



Quick Start Guide

**Replication Server[®] Option for
Oracle 15.7.1 ESD #2**

Microsoft Windows, UNIX, and Linux

DOCUMENT ID: DC01845-01-1571-01

LAST REVISED: October 2012

Copyright © 2012 by Sybase, Inc. All rights reserved.

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

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

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

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

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

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

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

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

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

Contents

Conventions	1
Oracle Replication Setup	3
System Architecture	3
Prerequisites	4
Planning Your Installation	4
System Requirements	4
Identify the Sybase Installation Directory	5
Accessing and Downloading Products from SPDC	5
Installing Replication Components	6
Installing Replication Server	7
Installing the Oracle Instant Client Libraries	11
Installing Replication Agent	13
Configuring Replication Components	16
Configuring Replication Server	16
Configuring ExpressConnect for Oracle	20
Configuring Replication Agent	22
Testing Replication	33
Resetting the Primary Oracle Database for Replication	33
Obtaining Help and Additional Information	34
Technical Support	34
Downloading Sybase EBFs and Maintenance Reports	35
Sybase Product and Component Certifications	35
Creating a MySybase Profile	35
Accessibility Features	36
Glossary	36
Index	43

Contents

Conventions

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

Style conventions

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

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

Run:

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

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

Syntax conventions

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

Case-sensitivity

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

Terminology

Replication Agent™ is a generic term used to describe Replication Agent for Oracle for Linux, Unix and Windows.

Oracle Replication Setup

Implement a sample replication system for Oracle.

Note: Use the tasks in the *Replication Server Option for Oracle 15.7.1 ESD #2 Quick Start Guide* to set up replication only for proof-of-concept or testing, not for production. Only basic Oracle features are addressed in the example scenario. In particular, Oracle Automatic Storage Management (ASM) and Real Application Clusters (RAC) are not addressed in this guide, although Replication Agent does support these features. See the *Replication Agent Primary Database Guide*.

The replication of sequences is not discussed in the *Replication Server Option for Oracle 15.7.1 ESD #2 Quick Start Guide*. To set up sequence replication, see the *Replication Agent Installation Guide* and the *Replication Agent Primary Database Guide*.

See also


- *Planning Your Installation* on page 4
- *Installing Replication Components* on page 6
- *Configuring Replication Components* on page 16

System Architecture

There are multiple components required to implement a replication system for Oracle.

- A primary Oracle data server
- A replicate Oracle data server
- A Replication Server®
- The Replication Server Options components:
 - Replication Agent
 - ExpressConnect for Oracle

Figure 1: Sybase replication system for Oracle



Prerequisites

Review the prerequisites before you begin setting up replication for Oracle.

Make sure:

- You have an understanding of Sybase replication and are familiar with Oracle data servers.
- You have valid licenses for Replication Server and Replication Server Options.
- Oracle Database 10g or 11g Enterprise Edition is installed and configured to serve as the source (primary) database from which Replication Agent replicates transactions and as the target (replicate) database to which Replication Server with ExpressConnect for Oracle replicates transactions.
- Install the Oracle JDBC™ driver on the host machine on which Replication Agent resides. If Oracle 10g or 11g Release 1 is the primary database, install the Oracle JDBC thin driver for Oracle 10g or 11g and for JDK 1.4 and 1.5. If the primary database is Oracle 11g Release 2, install the Oracle JDBC thin driver 11.2 for JDK 1.6.
- TCP/IP connectivity is available.
- Any operating system patches required for Java 7.0 have been installed on the host on which Replication Agent resides.

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

Planning Your Installation

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

System Requirements

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

For Windows:

Component	Memory (RAM)	Disk space
Replication Server	512MB	380MB
Replication Agent	512MB	300MB

For UNIX:

Component	Memory (RAM)	Disk space
Replication Server	512MB	520MB
Replication Agent	512MB	300MB

For Linux:

Component	Memory (RAM)	Disk space
Replication Server	512MB	400MB
Replication Agent	512MB	300MB

To improve performance by avoiding multiple network hops, install Replication Server, ExpressConnect for Oracle, and the target database on the same machine.

See the *Replication Server Option for Oracle 15.7.1 ESD #2 Release Bulletin*.

Identify the Sybase Installation Directory

Determine the directory in which to install the replication components.

Replication Agent, Replication Server, and ExpressConnect for Oracle are installed in the same base directory, which is identified by the SYBASE environment variable. Select a path on a host drive to be the recipient of the Sybase installation and configuration activities. Replication Agent must be installed on the same type of host as the primary Oracle instance and must be able to directly access the primary Oracle redo and archive logs.

Replication Agent Installation Directory

Replication Agent for Oracle 15.7.1 ESD #2 uses the RAX-15_5 directory name. If you have an existing Replication Agent 15.5 or 15.6 installation and want to keep this version separate, install Replication Agent for Oracle 15.7.1 ESD #2 into a different directory. Replication Agent for Oracle 15.7.1 ESD #2 installer does not overwrite any previous version of Replication Agent installed in the existing RAX-15_5 directory.

Accessing and Downloading Products from SPDC

If you are not using the installation CD, download replication components from the Software Product Download Center (SPDC).

Prerequisites

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

Task

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

1. Go to <https://sybase.subscribenet.com/control/sybs/login>, and enter your login ID and password.
2. Select **Replication Server**.
3. From the high-level Product Information page, select **Replication Server** again.
4. Select the Replication Server software you need by version and platform.

The Software Terms and Conditions page appears, listing countries and regions.

5. Either:
 - Select **I Agree**, go to the Product Download Page, and continue with step 6, or
 - Select your country or region to read the terms.

The software licence agreement for the software appears. Read the terms and either:

 - Select **I Agree** to accept the terms, go to the Product Download Page, and continue with step 6, or
 - Select **Cancel** to return to the high-level Product Information page.
6. The Product Download Page lists the software and documentation for the product. Select **FTP Download** or **Advanced Download Options** for the Replication Server software to download and save the compressed files to your local drive.

Note: For each product, the documentation download contains installation-related documents, and the rest of the documentation set can be accessed online.

7. Return to the detailed Product Information page that you accessed in step 3, and select the RSO software by primary database, version, and platform.

The Software Terms and Conditions page appears, listing countries and regions.
8. Navigate the Software Terms and Conditions pages as you did in step 5.
9. From the Product Download Page, select **FTP Download** or **Advanced Download Options** for the components to download, and save the compressed files to your local drive.

For RSO, download Replication Agent 15.7.1 ESD #2.
10. Uncompress and extract all the installation images to your local drive.

Installing Replication Components

Review the procedures for installing the replication components.

See also

- *Configuring Replication Components* on page 16

Installing Replication Server

Install Replication Server using the setup program.

Prerequisites

- Allocate a disk partition of at least 20MB for each Replication Server you are installing. You can add more partitions later, if necessary. Check each partition to make sure it is available and has write permissions.
- Allocate the entire partition to the Replication Server. If you allocate only a portion of the partition to Replication Server, you cannot use the remainder for any other purpose.

Task

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

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

If the installer does not start automatically, double-click `setup.exe` or select **Start > Run** and enter the following, 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 the *directory_name*, include its full path.

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

```
./setup.bin
```

If you are installing using the CD, mount it and start the installer. The location of the **mount** command is site-specific and may differ from the instructions shown here. If you cannot mount the CD drive using the path shown, check your operating system documentation or contact your system administrator.

- On HP-UX Itanium:
 - a. Log out, log in as “root”, and mount the CD using:

```
mount -F cdfs -o ro device_name /mnt/cdrom
```

- b. Log out as “root” and log in again as “sybase.”

- c. Start the installer:

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

- On IBM AIX:

- a. Log in as “sybase” and mount the CD:

```
mount -v 'cdrfs' -r device_name /mnt/cdrom
```

- b. Start the installer:

```
cd /device_name
./setup.bin
```

- On Solaris:

- a. The operating system automatically mounts the CD. Log in as “sybase.”

Note: If you get CD-reading errors, check your operating system kernel to make sure that the ISO 9660 option is turned on.

If you have previously installed a Sybase CD on your system, the # sign interferes with the installation process. Before installing your current CD, either:

- Restart your system, or
- Eject the CD. Delete the *Volume Label* file in */vol/dsk*, and reinsert the CD.

-
- b. Start the installer:

```
cd /cdrom/Volume Label
./setup.bin
```

- On Linux and Linux on POWER:

- a. Log in as “sybase” and mount the CD:

```
# mount -t iso9660 /dev/cdrom /mnt/cdrom
```

- b. Start the installer:

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

where:

- *cdrom* and */mnt/cdrom* is the directory (mount point) you specified when mounting the CD drive.
- *device_name* is the CD device drive.
- *setup.bin* is the executable file name for installing Replication Server.

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

4. In the Introduction window, click **Next**.

5. Specify where to install Replication Server:

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

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

6. Select the type of installation:

Option	Description
Typical	Installs the default components. This is recommended for most users.
Full	Installs every Replication Server component, including all the supported language modules.
Custom	Lets you select the components to install. Some components are automatically installed if they are required to run your selected components.

Click **Next**.

7. Select the geographic location, agree to the license agreement, then click **Next**.
8. On the Sybase Software Asset Management License Server window:

Option	Choose
Specify license keys	Either: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 .
Continue installation without a license key	If you do not have licenses for any of the Replication Server components, select this option, then click Next to proceed. The installer allows you to install and use Replication Server 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 Replication Server License Installer.

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

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. On the Start Sample Replication Server window, select:

Option	Description
Yes	To 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: <ul style="list-style-type: none"> • single-byte character – enter 6 to 30 characters. • double-byte character – enter 3 to 15 characters.
No	To complete the installation. You can manually configure a full-featured Replication Server and start a sample Replication Server after installation.

Click **Next**.

12. When the installation is complete, click **Done**.

Verifying the Installation

Verify that Replication Server has been successfully installed.

1. In a command window, change to the directory where you installed Replication Server.
2. Set the environment variables by sourcing the `SYBASE.csh` file (UNIX or Linux) or by executing the `SYBASE.bat` file (Windows).
3. Log in to Replication Server:

```
isql -Username -Ppassword -SSAMPLE_RS
```

4. Verify the Replication Server version:

```
admin version
go
```

5. Exit the `isql` session with the `quit` command.

Installing the Oracle Instant Client Libraries

Install the Oracle Instant Client libraries.

As of version 15.7.1, ExpressConnect for Oracle no longer ships with the Oracle Instant Client libraries. Download and install these libraries after you install ExpressConnect for Oracle.

1. Create a temporary directory `<tempdir>`.
2. Go to the Oracle Web site at <http://www.oracle.com>.
3. Select **Downloads > Databases > Instant Client**.
4. Select the download link for your platform:

Table 1. Download Links By Platform

Platform	Instant Client Download Link
Windows x86 32-bit	Instant Client for Microsoft Windows (32-bit)
Windows x86-64 64-bit	Instant Client for Microsoft Windows (64-bit)
Linux x86-64 64-bit	Instant Client for Linux x86-64
Linux pSeries 64-bit	Instant Client for Linux on Power (64-bit)
Solaris SPARC 64-bit	Instant Client for Solaris Operating System (SPARC) (64-bit)
Solaris x86 64-bit	Instant Client for Solaris x86-64
HP-UX Itanium 64-bit	Instant Client for HP-UX Itanium (64-bit)
IBM AIX pSeries 64-bit	Instant Client for AIX5L (64-bit)

5. Select **Accept License Agreement for You must accept the OTN Development and Distribution License Agreement for Instant Client to download this software.**
6. Go to the specified version and download the Instant Client libraries package to the temporary directory <tempdir>:

Table 2. Instant Client Libraries Packages

Platform	Version	Instant Client Libraries Package
Windows x86 32-bit	10.2.0.4	instantclient-ba- sic- win32-10.2.0.4.zip
Windows x86-64 (64-bit) running Windows-XP, Windows Server 2000, Windows Server 2003, Windows Server 2008, Windows Vista, or Windows 7	10.2.0.5	instantclient-ba- sic- win64-10.2.0.5.zip
Linux x86-64 64-bit	10.2.0.4	basic-10.2.0.4.0- linux-x86_64.zip
Linux pSeries 64-bit	10.2.0.4	basic-10.2.0.4.0- linux-ppc64.zip
Solaris SPARC 64-bit	10.2.0.4	basic-10.2.0.4.0- solaris- sparc64.zip

Platform	Version	Instant Client Libraries Package
Solaris x86 64-bit	10.2.0.4	basic-10.2.0.4.0-solaris-x86-64.zip
HP-UX Itanium 64-bit	10.2.0.4	basic-10.2.0.4.0-hpux-ia64.zip
IBM AIX pSeries 64-bit	10.2.0.4	basic-10.2.0.4.0-aix-ppc64.zip

7. Extract the Oracle Instant Client package into the temporary directory <tempdir> using a zip utility (for Windows) or the **UnZip** utility (for UNIX platforms.)
For UNIX platforms, the UnZip utility can be downloaded from: <https://updates.oracle.com/unzips/unzips.html>.
8. Copy the extracted Oracle Instant Client library files located in <tempdir> \instantclient_10_2 directory to the \$SYBASE\REP-15_5\connector \lib directory.

Installing Replication Agent

Install Replication Agent using the GUI wizard.

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

1. Log in to the Replication Agent host machine using an operating system user account with authority to start, stop, and administer the Replication Agent instance (for example, the “sybase” user).
2. Close all nonessential applications, and minimize any open windows.
3. Insert the Replication Agent distribution media in the appropriate drive.
4. Start the installation program:
 - On Microsoft Windows platforms, the installation program should start automatically. If it does not, select **Start > Run**, and enter:

```
x:\setup.exe
```

where *x*: is your media drive.

You can also start the installation program from Microsoft Windows Explorer by double-clicking the `setup.exe` file.
 - On UNIX platforms, enter:

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

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

6. Select your geographic location.
7. Read the Sybase license agreement and select **I agree to the terms of the Sybase license for the install location specified.**

Note: You must agree to the terms of the software license before you can continue.

Click **Next**.

8. Specify an installation directory. The default installation directory is:
 - Existing `%SYBASE%` or `c:\sybase` on Microsoft Windows platforms
 - Existing `$SYBASE` or `/opt/sybase` on UNIX platforms

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

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

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

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

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

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

9. Select one of these options to enter the license:
 - **Specify license keys** – browse to or specify the license file.
 - **Use previously deployed license server** – use a previously deployed license server. Enter the host name of the machine where the license server is running and the port number if the port number you are using is not the default.
 - **Continue installation without a license key** – install and use Replication Agent without a license for a grace period of 30 days. To continue using Replication Agent after the end of the grace period, obtain a valid license from the Sybase Product Download Center at <http://www.sybase.com/detail?id=1025266>, and install it. See the *SySAM Users Guide*.

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

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

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

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

Click **Next**.

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

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

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

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

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

Verifying the Installation

Verify that Replication Agent has been successfully installed.

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

- For Windows:

```
rao -v
```

- For UNIX or Linux:

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

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

Configuring Replication Components

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

See also

- *Testing Replication* on page 33

Configuring Replication Server

Review the procedures for configuring Replication Server for Oracle.

Configuring Replication Server for Replication to the Target Data Server

Create a Replication Server connection to the replicate data server.

1. Navigate to the `$SYBASE/RAX-15_5/scripts/oracle` (UNIX or Linux) or `%SYBASE%\RAX-15_5\scripts\oracle` (Windows) directory.
2. Make a copy of the `oracle_create_rs_standby_connection.sql` script named `my_oracle_create_rs_standby_connection.sql`.
3. Before executing the `my_oracle_create_rs_standby_connection.sql` script against your Replication Server, change all occurrences of `{rds}`, `{rdb}` to the name of the connection that Replication Server uses to connect to ExpressConnect for Oracle where:
 - `rds` is the replicate Oracle System ID (SID) name.
 - `rdb` is any valid identifier. Sybase recommends that you use the replicate Oracle SID name.
 - `maintuser` and `password` are the user name and password created in ExpressConnect for Oracle.

For example:

```
create connection to OracleSID_name.OracleSID_name
using profile rs_oracle_to_oracle;eco
set username to maintuser
set password to "password"
go
```

Note: In the previous example, **password** is a Replication Server reserved word and therefore, must be enclosed in double quotes.

If you are not using an ExpressConnect for Oracle connection profile, set **dsi_proc_as_rpc** to 'on' in the **create connection** command. For example:

```
create connection to OracleSID_name.OracleSID_name
set error class <error_class>
set function string class <function_class>
```

```
set username <userid>
set password <password>
set dsi_dataserver_make to 'ora'
set dsi_dataserver_type to 'oci'
set batch to 'off'
set dsi_proc_as_rpc to 'on'
```

If you are using an ExpressConnect for Oracle connection profile, **dsi_proc_as_rpc** is enabled by default.

- For configuring the server connectivity for ExpressConnect for Oracle, copy the `tnsnames.ora` file to the `\RS_installation_directory\REP-15_5\connector\oraoci\network\admin` directory.

- Create the connection to the replicate database:

```
isql -Uusername -Ppassword -SSAMPLE_RS -i
my_oracle_create_rs_standby_connection.sql
```

- Log in to `SAMPLE_RS`, and verify the Replication Server connection to the standby database:

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

You should see a message returned for the DSI connection `{rds}.{rdb}`. Verify that the status is “Awaiting Message” or “Awaiting Command.”

- Exit the `isql` session with the `quit` command.

Configuring Replication Server for Replication from the Primary Database

Configure Replication Server for replication from the primary database.

Creating a Replication Server Connection to the Primary Database

Create a Replication Server connection to the primary database.

- Go to the `$SYBASE/RAX-15_5/scripts/oracle` (UNIX or Linux) or `%SYBASE%\RAX-15_5\scripts\oracle` directory (Windows).

- Make a copy of the `oracle_create_rs_primary_connection.sql` script.

```
cp oracle_create_rs_primary_connection.sql
my_oracle_create_rs_primary_connection.sql
```

- Before executing the `my_oracle_create_rs_primary_connection.sql` script against your Replication Server, change all occurrences of `{pds}`. `{pdb}` to the name of the Replication Server connection used for the connection from Replication Agent for Oracle, where:

- `pds` is the name of the primary Oracle server.
- `pdb` is the primary Oracle System ID (SID).

For example, `NY.NYora92`.

Note: Save these values to be used later for the **rs_source_ds** and **rs_source_db** parameters in the Replication Agent `oracle.rs` file.

4. Change `sys` and `sys_pwd` to the user ID and password of the Oracle user who must have permission to apply DML operations against all user tables to be replicated, where:
 - `sys` is the user ID of the Oracle user (**sys**).
 - `sys_pwd` is the password of the Oracle user (**change_on_install**).

Note: This user must be a valid user in the Oracle database, and it must be different from the user defined by the Replication Agent **pds_username** parameter, which is used for Replication Agent connections to the primary database.

This example creates a Replication Server connection to the primary database:

```
create connection to NY.NYora92
set error class rs_oracle_error_class
set function string class rs_oracle_function_class
set username sys
set password change_on_install
with log transfer on, dsi_suspended
go
```

5. Execute the script in Replication Server:

```
isql -Uusername -Ppassword -SSAMPLE_RS -i
<script_path>\my_oracle_create_rs_primary_connection.sql
```

where `<script_path>` is the full path to `my_oracle_create_rs_primary_connection.sql`.

You see a message indicating that the Replication Server connection to the primary database has been created.

See also

- *Granting create object Permission to the rs_username User* on page 18
- *Creating a Database Replication Definition* on page 19
- *Creating a Database Replication Subscription* on page 20

Granting create object Permission to the rs_username User

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

The **rs_username** user must have **create object** permission before it can be used to create a replication definition. You must grant this permission manually from the RSSD.

1. Log in to Replication Server:

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

2. Enter:

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

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

3. Log in to the RSSD:

```
isql -USAMPLE_RS_RSSD_prim -PSAMPLE_RS_RSSD_prim_ps
-SSAMPLE_RS_RSSD
```

where `SAMPLE_RS_RSSD`, `SAMPLE_RS_RSSD_prim`, `SAMPLE_RS_RSSD_prim_ps` are the RSSD database name, user ID, and password needed to log in to the RSSD.

4. Verify that the `rs_username` user has **create object** permission:

```
select * from rs_users where username = 'SAMPLE_RS_ra'
go
```

The `rs_users` system table should indicate that the `rs_username` user has **create object** permission.

See also

- *Creating a Database Replication Definition* on page 19
- *Creating a Database Replication Subscription* on page 20

Creating a Database Replication Definition

Create a database replication definition.

1. Navigate to the `%SYBASE%\RAX-15_5\scripts\sybase` (Windows) or `$SYBASE/RAX-15_5/scripts/sybase` (UNIX or Linux) directory.
2. Navigate to the `$SYBASE/RAX-15_2/scripts/sybase` (UNIX or Linux) or `%SYBASE%\RAX-15_2\scripts\sybase` (Windows) directory.
3. Make a copy of the `rs_create_test_db_repdef.sql` script named `my_rs_create_db_repdef.sql`.

```
cp rs_create_test_db_repdef.sql
my_rs_create_db_repdef.sql
```

4. Before executing the `my_rs_create_db_repdef.sql` script, change all occurrences of `{pds}`.`{pdb}` to the name of the connection string you defined for the primary database, where:
 - `pds` is the name of the primary Oracle server.
 - `pdb` is the primary Oracle System ID (SID).

For example, `NY.NYora92`.

Note: Save these values to be used later for the `rs_source_ds` and `rs_source_db` parameters in the Replication Agent `oracle.rs` file.

This is a database replication definition example:

```
create database replication definition NY_repdef1
with primary at NY.NYora92
```

Oracle Replication Setup

```
replicate DDL
go
```

5. Connect to the Replication Server and enter:

```
isql -Username -Ppassword -SSAMPLE_RS -i
my_rs_create_db_repdef.sql
```

You see a message indicating that the database replication definition has been created.

See also

- *Creating a Database Replication Subscription* on page 20

Creating a Database Replication Subscription

Create a database replication subscription.

1. Make a copy of the `rs_create_test_db_sub.sql` script named `my_rs_create_db_sub.sql`.
2. Before executing the `my_rs_create_db_sub.sql` script, change all occurrences of `{pds}`, `{pdb}` and `{rds}`, `{rdb}` to the appropriate connection name, where:
 - `pds` is the name of the primary Oracle server.
 - `pdb` is the primary Oracle System ID (SID).
 - `rds` is the replicate Oracle SID.
 - `rdb` is any valid identifier. Sybase recommends that you use the replicate Oracle SID.

Note: Save the values of `pds` and `pdb` to be used later for the `rs_source_ds` and `rs_source_db` parameters in the Replication Agent `oracle.rs` file.

Here is an example of how to create a database replication subscription:

```
create subscription NY_sub1
for database replication definition NY_repdef1
with primary at NY.NYora92
with replicate at orcl1.oratest2
without materialization
go
```

3. Connect to the Replication Server and enter:

```
isql -Username -Ppassword -SSAMPLE_RS -i my_rs_create_db_sub.sql
```

You see a message indicating that the subscription is being created.

Configuring ExpressConnect for Oracle

Configure ExpressConnect for Oracle to set up connections between Oracle and Replication Server.

1. Copy the Oracle server `tnsnames.ora` file to the `$SYBASE`
`\REP-15_5\connector\oraoci\network\admin` directory.

2. Determine the Oracle user ID and password used to connect from Replication Server. See “Oracle replicate database permissions” in the *Replication Server 15.7.1 Heterogeneous Replication Guide*.
3. Restart Replication Server.
4. Using **isql** create a connection to Replication Server using the alias name defined in the Oracle `tnsnames.ora` file, user ID, and password combination. For example:

```
create connection to
<tnsnames_alias>.<ora_rdb_name>
using profile rs_oracle_to_oracle;eco
set username <userid>
set password <password>
set batch to 'off'
```

where:

- **tnsnames_alias** is the case-sensitive name identifying the replicate Oracle database in the `tnsnames.ora` file. For example:

```
<tnsnames_alias> =
(DESCRIPTION =
  (ADDRESS = (PROTOCOL = TCP) (HOST = hostname) (PORT = 1521))
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = orcl)
  )
)
```

- **ora_rdb_name** can be any name which best describes the replicate Oracle database. For example, `orcl11g`.

If you are not creating the connection using one of the Replication Server connection profiles for ExpressConnect for Oracle, set **dsi_proc_as_rpc** to `on` in the **create connection** command. For example:

```
create connection to <tnsnames_alias>.<ora_rdb_name>
set error class <error_class>
set function string class <function_class>
set username <userid>
set password <password>

set batch to 'off'
set dsi_proc_as_rpc to 'on'
```

If you are using one of the ExpressConnect for Oracle connection profiles, **dsi_proc_as_rpc** is enabled by default.

See *Replication Server Reference Manual* for more information on the **create connection** command. For more information about configuring ExpressConnect for Oracle, see the *ExpressConnect for Oracle Configuration Guide*.

Configuring Replication Agent

Review the procedures for configuring Replication Agent for Oracle.

Primary Oracle Database Configuration for Replication

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

Verifying the Current Archive Setting of the Redo Logs

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

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

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

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

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

3. To turn on log archiving:

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

See the *Replication Agent Primary Database Guide*.

Verifying the Supplemental Logging of Primary Key Data

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

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

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

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

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

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

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

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

Configuring Oracle LogMiner

Configure Oracle LogMiner on the primary Oracle database.

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

```
@dbms1m.sql
```

2. After LogMiner is installed, create a public synonym so that you do not have to log in as the owner to execute LogMiner functions:

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

Creating Oracle User and Grant Permissions

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

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

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

```
CREATE USER RA_USER PROFILE "DEFAULT" IDENTIFIED
BY "sybase" DEFAULT TABLESPACE "USERS" ACCOUNT
UNLOCK;
GRANT ALTER ANY PROCEDURE TO RA_USER;
GRANT ALTER ANY TABLE TO RA_USER;
GRANT ALTER SESSION TO RA_USER;
GRANT ALTER SYSTEM TO RA_USER;
GRANT CONNECT TO RA_USER;
GRANT CREATE PROCEDURE TO RA_USER;
GRANT CREATE PUBLIC SYNONYM TO RA_USER;
GRANT CREATE SESSION TO RA_USER;
GRANT CREATE TABLE TO RA_USER;
GRANT DROP ANY PROCEDURE TO RA_USER;
GRANT DROP ANY TABLE TO RA_USER;
GRANT DROP PUBLIC SYNONYM TO RA_USER;
GRANT EXECUTE_CATALOG_ROLE TO RA_USER;
GRANT EXECUTE ON SYS.DBMS_FLASHBACK TO RA_USER;
GRANT EXECUTE ON SYS.DBMS_LOCK TO RA_USER;
GRANT RESOURCE TO RA_USER;
GRANT SELECT ANY TRANSACTION TO RA_USER;
GRANT SELECT_CATALOG_ROLE TO RA_USER;
GRANT SELECT ON SYS.ARGUMENT$ TO RA_USER;
GRANT SELECT ON SYS.ATTRIBUTE$ TO RA_USER;
GRANT SELECT ON SYS.CCOL$ TO RA_USER;
```

Oracle Replication Setup

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

See the *Replication Agent Primary Database Guide*.

Verifying the Oracle User Roles Created by the Replication Agent

Validate the required Oracle settings.

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

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

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

Creating the Replication Agent Instance

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

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

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

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

- For Oracle 10g:

```
$ORACLE_HOME/jdbc/lib/ojdbc14.jar
```

- For Oracle 11g Release 1:

```
$ORACLE_HOME/jdbc/lib/ojdbc5.jar
```

- For Oracle 11g Release 2:

```
$ORACLE_HOME/jdbc/lib/ojdbc6.jar
```

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

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

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

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

For Windows, add this to the CLASSPATH environment variable:

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

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

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

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

4. (Optional) If your operating system has a Java 6.0.x-compatible JRE or JDK installed, you can use the Oracle **iSQL** (the browser-based interface to SQL*PLUS) demo items to connect to Oracle using the JDBC driver and to verify the connection information.

See "Verifying a JDBC Client Installation" in the *Oracle JDBC Developer's Guide and Reference Guide*.

5. Obtain a local copy of the Oracle `timezone` file, so Replication Agent can correctly process the Oracle `timestamp with timezone datatype`.

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

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

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

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

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

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

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

The majority of configuration values required to create and initialize a Replication Agent can be recorded and stored in a resource file. Using a resource file provides a means to record or retain the configuration information for a Replication Agent instance, allowing an instance to be removed and re-created.

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

7. Create an instance resource file:

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

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

- Set the value of `rs_source_ds` to the value you chose for `pds`.

- Set the value of **rs_source_db** to the value you chose for *pdb*.

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

Warning! The **rs_source_ds** and the **rs_source_db** values must match the *{pds.pdb}* values of your Replication Server primary connection name that you configured in step 3, in *Creating a Replication Server Connection to the Primary Database* on page 17.

Table 3. Resource File Parameters

Parameter	Description	Example Values
instance name	Any valid name.	myra
admin_port	Port number that Replication Agent uses.	9030 (if in use, select a different port number)
pds_tns_connection	Connection name found in the <code>tnsnames.ora</code> file which identifies the connection information for the primary database.	ORA102.JDOE_HOST.COM
pds_tns_filename	File name identifying the Oracle <code>tnsnames.ora</code> file to be used to identify the connection information for the primary database.	A valid Oracle <code>tnsnames.ora</code> file. For example: <code>/opt/oracle/network/admin/tnsnames.ora</code>
pds_username	User ID that Replication Agent uses to access primary data server.	RA_USER
pds_password	Password for pds_username .	sybase
rs_host_name	Machine where Replication Server is installed.	jdoe_host1
rs_port_number	Port where Replication Server is installed.	11752
rs_username	Replication Server user with connect source and create object permissions.	SAMPLE_RS_ra
rs_password	Password for rs_username .	SAMPLE_RS_ra_ps

Parameter	Description	Example Values
rs_charset	Character set that Replication Server is using. Note: The value defined for the rs_charset configuration parameter must match the RS_charset value in the Replication Server configuration file, <code>\$(SYBASE)/REP-15_5/install/<server>.cfg</code> (UNIX or Linux) or <code>\$(SYBASE)\REP-15_5\install\<server>.cfg</code> (Windows).	<ul style="list-style-type: none"> Windows: cp850 UNIX: iso_1
rs_source_ds	Valid name representing data server of Oracle primary database.	NY
rs_source_db	Valid name representing Oracle primary database.	NYora92
rssid_host_name	Machine where RSSD resides.	jdoe_host1
rssid_port_number	Port number where RSSD resides.	11751
rssid_database_name	Database name for RSSD.	SAMPLE_RS_ERSSD
rssid_username	Valid user for RSSD.	SAMPLE_RS_RSSD_maint
rssid_password	Password for rssid_username .	SAMPLE_RS_RSSD_maint_ps
pdb_timezone_file	Path of the Oracle timezone file.	/software/oracle/Ora10g/oracore/zoneinfo/timezone.dat
start_instance	Start the instance that was created.	yes
initialize_instance	Initialize the Replication Agent instance.	yes
pdb_include_archives	Enables or disables the use of Oracle archive log files.	USE_DEFAULT
pdb_archive_path	Identifies directory path where Replication Agent expects to find archived Oracle redo log files.	A valid directory path on the machine hosting Replication Agent that points to a location where Oracle puts the archived redo log files.

8. Create and execute the new instance resource file.

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

a) Validate the settings in the resource file:

- On Windows:

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

- On UNIX or Linux:

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

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

- Response-file processing completed

Or

- Response-file processing completed with errors

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

b) Create and configure the Replication Agent instance:

- On Windows:

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

- On UNIX or Linux:

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

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

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

c) Results are returned as either:

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

See the *Replication Agent Administration Guide*.

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

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

- On Windows:

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

- On UNIX or Linux:

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

Execute the RUN file in the background.

Your Replication Agent for Oracle is now running.

See also

- *Verifying the Replication Agent Instance Installation* on page 30
- *Initializing the Replication Agent Instance* on page 31

Verifying the Replication Agent Instance Installation

Verify that the Replication Agent instance has been successfully installed.

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

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

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

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

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

3. Verify the Replication Agent connection to Replication Server by entering:

```
test_connection RS
go
```

This result is returned:

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

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

See the *Replication Agent Administration Guide*.

4. Verify the Replication Agent connection to the primary database:
 - a) Enter:

```
test_connection PDS
go
```

You see:

```
Type Connection
----
```

```
PDS succeeded
(1 row affected)
```

b) If the result indicates a failure:

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

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

See "Testing network connectivity" in the *Replication Agent Administration Guide*.

See also

- *Initializing the Replication Agent Instance* on page 31

Initializing the Replication Agent Instance

Initialize the Replication Agent instance.

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

1. The `ra_admin init` command verifies that the primary Oracle database is correctly configured to provide archived logging and supplemental logging, and that the Oracle user ID used by the Replication Agent has the necessary permissions. In addition to creating objects in the database to support stored procedure replication, the `ra_admin init` command also initializes the Replication Agent System Database (RASD) by reading schema information, and redo log location information from the primary Oracle database.

To initialize the Replication Agent instance, run:

```
ra_admin init
go
```

You see a message indicating that the procedure was successful.

2. Enter:

```
resume
go
```

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

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

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

```
1> resume
2> go
Msg 32000, Level 20, State 0:
Server 'rao', Procedure 'resume', Line 1:
```

Oracle Replication Setup

```
Command <resume> failed - Desired state <REPLICATE>
could not be achieved. Current state: <REPLICATION DOWN>
```

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

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

```
ra_status
go
```

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

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

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

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

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

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

Note: Be aware that:

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

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

```
Spid Name      State          Info
-----
-----
13  DSI EXEC
Awaiting Command 101(1) SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
9  DSI           Awaiting Message 101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
8  SQM
Awaiting Message 101:0 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD
54  DSI EXEC     Awaiting Command 102(1) orcl1.oratest2
53  DSI           Awaiting Message 102 orcl1.oratest2
17  SQM           Awaiting Message 102:0 orcl1.oratest2DSI
20  DSI EXEC     Suspended      103(1) NY.NYora92
21  DSI           Suspended      103 NY.NYora92
24  DIST         Awaiting Wakeup 103 NY.NYora92
25  SQT          Awaiting Wakeup 103:1 DIST NY.NYora92
23  SQM          Awaiting Message 103:1 NY.NYora92
```

```

22 SQM      Awaiting Message 103:0 NY.NYora92
62 REP AGENT Awaiting Command NY.NYora92

```

Testing Replication

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

1. Connect to the primary Oracle instance as a regular user, rather than as the maintenance user. Make sure the regular user also exists in the replicate database.

2. Create a test table to replicate in the primary database:

```
create table T1 (a int, b char(10));
```

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

```
grant all on T1 to public;
```

4. Connect to the replicate Oracle instance as a regular user, rather than as the maintenance user.

5. Create a test table in the replicate database:

```
create table T1 (a int, b char(10));
```

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

```
grant all on T1 to public;
```

7. Connect to Replication Agent, and mark the table:

```
pdb_setreptable T1, mark
```

8. After the table is marked and the Replication Agent is in a replicating state, insert test data into the test table and commit it.

By default, the work performed by the maintenance user of the primary connection is not replicated. The user ID used to insert transactions cannot be the same as the maintenance user defined in the primary connection.

Replication Agent applies only committed transactions to a replicate database. However, because Oracle expects a commit command to be issued, follow simple test commands with an explicit **commit** command:

```
insert into T1 values ( 42, 'foo' );
commit;
```

9. Use your preferred query tool to examine the replicate site for results and compare the contents of your test table from both the primary and replicate sites.

Resetting the Primary Oracle Database for Replication

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

1. Archive all current redo log data:

Obtaining Help and Additional Information

```
ra_locator move_truncpt
```

2. Reinitialize Replication Agent, which refreshes the Replication Agent repository rather than overwriting it:

```
ra_admin refresh
```

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

3. Reset the locator stored in Replication Server:

```
isql -USAMPLE_RS_RSSD_prim -PSAMPLE_RS_RSSD_prim_ps  
-SSAMPLE_RS_ERSSD -DSAMPLE_RS_ERSSD  
rs_zeroltm NY, NYora92  
go
```

You see:

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

Obtaining Help and Additional Information

Use the Sybase Getting Started CD, Product Documentation site, and online help to learn more about this product release.

- The Getting Started CD (or download) – contains release bulletins and installation guides in PDF format, and may contain other documents or updated information.
- Product Documentation at <http://sybooks.sybase.com/> – is an online version of Sybase documentation that you can access using a standard Web browser. You can browse documents online, or download them as PDFs. In addition to product documentation, the Web site also has links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, Community Forums/Newsgroups, and other resources.
- Online help in the product, if available.

To read or print PDF documents, you need Adobe Acrobat Reader, which is available as a free download from the *Adobe* Web site.

Note: A more recent release bulletin, with critical product or document information added after the product release, may be available from the Product Documentation Web site.

Technical Support

Get support for Sybase products.

If your organization has purchased a support contract for this product, then one or more of your colleagues is designated as an authorized support contact. If you have any questions, or if

you need assistance during the installation process, ask a designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

Downloading Sybase EBFs and Maintenance Reports

Get EBFs and maintenance reports from the Sybase Web site or the SAP® Service Marketplace (SMP). The location you use depends on how you purchased the product.

- If you purchased the product directly from Sybase or from an authorized Sybase reseller:

- a) Point your Web browser to <http://www.sybase.com/support>.
- b) Select **Support > EBFs/Maintenance**.
- c) If prompted, enter your MySybase user name and password.
- d) (Optional) Select a filter, a time frame, or both, and click **Go**.
- e) Select a product.

Padlock icons indicate that you do not have download authorization for certain EBF/Maintenance releases because you are not registered as an authorized support contact. If you have not registered, but have valid information provided by your Sybase representative or through your support contract, click **My Account** to add the “Technical Support Contact” role to your MySybase profile.

- f) Click the **Info** icon to display the EBF/Maintenance report, or click the product description to download the software.
- If you ordered your Sybase product under an SAP contract:
 - a) Point your browser to <http://service.sap.com/swdc> and log in if prompted.
 - b) Select **Search for Software Downloads** and enter the name of your product. Click **Search**.

Sybase Product and Component Certifications

Certification reports verify Sybase product performance on a particular platform.

To find the latest information about certifications:

- For partner product certifications, go to http://www.sybase.com/detail_list?id=9784
- For platform certifications, go to <http://certification.sybase.com/ucr/search.do>

Creating a MySybase Profile

MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

1. Go to <http://www.sybase.com/mysybase>.

2. Click **Register Now**.

Accessibility Features

Accessibility ensures access to electronic information for all users, including those with disabilities.

Documentation for Sybase products is available in an HTML version that is designed for accessibility.

Vision impaired users can navigate through the online document with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Sybase HTML documentation has been tested for compliance with accessibility requirements of Section 508 of the U.S Rehabilitation Act. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

Note: You may need to configure your accessibility tool for optimal use. Some screen readers pronounce text based on its case; for example, they pronounce ALL UPPERCASE TEXT as initials, and MixedCase Text as words. You might find it helpful to configure your tool to announce syntax conventions. Consult the documentation for your tool.

For information about how Sybase supports accessibility, see the Sybase Accessibility site: <http://www.sybase.com/products/accessibility>. The site includes links to information about Section 508 and W3C standards.

You may find additional information about accessibility features in the product documentation.

Glossary

This glossary describes Replication Server Options terms.

- **Adaptive Server®** – the brand name for Sybase relational database management system (RDBMS) software products.
 - Adaptive Server® Enterprise manages multiple, large relational databases for high-volume online transaction processing (OLTP) systems and client applications.
 - Sybase®IQ manages multiple, large relational databases with special indexing algorithms to support high-speed, high-volume business intelligence, decision support, and reporting client applications.
 - SQL Anywhere® (formerly Adaptive Server Anywhere) manages relational databases with a small DBMS footprint, which is ideal for embedded applications and mobile device applications.

See also *DBMS* and *RDBMS*.

- **atomic materialization** – a materialization method that copies subscription data from a primary database to a replicate database in a single, atomic operation. No changes to primary data are allowed until the subscription data is captured at the primary database. See also *bulk materialization* and *nonatomic materialization*.
- **BCP utility** – a bulk copy transfer utility that provides the ability to load multiple rows of data into a table in a target database. See also *bulk copy*.
- **bulk copy** – an Open Client™ interface for the high-speed transfer of data between a database table and program variables. Bulk copying provides an alternative to using SQL **insert** and **select** commands to transfer data.
- **bulk materialization** – a materialization method whereby subscription data in a replicate database is initialized outside of the replication system. You can use bulk materialization for subscriptions to table replication definitions or function replication definitions. See also *atomic materialization* and *nonatomic materialization*.
- **client** – in client/server systems, the part of the system that sends requests to servers and processes the results of those requests. See also *client application*.
- **client application** – software that is responsible for the user interface, including menus, data entry screens, and report formats. See also *client*.
- **commit** – an instruction to the DBMS to make permanent the changes requested in a transaction. See also *transaction*. Contrast with *rollback*.
- **database** – a collection of data with a specific structure (or schema) for accepting, storing, and providing data for users. See also *data server*, *DBMS*, and *RDBMS*.
- **database connection** – a connection that allows Replication Server to manage the database and distribute transactions to the database. Each database in a replication system can have only one database connection in Replication Server. See also *Replication Server* and *route*.
- **data client** – a client application that provides access to data by connecting to a data server. See also *client*, *client application*, and *data server*.
- **data distribution** – a method of locating (or placing) discrete parts of a single set of data in multiple systems or at multiple sites. Data distribution is distinct from data replication, although a data replication system can be used to implement or support data distribution. Contrast with *data replication*.
- **data replication** – the process of copying primary data to remote locations and synchronizing the copied data with the primary data. Data replication is different from data distribution. Replicated data is a stored copy of data at one or more remote sites throughout a system, and it is not necessarily distributed data. Contrast with *data distribution*. See also *transaction replication*.
- **data server** – a server that provides the functionality necessary to maintain the physical representation of a table in a database. Data servers are usually database servers, but they can also be any data repository with the interface and functionality a data client requires. See also *client*, *client application*, and *data client*.

- **datatype** – a keyword that identifies the characteristics of stored information on a computer. Some common datatypes are: *char*, *int*, *smallint*, *date*, *time*, *numeric*, and *float*. Different data servers support different datatypes.
- **DBMS** – an abbreviation for database management system, a computer-based system for defining, creating, manipulating, controlling, managing, and using databases. The DBMS can include the user interface for using the database, or it can be a standalone data server system. Compare with *RDBMS*.
- **ERSSD** – an abbreviation for Embedded Replication Server System Database, which manages replication system information for a Replication Server. See also *Replication Server*.
- **failback** – a procedure that restores the normal user and client access to a primary database, after a failover procedure switches access from the primary database to a replicate database. See also *failover*.
- **failover** – a procedure that switches user and client access from a primary database to a replicate database, particularly in the event of a failure that interrupts operations at the primary database, or access to the primary database. Failover is an important fault-tolerance feature for systems that require high availability. See also *failback*.
- **function** – a data server object that represents an operation or set of operations. Replication Server distributes operations to replicate databases as functions. See also *stored procedure*.
- **function string** – a string that Replication Server uses to map a function and its parameters to a data server API. Function strings allow Replication Server to support heterogeneous replication, in which the primary and replicate databases are different types, with different SQL extensions and different command features. See also *function*.
- **gateway** – connectivity software that allows two or more computer systems with different network architectures to communicate.
- **inbound queue** – a stable queue managed by Replication Server to spool messages received from a Replication Agent. See also *outbound queue* and *stable queue*.
- **interfaces file** – a file containing information that Sybase Open Client and Open Server™ applications need to establish connections to other Open Client and Open Server applications. See also *Open Client* and *Open Server*.
- **isql** – an Interactive SQL client application that can connect and communicate with any Sybase Open Server application, including Adaptive Server, Replication Agent, and Replication Server. See also *Open Client* and *Open Server*.
- **Java** – an object-oriented programming language developed by Sun Microsystems. A platform-independent, “write once, run anywhere” programming language.
- **Java VM** – the Java Virtual Machine. The Java VM (or JVM) is the part of the Java Runtime Environment (JRE) that is responsible for interpreting Java byte codes. See also *Java* and *JRE*.
- **JDBC** – an abbreviation for Java Database Connectivity. JDBC is the standard communication protocol for connectivity between Java clients and data servers. See also *data server* and *Java*.

- **JRE** – an abbreviation for Java Runtime Environment. The JRE consists of the Java Virtual Machine (Java VM or JVM), the Java Core Classes, and supporting files. The JRE must be installed on a machine to run Java applications, such as Replication Agent. See also *Java VM*.
- **LAN** – an abbreviation for “local area network,” a computer network located on the user premises and covering a limited geographical area (usually a single site). Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary can be subject to some form of regulation. Contrast with *WAN*.
- **latency** – in transaction replication, the time it takes to replicate a transaction from a primary database to a replicate database. Specifically, latency is the time elapsed between committing an original transaction in the primary database and committing the replicated transaction in the replicate database.

In disk replication, latency is the time elapsed between a disk write operation that changes a block or page on a primary device and the disk write operation that changes the replicated block or page on a replicate device.

See also *transaction replication*.

- **LOB** – an abbreviation for large object, a large collection of data stored as a single entity in a database.
- **Log Reader** – an internal component of Replication Agent that interacts with the primary database to capture transactions for replication. See also *Log Transfer Interface* and *Log Transfer Manager*.
- **Log Transfer Interface** – an internal component of Replication Agent that interacts with Replication Server to forward transactions for distribution to Replication Server. See also *Log Reader* and *Log Transfer Manager*.
- **Log Transfer Language** – the proprietary protocol used between Replication Agent and Replication Server to replicate data from the primary database to Replication Server. See also *Log Reader* and *Log Transfer Interface*.
- **Log Transfer Manager** – an internal component of Replication Agent that interacts with the other Replication Agent internal components to control and coordinate Replication Agent operations. See also *Log Reader* and *Log Transfer Interface*.
- **maintenance user** – a special user login name in the replicate database that Replication Server uses to apply replicated transactions to the database. See also *replicate database* and *Replication Server*.
- **materialization** – the process of copying the data from a primary database to a replicate database, initializing the replicate database so that the replication system can begin replicating transactions. See also *atomic materialization*, *bulk materialization*, and *nonatomic materialization*.
- **Multi-Path Replication**[™] – Replication Server feature that improves performance by enabling parallel paths of data from the source database to the target database. These multiple paths process data independently and are applicable when sets of data can be processed in parallel without transactional consistency requirements between them.

- **nonatomic materialization** – a materialization method that copies subscription data without a lock on the primary database. Changes to primary data are allowed during data transfer, which may cause temporary inconsistencies between the primary and replicate databases. Contrast with *atomic materialization*. See also *bulk materialization*.
- **ODBC** – an abbreviation for Open Database Connectivity, an industry-standard communication protocol for clients connecting to data servers. See also *client*, *data server*, and *JDBC*.
- **Open Client** – a Sybase product that provides customer applications, third-party products, and other Sybase products with the interfaces needed to communicate with Open Server applications. See also *Open Server*.
- **Open Client application** – An application that uses Sybase Open Client libraries to implement Open Client communication protocols. See also *Open Client* and *Open Server*.
- **Open Server** – a Sybase product that provides the tools and interfaces required to create a custom server. See also *Open Client*.
- **Open Server application** – a server application that uses Sybase Open Server libraries to implement Open Server communication protocols. See also *Open Client* and *Open Server*.
- **outbound queue** – a stable queue managed by Replication Server to spool messages to a replicate database. See also *inbound queue*, *replicate database*, and *stable queue*.
- **primary data** – the data source used for replication. Primary data is stored and managed by the primary database. See also *primary database*.
- **primary database** – the database that contains the data to be replicated to another database (the replicate database) through a replication system. The primary database is the source of replicated data in a replication system. Sometimes called the active database. Contrast with *replicate database*. See also *primary data*.
- **primary key** – a column or set of columns that uniquely identifies each row in a table.
- **primary site** – the location or facility at which primary data servers and primary databases are deployed to support normal business operations. Sometimes called the active site or main site. See also *primary database* and *replicate site*.
- **primary table** – a table used as a source for replication. Primary tables are defined in the primary database schema. See also *primary data* and *primary database*.
- **primary transaction** – a transaction that is committed in the primary database and recorded in the primary database transaction log. See also *primary database*, *replicated transaction*, and *transaction log*.
- **quiesce** – to cause a system to go into a state in which further data changes are not allowed. See also *quiescent*.
- **quiescent** – in a replication system, a state in which all updates have been propagated to their destinations. Some Replication Agent and Replication Server commands require that you first quiesce the replication system.

In a database, a state in which all data updates are suspended so that transactions cannot change any data, and the data and log devices are stable.

This term is interchangeable with quiesced and in quiesce. See also *quiesce*.

- **RASD** – an abbreviation for Replication Agent System Database. Information in the RASD is used by the primary database to recognize database structure or schema objects in the transaction log.
- **RCL** – an abbreviation for Replication Command Language, the command language used to manage Replication Server. See also *Replication Server*.
- **RDBMS** – an abbreviation for relational database management system, an application that manages and controls relational databases. Compare with *DBMS*. See also *relational database*.
- **relational database** – a collection of data in which data is viewed as being stored in tables, which consist of columns (data items) and rows (units of information). Relational databases can be accessed by SQL requests. Compare with *database*. See also *SQL*.
- **replicate data** – A set of data that is replicated from a primary database to a replicate database by a replication system. See also *primary database*, *replication system*, and *replicate database*.
- **replicate database** – a database that contains data replicated from another database (the primary database) through a replication system. The replicate database is the database that receives replicated data in a replication system. Contrast with *primary database*. See also *replicate data*, *replicated transaction*, and *replication system*.
- **replicated transaction** – a primary transaction that is replicated from a primary database to a replicate database by a transaction replication system. See also *primary database*, *primary transaction*, *replicate database*, and *transaction replication*.
- **replicate site** – the location or facility at which replicate data servers and replicate databases are deployed to support normal business operations during scheduled downtime at the primary site. Contrast with *primary site*. See also *replicate database*.
- **Replication Agent** – an application that reads a primary database transaction log to acquire information about data-changing transactions in the primary database, processes the log information, and then sends it to a Replication Server for distribution to a replicate database. See also *primary database* and *Replication Server*.
- **replication definition** – a description of a table or stored procedure in a primary database, for which subscriptions can be created. The replication definition, maintained by Replication Server, includes information about the columns to be replicated and the location of the primary table or stored procedure. See also *Replication Server* and *subscription*.
- **Replication Server** – a Sybase software product that provides the infrastructure for a transaction replication system. See also *Replication Agent*.
- **replication system** – a data processing system that replicates data from one location to another. Data can be replicated between separate systems at a single site, or from one or more local systems to one or more remote systems. See also *transaction replication*.
- **rollback** – an instruction to a database to back out of the changes requested in a unit of work (called a transaction). Contrast with *commit*. See also *transaction*.
- **route** – A one-way message stream from a primary Replication Server to a replicate Replication Server. Routes carry data-changing commands (including those for RSSDs)

and replicated functions (database procedures) between separate Replication Servers. See also *Replication Server*.

- **RSSD** – an abbreviation for Replication Server System Database, which manages replication system information for a Replication Server. See also *Replication Server*.
- **SQL** – an abbreviation for Structured Query Language, a nonprocedural programming language used to process data in a relational database. ANSI SQL is an industry standard. See also *transaction*.
- **stable queue** – a disk device-based, store-and-forward queue managed by Replication Server. Messages written into the stable queue remain there until they can be delivered to the appropriate process or replicate database. Replication Server provides a stable queue for both incoming messages (the inbound queue) and outgoing messages (the outbound queue). See also *database connection*, *Replication Server*, and *route*.
- **stored procedure** – a data server object that represents an operation or set of operations. This term is often used interchangeably with *function*.
- **subscription** – a request for Replication Server to maintain a replicated copy of a table, or a set of rows from a table, in a replicate database at a specified location. See also *replicate database*, *replication definition*, and *Replication Server*.
- **table** – in a relational DBMS, a two-dimensional array of data or a named data object that contains a specific number of unordered rows composed of a group of columns that are specific for the table. See also *database*.
- **transaction** – a unit of work in a database that can include zero, one, or many operations (including **insert**, **update**, and **delete** operations), and that is either applied or rejected as a whole. Each SQL statement that modifies data can be treated as a separate transaction, if the database is so configured. See also *SQL*.
- **transactional consistency** – A condition in which all transactions in the primary database are applied in the replicate database, and in the same order that they were applied in the primary database.
- **transaction log** – generally, the log of transactions that affect the data managed by a data server. Replication Agent reads the transaction log to identify and acquire the transactions to be replicated from the primary database. See also *Replication Agent*, *primary database*, and *Replication Server*.
- **transaction replication** – a data replication method that copies data-changing operations from a primary database to a replicate database. See also *data replication*.
- **UDB** – IBM DB2 Universal Database (formerly IBM DB2 for Linux, UNIX, and Windows).
- **WAN** – an abbreviation for “wide area network,” a system of local-area networks (LANs) connected together with data communication lines. Contrast with *LAN*.

Index

A

- accessing and downloading from SPDC 5
- Adaptive Server Enterprise
 - verifying Replication Server installation 11

C

- configuring
 - Primary Oracle database for replication 22
 - Replication Agent for Oracle 22
 - replication components 16
 - Replication Server for Oracle 16
- configuring ExpressConnect for Oracle 20
- configuring Replication Server
 - for replication from the primary database 16
 - for replication to standby 16
- conventions
 - style 1
 - syntax 1
- creating
 - database replication definitions 19
 - database replication subscriptions 20
 - Oracle user 23
 - Replication Agent instance 25
 - Replication Server connection to standby
 - database 16
 - Replication Server from primary database 17

E

- errors
 - while starting installer 7
- ExpressConnect for Oracle
 - configuring 20

I

- initializing
 - Replication Agent instance 31
- installation
 - Replication Server CD, mounting 8
- installation guidelines
 - for Replication Agent 13
 - for Replication Server 7

- installation program
 - Replication Agent 13
- installer
 - errors in starting 7
- installing Oracle Instant Client libraries 11

L

- LogMiner
 - configuration 23

M

- mount commands by platform 7
- mounting the Replication Server installation media
 - 8

O

- Oracle
 - configuring Replication Server 16
 - configuring Replication Server for replication
 - from the primary database 17
 - installing replication components 6
 - installing Replication Server 7
 - system requirements 4
 - testing replication 33
 - verifying the Replication Agent instance 30
- Oracle Instant Client libraries
 - installing 11

P

- planning your installation
 - replication components 4
- platform
 - mount commands 7
- prerequisites
 - Oracle 5

R

- replication
 - components 3

Index

- overview 3
- Replication Agent
 - installing 13
- Replication Agent for Oracle
 - configuring 22
- Replication Agent installation directory 5
- resetting primary Oracle database for replication 33

S

- setting up replication for Oracle 3
- setup prerequisites 4
- system requirements 4

T

- testing
 - Oracle replication 33

V

- verifying
 - current archive setting of the redo logs 22
 - installation 15
 - primary Oracle configuration for replication 24
 - Replication Agent instance in Oracle 30
 - supplemental logging of primary key data 22