activity cannot be stopped at the primary database while the database dump is in progress. This scenario uses the primary database to populate the replicate database and uses database replication definition and subscription.

If the primary database is replicating to multiple replicate databases, the complete process including defining the subscription, dumping the primary database, and loading the replicate database must be completed for each replicate database, defining the subscription for the next replicate database.

Before You Begin

- At the RSSD, obtain the maintenance users for the connections:

```
use PRS2_RSSD
go
rs_helpuser
go
User Name      Permission(s) Name
------------ ------------------------------
PRS2_id_user    no grants
sa              sa
PRS2_ra         connect source
PRS2_rsi        connect source
PRS2_RSSD_prim  source, primary subscr

User name       Destination DS.DB
------------ -------------------------------------
PRS2_RSSD_maint  sunak1505x.PRS2_RSSD
pri_maint        sunak1505x.pri
pri_maint        sunak1505x.rep
(return status = 0)
```

In this example, “pri_maint” is the maintenance user for both the connection to the primary database, sunak1505x.pri. and replicate database, sunak1505x.rep.

Begin the Resynchronization

1. If the database subscription exists, drop the subscription:

```
isql -Usa -P -Ssunak1505x
use PRS2_RSSD
go
rs_helpdbsub
go
DBSub.Name ReplicateDS.DB ReplicateRS Status at RRS DBRep.Def.Name
---------- -------------- ----------- ------------  -----------
rep_db_sub sunak1505x       PRS2        Valid           pri_db_rep
```

```
PrimaryDS.DB PrimaryRS Status at PRS Method Trunc.Table Creation
Date
------------ --------- -----------  ---------- -------------
sunak1505x      PRS2    Valid Use  Dump Marker  Yes Dec  9 2009
3:38PM
```

Note: The connection to the replicate database must be "Awaiting Command" prior to dropping the subscription.

```
isql –Usa –P<i>sa_pass</i> –SPRS2
drop subscription rep_db_sub
for database replication definition pri_db_rep
with primary at sunak1505x.pri
with replicate at sunak1505x.rep
without purge
go
Subscription 'rep_db_sub' is in the process of being
dropped.
```

2. Validate that the primary database is marked, enter:

```
isql -Usa -P -Ssunak1505x
use pri
go
sp_reptostandby pri
go
The replication status for database 'pri' is 'ALL'.
The replication mode for database 'pri' is 'off'.
(return status = 0)
```

Note: In this scenario, the database "pri" is marked to replicate both DML and DDL and is not marked for SQL statement replication.

3. Validate that the RepAgent on the primary database is running and the connection to the replicate database exists and is not suspended:

```
isql -Usa -P<i>sa_pass</i> -SPRS2
admin who
go
Spid Name       State              Info
---- ---------- --------------    ------------------
62   DSI EXEC    Awaiting Command 104(1) sunak1505x.pri
57   DSI         Awaiting Message 104 sunak1505x.pri
59   DIST        Awaiting Wakeup  104 sunak1505x.pri
60   SQT         Awaiting Wakeup  104:1 DIST sunak1505x.pri
58   SQM         Awaiting Message 104:1 sunak1505x.pri
56   SQM         Awaiting Message 104:0 sunak1505x.pri
61   REP AGENT   Awaiting Command       sunak1505x.pri
63   NRM         Awaiting Command       sunak1505x.pri
68   DSI EXEC    Awaiting Command 105(1) sunak1505x.rep
67   DSI         Awaiting Message 105 sunak1505x.rep
66   SQM         Awaiting Message 105:0 sunak1505x.rep
```

4. Define the subscription with the option that we are going to use a dump to syncronize the data:

```
define subscription rep_db_sub
for database replication definition pri_db_rep
with primary at sunak1505x.pri
with replicate at sunak1505x.rep
subscribe to truncate table
use dump marker
go
Subscription 'rep_db_sub' is in the process of being
defined.
```

The connection to the replicate database is still not suspended.

```
admin who
go
Spid Name      State                Info
---- ----      ---------------- -------------------
68   DSI EXEC Awaiting Command 105(1) sunak1505x.rep
67   DSI      Awaiting Message 105 sunak1505x.rep
66   SQM      Awaiting Message 105:0 sunak1505x.rep
```

The above **admin who** command displays only the connections that are affected and not the complete list of connections.

5. Dump the primary database:

```
dump database pri to 'pri.dmp'
go
```

The connection to the replicate database is now suspended.

```
admin who
go
Spid Name     State         Info
---- ---- ---------------- -------------------
68   DSI EXEC Suspended        105(1) sunak1505x.rep
67   DSI      Suspended        105 sunak1505x.rep
66   SQM      Awaiting Message 105:0 sunak1505x.rep
```

6. Load the replicate database:

```
load database rep from '/c11014900/sybase1520x/ASE-
15_0/bin/pri.dmp'
go
```

7. Bring the replicate database online:

```
online database rep
go
```

8. Add the maintenance user of the replicate connection to the replicate database and grant the appropriate privileges before the resynchronization.

```
use rep
go
sp_adduser pri_maint
go
New user added.
(return status = 0)
grant execute on rs_get_lastcommit to pri_maint
go
grant all on rs_lastcommit to pri_maint
go
```

Note: If the maintenance user is new to this server, add the maintenance user login to this server, and grant the replication role:

```
sp_addlogin <maintenance user>, <maintenance userpassword>,
<database>
go
sp_role "grant", replication_role, <maintenance
user>
go
```

If the password to the maintenance user is unknown, set a new password in Adaptive Server. In the Replication Server, change the password of the maintenance user to match with what was assigned in the Adaptive Server:

```
alter connection to <dataserver>.<database>
set password to <new password>
go
```

9. Truncate the `rs_lastcommit` table at the replicate database:

```
use rep
go
truncate table rs_lastcommit
go
```

10. Release the secondary truncation point from the replicate database:

```
dbcc settrunc (ltm,ignore)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- ---------------------- -----
2668                          0                  172
generation id database id database name ltl version
------------- ------------ ------------- -------
0                  8               rep         720
```

11. Resume connection to the replicate database. If the DSI stays suspended, look at the Replication Server logs for any errors encountered while resuming the DSI:

```
resume connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is resumed
```

Scenario 4

You can materialize a primary database from a replicate database where user activity can be stopped and the replicate database is a duplicate of the primary database.

Before You Begin

Use this materialization method if a problem ocuurs to your primary database and the replicate database, which is the duplicate of your primary database, is old data. In the current primary database, validate the objects that have been marked for replication. If the database has text and image columns and these columns will be replicated, decide whether to mark the database

for replication prior to the copy or after replication from the primary database has been established.

1. Check whether the database has been marked for replication:

```
use pri
go
sp_reptostandby pri
go
The replication status for database 'pri' is 'ALL'.
The replication mode for database 'pri' is 'off'.
(return status = 0)
```

Note: In this scenario, the database has been marked for database replication and has not been marked for SQL statement replication.

2. If the database has not been marked for replication, check to see whether the tables and stored procedures are marked for replication:

   For tables, execute:

```
use pri
go
sp_setreptable
go
Name       Type
------- ----------------
t1         user table
t2         user table
(2 rows affected)
(return status = 0)
```

   For stored procedures, execute:

```
use pri
go
sp_setrepproc
go
Name       Type
------- ----------------
rs_marker  stored procedure
(1 rows affected)
(return status = 0)
```

Note: Any table or stored procedure beginning with "rs_" is created by Replication Server. Therefore, in this scenario, t1 and t2 are application tables, which were marked for replication and **rs_marker** is a stored procedure created by adding the database to the replication system.

3. Get the current generation number of the primary database.

```
use pri
go
dbcc gettrunc
go
secondary trunc page secondary trunc state dbrepstat
-------------------- -------------------- -------
2669                            1            173
```

```
generation id database id database name ltl version
------------- ------------ ------------- ---------
0               7            pri           720
```

4. At the RSSD, obtain the maintenance users for the connections:

```
use PRS2_RSSD
go
rs_helpuser
go

User Name      Permission(s) Name
------------ ------------------------------
PRS2_id_user    no grants
sa              sa
PRS2_ra         connect source
PRS2_rsi        connect source
PRS2_RSSD_prim  source, primary subscr

Maintenance Users
User name       Destination DS.DB
------------ -------------------------------------
PRS2_RSSD_maint  sunak1505x.PRS2_RSSD
pri_maint        sunak1505x.pri
pri_maint        sunak1505x.rep
(return status = 0)
```

In this example, “pri_maint” is the maintenance user for both the connection to the primary database, sunak1505x.pri. and replicate database, sunak1505x.rep.

Begin the Resynchronization

You need not stop user activity at the source database; because when the new database goes online, Adaptive Server creates a database that is transactionally consistent at the time of the dump.

1. Stop all user activity of the replicate database including the existing DSI connection:

```
isql –Usa –Psa_pass –SPRS2
suspend connection to sunak1505.rep
go
Connection to 'sunak1505x.rep' is suspended.
```

2. Stop all user activity of the primary database including the existing Replication Agent and DSI:

```
sp_stop_rep_agent pri
go
The Replication Agent thread for database 'pri' is
being stopped.
(return status = 0)
isql -Usa -Psa_pass –SPRS2
suspend connection to sunak1505x.pri
go
Connection to 'sunak1505x.pri' is suspended.
```

3. Validate that the Replication Server is not actively using either the primary or the replicate database connection:

```
admin who
go
```

```
Spid    Name     State                Info
----  ---------- -----------------    ------------
     DSI EXEC   Suspended            104(1)sunak1505x.pri
     DSI        Suspended            104 sunak1505x.pri
129  DIST       Awaiting Wakeup      104 sunak1505x.pri
130  SQT        Awaiting Wakeup      104:1 DIST sunak1505x.pri
123  SQM        Awaiting Message     104:1 sunak1505x.pri
56   SQM        Awaiting Message     104:0 sunak1505x.pri
     REP AGENT  Down                 sunak1505x.pri
     NRM        Down                 sunak1505x.pri
     DSI EXEC   Suspended            106(1) sunak1505x.rep
     DSI        Suspended            106 sunak1505x.rep
124  SQM       Awaiting Message      106:0 sunak1505x.rep
```

4. Dump the replicate database:

```
dump database diffprim to
'/c11014900/sybase1520x/rep.dmp'
go
```

5. Load the replicate database dump to the primary database:

```
load database pri from
'/c11014900/sybase1520x/rep.dmp'
go
```

6. Bring the primary database online:

```
online database pri
go
```

7. Add the maintenance user of the primary connection to the primary database and grant the appropriate privileges. If the maintenance user already exists on this server, synchronize the “suid” of the maintenance user and any other user that will be logging in to the primary database.

```
use pri
go
sp_adduser pri_maint
go
New user added.
(return status = 0)
grant execute on rs_get_lastcommit to pri_maint
go
grant all on rs_lastcommit to pri_maint
go
```

Note: If the maintenance user is new to this server, add the maintenance user login to this server, and grant the replication role:

```
sp_addlogin <maintenance user>, <maintenance userpassword>,
<database>
go
sp_role "grant", replication_role, <maintenance
user>
go
```

If the password to the maintenance user is unknown, set a new password in Adaptive Server. In the Replication Server, change the password of the maintenance user to match with what was assigned in the Adaptive Server:

```
alter connection to <dataserver>.<database>
set password to <new password>
go
```

8. If the replicate database dump had a Replication Agent, remove the secondary truncation point and the existing RepAgent from the new primary database:

```
use pri
go
dbcc settrunc (ltm,ignore)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- --------------------- --------
2667                        0                    166
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                 7               pri          720
```

Note: A secondary truncation state of 0 indicates that the secondary truncation point is inactive.

```
sp_config_rep_agent pri,'disable'
go
Replication Agent disabled for database 'pri'. The
secondary truncation point
in the database is no longer active.
(return status = 0)
```

9. Enable the RepAgent at the primary database:

```
sp_config_rep_agent pri,
'enable','PRS2','PRS2_ra','PRS2_ra_ps'
go
Replication Agent enabled for database 'pri'. The
Replication Agent thread
needs to be started using sp_start_rep_agent.
(return status = 0)
```

Note: Any Replication Server user with **connect source** privilege can be used. If the password of the Replication Server is unknown, it can be reset using the alter user Replication Server command.

10. Ensure that either the database or the tables on the updated primary database are marked for replication.
11. Modify the settings in the RSSD and the primary database to have the Replication Agent start at the end of the transaction log:

```
use PRS2_RSSD
go
rs_zeroltm sunak1505x,pri
go
Locater has been reset to zero.
(return status = 0)
isql –Usa –P –Ssunak1505x
use pri
go
dbcc settrunc (ltm,valid)
```

```
go
secondary trunc page secondary trunc state dbrepstat
-------------------- ---------------------- --------
2670                            1               167
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                 7          pri             720
```

12. Increase the generation number by 1 in case the log pages of the new database are numerically less than the log pages of the previous copy of the database:

```
dbcc settrunc (ltm,gen_id,1)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- ---------------------- --------
2670                            1               167
generation id database id database name ltl version
------------- ------------ ------------- ----------
0                 7          pri             720
```

Note: The generation number was obtained prior to beginning the resynchronization.

13. Truncate the rs_lastcommit table at the replicate database:

```
use rep
go
truncate table rs_lastcommit
go
```

14. Purge queues of any existing work from primary to replicate:

```
isql -Usa -Psa_pass -SPRS2
sysadmin hibernate_on
go
The Replication Server has now entered hibernation mode.
```

Note: The **sysadmin sqm_purge_queue** command requires the Replication Server to be in hibernate or standalone mode. When the Replication Server is in hibernate or standalone mode, no work is performed.

```
sysadmin sqm_purge_queue, 104,1
go
sysadmin sqm_purge_queue, 106,0
go
```

Note: The queue number used in the **sysadmin sqm_purge_queue** command, is the connection number for the connection, as shown in the **admin who** command. The queue type is either 1 for inbound or 0 for outbound and is available from the admin who command.

```
sysadmin hibernate_off
go
The Replication Server has now finished hibernation
mode.
```

15. Turn of the hibernation mode after the queues are purged:

```
isql -Usa -Psa_pass -SPRS2
sysadmin hibernate_off
go
```

16. Resume the connection to the primary and replicate database:

```
resume connection to sunak1505x.pri
go
Connection to 'sunak1505x.pri' is resumed.
resume connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is resumed
```

17. On the primary database, start the RepAgent:

```
sp_start_rep_agent pri
go
Replication Agent thread is started for database
'pri'.
(return status = 0)
```

18. In the Replication Server, ensure that both the RepAgent and the DSI threads have been successfully started:

```
admin who
go
Spid Name     State             Info
---- -------- ------------     ------------------
139  DSI EXEC Awaiting Command 104(1) sunak1505x.pri
135  DSI      Awaiting Message 104 sunak1505x.pri
129  DIST     Awaiting Wakeup  104 sunak1505x.pri
130  SQT      Awaiting Wakeup  104:1 DIST sunak1505x.pri
123 SQM       Awaiting Message 104:1 sunak1505x.pri
56  SQM       Awaiting Message 104:0 sunak1505x.pri
143 REP AGENT Awaiting Command sunak1505x.pri
145 NRM       Awaiting Command sunak1505x.pri
```

19. Validate that replication is available and working.
20. Allow users on the primary database.

Scenario 5

You can materialize a table to the replicate database where the Replication Server is doing the materialization. In this scenario there are either no users using the table or the primary table can be locked while the materialization is taking place and there are few rows in the table. This scenario lets Replication Server populate the replicate table.

Ensure that the password of the “sa” user is the same for the primary ASE and the Replication Servers. You cannot use this option if these passwords are not identical.

Before You Begin

The maintenance user for the replicate database must have access to insert data into the replicate table. The table has already been marked using **sp_setreptable**.

Begin the Resynchronization

1. Create the replication definition:

```
create replication definition t1_repdef
with primary at sunak1505x.pri
with all tables named t1
(a char(10),
b char(10))
primary key (a)
go
Replication definition 't1_repdef' is created
```

2. Create the subscription:

```
create subscription t1_sub
for t1_repdef
with replicate at sunak1505x.rep
incrementally
subscribe to truncate table
go
Subscription 't1_sub' is in the process of being
created.
```

Note: This option places a hold on the table at the primary database while the selection of the rows for materialization is taking place.

3. Check if the subscription is valid at the primary and replicate site using the **check subscription** command:

```
check subscription t1_sub
for t1_repdef
with replicate at sunak1505x.rep
go
Subscription t1_sub is VALID at the replicate.
Subscription t1_sub is VALID at the primary.
```

4. If the subscription is not valid, check the Replication Server to see if a materialization queue has been created.

```
admin who
go
Spid   Name    State         Info
---- ------ ------------------- -------------
2      SQM     Awaiting Message 106:-2147483541
t1_repdef.t1_sub
```

If the subscription has not materialized within an acceptable time, and the materialization queue still exists, look in the Replication Server log for any error messages, correct the error, drop the materialization queue (**sysadmin drop_queue**), drop the subscription, and re-create the subscription.

Scenario 6

You can materialize a table where activity to the primary table is going on while the materialization takes place. Use this option when there is a more than a small amount of data to synchronize between the primary and replicate table. This scenario assumes that the table has already been marked using **sp_setreptable**.

Begin the Resynchronization

1. Create the replication definition:

```
create replication definition t1_repdef
with primary at sunak1505x.pri
with all tables named t1
(a char(10),
b char(10))
primary key (a,b)
go
Replication definition 't1_repdef' is created
```

   If the replication definition was previously created with **replicate minimal columns**, alter the replication definition to **replicate all columns**.

2. Define the subscription:

```
define subscription t1_sub
for t1_repdef
with replicate at sunak1505x.rep
subscribe to truncate table
go
Subscription t1_sub is in the process of being
defined.
```

3. Validate that all connections are successful:

```
admin who
go
Spid     Name          State               Info
----     ---------- -------------------- --------
139      DSI EXEC   Awaiting Command     104(1) sunak1505x.pri
135      DSI        Awaiting Message     104 sunak1505x.pri
129      DIST       Awaiting Wakeup      104 sunak1505x.pri
130      SQT        Awaiting Wakeup      104:1 DIST sunak1505x.pri
123      SQM        Awaiting Message     104:1 sunak1505x.pri
56       SQM        Awaiting Message     104:0 sunak1505x.pri
143      REP AGENT  Awaiting Command     sunak1505x.pri
144      NRM        Awaiting Command     sunak1505x.pri
156      DSI EXEC   Awaiting Command     106(1) sunak1505x.rep
155      DSI        Awaiting Message     106 sunak1505x.rep
124      SQM        Awaiting Message     106:0 sunak1505x.rep
```

4. Activate subscription with suspension:

```
activate subscription t1_sub
for t1_repdef
with replicate at sunak1505x.rep
```

```
with suspension
go
Subscription 't1_sub' is in the process of being
activated.
```

Note: The with suspension clause suspends the DSI to the target to allow the initial materialization of the replicate database. Any changes to the primary table are placed in the outbound queue of the Replication Server, waiting for the DSI to be resumed.

5. Validate that the DSI is now suspended. Enter:

```
admin who
go
Spid      Name           State               Info
----      ----------     ---------------     --------
139      DSI EXEC       Awaiting Command     104(1) sunak1505x.pri
135      DSI             Awaiting Message     104 sunak1505x.pri
129      DIST            Awaiting Wakeup      104 sunak1505x.pri
130      SQT             Awaiting Wakeup      104:1 DIST
sunak1505x.pri
123      SQM             Awaiting Message     104:1 sunak1505x.pri
56       SQM             Awaiting Message     104:0 sunak1505x.pri
143      REP AGENT       Awaiting Command     sunak1505x.pri
144      NRM             Awaiting Command     sunak1505x.pri
156      DSI EXEC        Suspended            106(1)sunak1505x.rep
155      DSI             Suspended            106 sunak1505x.rep
124      SQM             Awaiting Message     106:0 sunak1505x.rep
```

6. Copy the data from the primary table:

```
% bcp pri..t1 out 't1.bcp' -Usa -P -Ssunak1505x -c
Starting copy...
5 rows copied.
Clock Time (ms.): total = 9 Avg = 1 (555.56 rows per
sec.)
```

7. Insert the data in to the target table:

```
% bcp rep..t1 in 't1.bcp' -Usa -P -Ssunak1505x -c
Starting copy...
5 rows copied.
Clock Time (ms.): total = 30 Avg = 6 (166.67 rows
per sec.
```

8. On the Replication Server, set autocorrection on:

```
set autocorrection on
for t1_repdef
with replicate at sunak1505x.rep
go
autocorrection' is modified for replication
definition 't1_repdef' with replicate at
'sunak1505x.rep'.
```

Note: If there are multiple replication definitions from the same primary table going to the same replication table, make sure that autocorrection is set on all the replication definitions for this table.

9. Resume the connection to let the data stored in the queues be applied to the target table:

```
resume connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is resumed.
```

Note: Because of the overhead associated with autocorrection, SAP recommends that you do not keep autocorrection on when doing normal replication activities.

10. Send through a marker on the primary so that we know when we no longer need autocorrection. In this example, we are going to use **rs_ticket**. On the replication server, validate that **rs_ticket** is available on the DSI. By default, it is available.

```
isql –Usa –Psa_pass –SPRS2
admin config, "connection",sunak1505x,rep,"ticket"
go
Configuration      Config Value       Run Value
--------------    ------------------ ------------
dsi_rs_ticket_report <server default> <server
default>
Default Value   Legal Values Datatype  Status
--------------- ------------   -------- ------------
on              list: on, off string Connection/route
restart required
```

Start the **rs_ticket** marker at the primary database.

```
isql –Usa –P –Ssunak1505x
use pri
go
rs_ticket "done t1"
go
(return status = 0)
```

To check whether **rs_ticket** made it to the replicate database:

```
use rep
go
select ticket from rs_ticket_history where h1 =
"done t1"
go
ticket
-------------------------------------
V=2;H1=done t1;PDB(pri)=11/30/09
12:14:26.253;EXEC(143)=11/30/09 12:14:26.261;B
(143)=19705;DIST(129)=11/30/09
12:14:27.273;DSI(158)=11/30/0912:14:28
294;DSI_T=12;DSI_C=15;RRS=PRS2
```

Note: rs_ticket version 2, which includes all stored procedures and tables to support it, is available in Replication Server 15.1 and later. If you have Replication Servers earlier than 15.1, read the Administration Guide for limitations.

11. Turn off autocorrection when the marker is seen at the replicate database:

```
set autocorrection off
for t1_repdef
with replicate at sunak1505x.rep
go
'autocorrection' is modified for replication
```

```
definition 't1_repdef' with
replicate at 'sunak1505x.rep'.
```

12. Validate the subscription:

```
validate subscription t1_sub
for t1_repdef
with replicate at sunak1505x.rep
go
```

Scenario 7

You can materialize a replicate database from a primary database where user activity cannot be stopped.

It is assumed that activity cannot be stopped at the primary database while the database dump is in progress. This scenario uses the primary database to populate the replicate database and can be used to populate multiple replicate databases with the same primary database dump.

This scenario is very similar to Scenario 3 with these additions:

- The SAP Adaptive Server Enterprise (SAP ASE) and SAP Replication Server must be at 15.5 or higher.
- This solution does not require an MSA replication definition or subscription.
- This solution suspends multiple DSI connections at the same time, so that the same primary database dump can be used to populate multiple replicate databases.

Before You Begin

- At the RSSD, obtain the maintenance users for the connections:

```
use PRS2_RSSD
go
rs_helpuser
go
User Name      Permission(s) Name
------------ ------------------------------
PRS2_id_user    no grants
sa              sa
PRS2_ra         connect source
PRS2_rsi        connect source
PRS2_RSSD_prim  source, primary subscr

User name       Destination DS.DB
------------ ------------------------------------
PRS2_RSSD_maint  sunak1505x.PRS2_RSSD
pri_maint        sunak1505x.pri
pri_maint        sunak1505x.rep
(return status = 0)
```

In this example, “pri_maint” is the maintenance user for both the connection to the primary database, sunak1505x.pri. and replicate database, sunak1505x.rep.

Begin the Resynchronization

1. On the primary SAP ASE, stop the Replication Agent on the primary database:

```
isql -Usa -P -Ssunak1505x
sp_stop_rep_agent pri
go
The Replication Agent thread for database 'pri' is
being stopped.
(return status = 0)
```

2. On the replicate SAP Replication Server, suspend the DSI to the replicate database:

```
isql -Usa -Psa_pass -SPRS2
admin who
go
Spid Name     State              Info
---- ---- ----------------     -------------------
68   DSI EXEC Suspended        105(1) sunak1505x.rep
67   DSI      Suspended        105 sunak1505x.rep
66   SQM      Awaiting Message 105:0 sunak1505x.rep
```

3. Resume the connection to the replicate database and instruct it to wait for the resync marker:

```
resume connection to sunak1505x.rep
skip to resync marker
go
Connection to 'sunak1505x.rep' is resumed.
```

4. Validate that the DSI thread for the replicate database connection is properly set:

```
admin who
go
Spid  Name     State                    Info
----  ----     ---------------- -------------------
120   DSI EXEC Awaiting Command 105(1) sunak1505x.rep
119   DSI      SkipUntil Resync 105    sunak1505x.rep
66    SQM      Awaiting Message 105:0  sunak1505x.rep
```

5. On the primary SAP ASE, start the Replication Agent with the resync option:

```
sp_start_rep_agent pri,'resync'
go
Replication Agent thread is started for database 'pri'.
```

Note: If you changed the location of the secondary truncation point, see the *Reference Manual* for other options on **sp_start_rep_agent** that are available.

6. On the SAP Replication Server, validate that the DSI thread status has now changed:

```
admin who
go
Spid   Name     State             Info
----   ----     ---------------- -------------------
120    DSI EXEC Awaiting Command 105(1) sunak1505x.rep
119    DSI      SkipUntil Dump   105 sunak1505x.rep
66     SQM      Awaiting Message 105:0 sunak1505x.rep
```

In the SAP Replication Server log, you see that the resync option has been acknowledged:

```
DSI for sunak1505x.rep received and processed Resync Database
Marker. Waiting for Dump Marker.
```

Note: If the DSI thread does not change to “SkipUntilDump”, restart the process from step 1.

7. Dump the primary database:

```
dump database pri to 'pri.dmp'
go
```

8. On the SAP Replication Server, validate that the DSI thread is now suspended:

```
admin who
go
Spid Name     State             Info
---- ---- ----------------     -------------------
68   DSI EXEC Suspended        105(1) sunak1505x.rep
67   DSI      Suspended        105 sunak1505x.rep
66   SQM      Awaiting Message 105:0 sunak1505x.rep
```

In the SAP Replication Server log, you see that the dump marker has been processed:

```
DSI for ‘sunak1505x.rep’ received and processed Dump Marker. DSI
is now suspended. Resume after database has been reloaded.
The DSI thread for database ‘sunak1505x.rep’ is shutdown.
```

9. Load the replicate database:

```
load database rep from '/c11014900/sybase1520x/ASE-
15_0/bin/pri.dmp'
go
```

10. Bring the replicate database online:

```
online database rep
go
```

11. Add the maintenance user of the replicate connection to the replicate database and grant the appropriate privileges before the resynchronization.

```
use rep
go
sp_adduser pri_maint
go
New user added.
(return status = 0)
grant execute on rs_get_lastcommit to pri_maint
go
grant all on rs_lastcommit to pri_maint
go
```

Note: If the maintenance user is new to this server, add the maintenance user login to this server, and grant the replication role:

```
sp_addlogin <maintenance user>, <maintenance userpassword>,
<database>
go
sp_role "grant", replication_role, <maintenance user>
go
```

If the password to the maintenance user is unknown, set a new password in Adaptive Server. In the Replication Server, change the password of the maintenance user to match with what was assigned in the Adaptive Server:

```
alter connection to <dataserver>.<database>
set password to <new_password>
go
```

12. Truncate the `rs_lastcommit` table at the replicate database:

```
use rep
go
truncate table rs_lastcommit
go
```

13. Release the secondary truncation point from the replicate database:

```
dbcc settrunc (ltm,ignore)
go
secondary trunc page secondary trunc state dbrepstat
-------------------- --------------------- -----
2668                        0                  172
generation id database id database name ltl version
------------- ------------ ------------- -------
0                  8               rep         720
```

14. Resume connection to the replicate database. If the DSI stays suspended, look at the Replication Server logs for any errors encountered while resuming the DSI:

```
resume connection to sunak1505x.rep
go
Connection to 'sunak1505x.rep' is resumed
```

Scenario 8

You can materialize a replicate table from a primary table where user and replication activity cannot be stopped and replication to other tables in the replicate database must continue.

In the direct_load method, the replicate SAP Replication Server logs directly into the primary SAP ASE database, and selects the rows to be materialized for the table. Datatype translations and custom function strings are utilized before applying the data to the replicate table. During materialization, replication activity continues from the primary database and the rows for the table are placed into a catchup queue in SAP Replication Server. The data in the catchup queue is automatically applied, once the rows from the initial select are applied at the target. At which time, the subscription is marked VALID, and activities from the primary table are replicated to the replicate table.

This method can only be used if the site version of SAP Replication Server is ‘1571100’ or later. To check for site version:

```
isql –Usa –Ppassword -Sreplicate_Replication_Server
sysadmin site_version
go
The current site version is 1571100
```

Before You Begin

- Ensure that replication is working from the primary table to the replicate table. Direct_load requires a table replication definition and table level subscription.
- On the replicate database, ensure that the database option, 'select into/bulkcopy/pllsort' is set to true.
- Ensure that an entry for the ASE server of the primary database, is in the `interfaces` file of the replicate Replication Server.

Begin the Resynchronization

1. If the table subscription exists, drop the subscription.

```
isql –Usa –Psa_pass –SPRS2

drop  subscription t1_sub
for  t1_repdef
with replicate at sunak1505x.rep
without purge
go
```

2. At the replicate table, remove the rows that will be replaced by the materialization.

```
isql –Uuser -Ppassword -Sreplicate_ASE_server -
Dreplicate_database

truncate table t1
go
```

3. Create subscription identifying the materialization type as direct_load.

```
isql –Usa –Psa_pass –SPRS2

create subscription t1_sub
for  t1_repdef
with replicate at sunak1505x.rep
without holdlock
direct_load
user puser password ppwd
go
```

where:

- sunak1505x.rep is the replicate dataserver and database.
- *puser* is the user who selects the data from the table in the primary SAP ASE database.
- *ppwd* is the password of the *puser*. You must enter a password if a user is specified in the subscription.

Note:

- The `interfaces` file that the replicate SAP Replication Server uses must contain an entry for the primary SAP ASE where the initial data resides.
- The *puser* must not be the SAP Replication Server maintenance user.

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

4. Check the subscription to get a status on the direct_load.

```
check subscription t1_sub
for  t1_repdef
with replicate at sunak1505x.rep
go
```

Message: `Subscription t1_sub is VALID at the replicate`. This indicates materialization is complete and the catchup queue is drained.

Message: `Subscription t1_sub encountered ERROR`. This indicates that the Replication Server encountered an error when completing the direct load. Check the Replication Server and ASE logs to see the problems encountered.

Messages similar to:

```
Subscription t1_sub has been MATERIALIZED at the replicate.
```

```
 Subscriptions t1_sub progress: catchup, 0% done, -1 commands
remaining.
```

Indicates that the direct_load has not completed or that there may be a problem in completing the direct_load process. This also indicates that the replicate DSI has not received the validation marker yet. In this case, make sure the replication path is working. If there is an error, drop the subscription and start the direct_load process over.

SAMPLE_RS

Many examples in the SAP Replication Server Quick Start Guides use the SAMPLE_RS as the SAP Replication Server.

The option to install and create the SAMPLE_RS can only be done at SAP Replication Server installation time. You can build a SAMPLE_RS if it was not built during SAP Replication Server installation.

This section of the document builds a SAMPLE_RS as if it were being done at the time of SAP Replication Server installation.

Creating a SAMPLE_RS

Create a SAMPLE_RS and SAMPLE_RS_ERSSD when the directories or files do not exist.

If the SAMPLE_RS directories already exist, go to *Refreshing a SAMPLE_RS* section.

1. Go to the SAP Replication Server installation directory. This is one directory level above where REP-15_5 resides. This is the value for $SYBASE.
2. Set the environment variables by sourcing the `SYBASE.csh` file.
3. Edit the `interfaces` files and add these entries:

   ```
   SAMPLE_RS
      master tcp ether localhost 11752
      query tcp ether localhost 11752
   #
   SAMPLE_RS_ERSSD
      master tcp ether localhost 11751
      query tcp ether localhost 11751
   ```

   Change *localhost* to the name of the host machine.

   Warning! There can only be one SAMPLE_RS_ERSSD on a host machine.
4. Go to `$SYBASE/REP*` directory.
5. Create a directory called `samp_repserver`.

   ```
   mkdir samp_repserver
   ```
6. Go to the `samp_repserver` directory.

   ```
   cd samp_repserver
   ```
7. Create a directory called `samp_partition`.

   ```
   mkdir samp_partition
   ```
8. Create a file called `SAMPLE_RS.res` like the following with these modifications:

In the values of the parameters, the `$SYBASE` must reflect the actual name of the path as **rs_init** cannot utilize environment variables.

Parameter	Description
sybinit.release_directory	The actual path or value of `$SYBASE`
rs.rs_idserver_pass	Valid password
rs.rs_rssd_prim_pass	Valid password
rs.rs_rssd_maint_pass	Valid password
rs.rs_ltm_rs_pass	Valid password
rs.rs_rs_sa_pass	Valid password
rs.rs_rs_run_file	Valid file and path that will used to start up SAMPLE_RS. `$SYBASE/REP-15_5/samp_repserver/RUN_SAMPLE_RS` is the suggested location and file name.
rs.rs_rs_errorlog	Valid file and path that will used SAMPLE_RS log file. `$SYBASE/REP-15_5/samp_repserver/SAMPLE_RS.log` is the suggested location and file name.
rs.rs_rs_cfg_file	Valid file and path that will used SAMPLE_RS configuration file. `$SYBASE/REP-15_5/samp_repserver/SAMPLE_RS.cfg` is the suggested location and file name.
rs.rs_erssd_database_dir	Valid path where the ERSSD database will be located. `$SYBASE/REP-15_5/samp_repserver/dbfile` is the suggested location.
rs.rs_erssd_translog_dir	Valid path where the ERSSD database transaction log will be located. `$SYBASE/REP-15_5/samp_repserver/log` is the suggested location.
rs.rs_erssd_backup_dir	Valid path where the ERSSD backups, including the mirror log, will be located. `$SYBASE/REP-15_5/samp_repserver/backup` is the suggested location.
rs.rs_erssd_errorlog_dir	Valid path where the ERSSD log will be located. `$SYBASE/REP-15_5/samp_repserver/errorlog` is the suggested location.
rs.rs_diskp_name	Valid file and path that will be used for the stable device. The file does not have to exist, however, the directory path must exist. `$SYBASE/REP-15_5/samp_repserver/samp_partition/partition1` is the suggested location.
rs.rs_rs_pass	Valid password

See the table for description of the parameters that do not have values in the following resource file and provide their values based on the description.

```
# --- rs_init resource file ----
#
sybinit.release_directory:
sybinit.product: rs
rs.rs_operation: rs_install
#
# --- ID SERVER INFORMATION ----
#
rs.rs_idserver_name: SAMPLE_RS
rs.rs_id_server_is_rs_server: yes
rs.rs_idserver_user: USE_DEFAULT
rs.rs_idserver_pass:
rs.rs_rssd_prim_pass:
rs.rs_rssd_maint_pass:
rs.rs_ltm_rs_pass:
#
# --- REPLICATION SERVER INFORMATION ----
#
rs.rs_name: SAMPLE_RS
rs.rs_rs_sa_user: USE_DEFAULT
rs.rs_rs_sa_pass:
rs.rs_erssd_requires_ltm: no
rs.rs_needs_repagent: yes
rs.rs_rs_run_file:
rs.rs_rs_errorlog:
rs.rs_rs_cfg_file:
rs.rs_charset: USE_DEFAULT
rs.rs_language: USE_DEFAULT
rs.rs_sortorder: USE_DEFAULT
#
# --- ERSSD ----
#
rs.rs_rssd_embedded: yes
rs.rs_erssd_name: SAMPLE_RS_ERSSD
rs.rs_erssd_database_dir:
rs.rs_erssd_translog_dir:
rs.rs_erssd_backup_dir:
rs.rs_erssd_errorlog_dir:
#
# --- DISK PARTITION INFORMATION ----
#
rs.rs_diskp_name:
rs.rs_diskp_lname: partition1
rs.rs_diskp_size: 20
rs.rs_diskp_vstart: 0
#
# --- REMOTE SITE CONNECTION INFORMATION ----
#
rs.rs_rs_user: USE_DEFAULT
rs.rs_rs_pass:

#
# --- ID SERVER INTERFACES INFORMATION ----
```

```
#
rs.do_add_id_server: no
rs.do_add_replication_server: no
```

9. Use **rs_init** to execute this resource file:

   ```
   $SYBASE/REP-15_5/install/rs_init -r SAMPLE_RS.res -T
   T_SEND_CLEARTEXT_PASSWORD
   ```

   If the creation of the SAMPLE_RS fails:

   a) Correct the error.
   b) Stop the SAP Replication Server if it is running.
   c) Remove the files in these directories:
      - `$SYBASE/REP*/samp_repserver/backup/SAMPLE_RS_ERSSD.db`
      - `$SYBASE/REP*/samp_repserver/backup/SAMPLE_RS_ERSSD.log`
      - `$SYBASE/REP*/samp_repserver/backup/SAMPLE_RS_ERSSD.mlg`
      - `$SYBASE/REP*/samp_repserver/dbfile/SAMPLE_RS_ERSSD.db`
      - `$SYBASE/REP*/samp_repserver/errorlog/SAMPLE_RS_ERSSD.out`
      - `$SYBASE/REP*/samp_repserver/log/SAMPLE_RS_ERSSD.log`
      - `$SYBASE/REP*/samp_repserver/RUN_SAMPLE_RS`
      - `$SYBASE/REP*/samp_repserver/SAMPLE_RS.cfg`
      - `$SYBASE/REP*/samp_repserver/SAMPLE_RS.log`
   d) Rerun **rs_init**.

10. Verify that the SAMPLE_RS is available. Log in to the SAMPLE_RS.

    ```
    isql -Usa -Ppassword -SSAMPLE_RS
    ```

    The password to the sa user ID is the value that was filled in the resource file for rs.rs_rs_sa_pass.

11. Enter:

    ```
    admin who
    go
    ```

12. End the session by exiting from the **isql** client.

    ```
    exit
    go
    ```

Refreshing a SAMPLE_RS

Refresh a SAMPLE_RS if the directories already exist.

1. Go to the SAP Replication Server installation directory. This is one directory level above where REP-15_5 resides. This is the value for $SYBASE.

2. Set the environment variables by sourcing the `SYBASE.csh` file.
3. In the `$SYBASE/interfaces` file, verify that there are entries for SAMPLE_RS and SAMPLE_RS_ERSSD for this host machine:

```
SAMPLE_RS
   master tcp ether localhost 11752
   query tcp ether localhost 11752
#
SAMPLE_RS_ERSSD
   master tcp ether localhost 11751
   query tcp ether localhost 11751
```

   Change *localhost* to the name of the host machine. The port numbers can be any available port on the host. If you want to use different port numbers than what is shown in the interfaces file, ensure that the SAMPLE_RS is shutdown before changing the port number of the SAMPLE_RS and SAMPLE_RS_ERSSD in the interfaces file.

   Note: The SAMPLE_RS is created with an ERSSD. One cannot have multiple ERSSDs (or ASA servers) of the same name, on the same host, running at the same time.

4. If SAMPLE_RS is currently running, shutdown the existing SAMPLE_RS.

```
isql -Usa -Ppassword -SSAMPLE_RS
shutdown
go
```

5. Verify that the following directories and files exist under `$SYBASE/REP*/samp_repserver`:

```
-rwxr-xr-x  1 sybase sybase   93 Dec 18 19:07 rs_init-
SAMPLE_RS.sh*
-rwxr-xr-x  1 sybase sybase 1667 Dec 18 19:08 SAMPLE_RS.res*
drwxr-xr-x  2 sybase sybase 4096 Dec 18 19:07 samp_partition/
```

6. If the following directories exist, ensure that the files in the directory do not contain the name of the SAMPLE_RS_ERSSD.

```
drwxr-x---  2 sybase sybase 4096 Dec 19 17:50 backup/
drwxr-x---  2 sybase sybase 4096 Dec 19 17:50 dbfile/
drwxr-x---  2 sybase sybase 4096 Dec 19 17:50 errorlog/
drwxr-x---  2 sybase sybase 4096 Dec 19 17:50 log/
```

   For example, in the `$SYBASE/REP*/samp_repserver/dbfile` directory, the file SAMPLE_RS_ERSSD.db must not exist.

   If the files exist, one either needs to choose another name for the SAMPLE_RS_ERSSD or delete the file under the directory.

   Note: These directories do not have to exist for the SAMPLE_RS to be successfully built.

7. If the SAMPLE_RS was previously created, the following files need to be renamed or deleted:

```
-rwxr-xr-x  1 sybase sybase  290 Dec 19 17:54 RUN_SAMPLE_RS*
-rw-------  1 sybase sybase 1149 Dec 19 17:54 SAMPLE_RS.cfg
-rw-r-----  1 sybase sybase 8906 Dec 19 17:54 SAMPLE_RS.log
```

Note: The `RUN` and `cfg` files contain all the information, including directories and filenames, about the previous SAMPLE_RS and SAMPLE_RS_ERSSD.

8. Validate the entries in `$SYBASE/REP*/samp_repserver/SAMPLE_RS.res`.
9. Set the environment variables by sourcing the `$SYBASE/SYBASE.csh` file.
10. Go to the `$SYBASE/REP*/samp_repserver` directory.
11. Execute **rs_init-SAMPLE_RS.sh**:

    ```
    ./ rs_init-SAMPLE_RS.sh
    ```

    If the file, `rs_init-SAMPLE_RS.sh` does not exist, execute **rs_init** using the `SAMPLE_RS.res` file:

    ```
    $SYBASE/REP-15_5/install/rs_init -r SAMPLE_RS.res -T
    T_SEND_CLEARTEXT_PASSWORD
    ```

12. Verify that the SAMPLE_RS is available. Log in to the SAMPLE_RS.

    ```
    isql -Usa -Ppassword -SSAMPLE_RS
    ```

    The password to the sa user ID is the value that was filled in the resource file for rs.rs_rs_sa_pass.

13. Enter:

    ```
    admin who
    go
    ```

14. End the session by exiting from the **isql** client.

    ```
    exit
    go
    ```

Removing the SAMPLE_RS from the System

Remove the SAMPLE_RS from the system if it is not part of a multiple SAP Replication Server environment.

1. Go to `$SYBASE`.
2. Source the `SYBASE.csh`.
3. If the SAMPLE_RS is running, shut it down:

    ```
    isql –Usa –Ppassword -SSAMPLE_RS
    shutdown
    go
    ```

4. View the file that starts up the SAMPLE_RS.

    The default installation is at `$SYBASE/REP*/samp_repserver/RUN_SAMPLE_RS`.

5. Make note of the file location identified by the –C parameter (the SAP Replication Server configuration file) the –E (the SAP Replication Server log) and the –I (the interfaces file).
6. Go to the file identified by the –C parameter (the configuration file). Record the files and directories in the following lines in the configuration file:

```
erssd_errorlog
erssd_dbfile
erssd_translog
erssd_logmirror
erssd_backup_dir
```

7. For the filename referenced by the **erssd_dbfile** parameter, note the name of the ERSSD.

   The name of the ERSSD will be the name in front of the `.db` file extension. If the entry looks like `/opt/sybase/REP-15_5/samp_repserver/dbfile/SAMPLE_RS_ERSSD.db`, SAMPLE_RS_ERSSD is the name of the ERSSD.

8. In the directory identified in the **erssd_backup_dir**, remove all files with the name of the ERSSD.

   For example, `SAMPLE_RS_ERSSD.db` must be removed.

9. Delete the file identified in the:
   - erssd_errorlog
   - erssd_dbfile
   - erssd_translog
   - erssd_logmirror

10. In the `$SYBASE/REP*/samp_repserver` directory, delete the RUN file, the configuration file, and the SAMPLE_RS log file.

11. Edit the `interfaces` file to remove the entries for the SAMPLE_RS and the SAMPLE_RS_ERSSD.

Index

O

P

R

S

T

U

V

W