



Quick Start Guide

Mirror Activator™

15.1

[Linux, Microsoft Windows, and UNIX]

DOCUMENT ID: DC00714-01-1510-01

LAST REVISED: September 2008

Copyright © 2008 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.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. 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.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. 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

About This Book	v
CHAPTER 1 Setting Up Replication for ASE.....	1
System configuration.....	1
Prerequisites	2
System requirements	3
Installing Mirror Activator components	3
Before you start	3
Installing Mirror Replication Agent.....	5
Installing Replication Server	8
Configuring Mirror Activator components for ASE.....	11
Configuring Replication Server.....	11
Configuring Mirror Replication Agent	17
CHAPTER 2 Setting Up Replication for Microsoft SQL Server	27
System configuration.....	27
Prerequisites	28
System requirements	29
Installing Mirror Activator components	30
Before you start	30
Installing Mirror Replication Agent.....	32
Installing Replication Server	35
Installing the ECDA Option for ODBC	38
Configuring Mirror Activator components.....	40
Configuring ECDA Option for ODBC.....	40
Configuring Replication Server.....	48
Configuring Mirror Replication Agent	54
Testing replication	70
Resetting the primary Microsoft SQL Server database for replication	71
CHAPTER 3 Setting Up Replication for Oracle.....	73
System configuration.....	73

Prerequisites	74
System requirements	75
Installing Mirror Activator components	76
Before you start	76
Installing Mirror Replication Agent.....	77
Installing Replication Server.....	80
Installing the ECDA Option for Oracle.....	84
Configuring Mirror Activator components.....	86
Configuring ECDA Option for Oracle.....	86
Configuring Replication Server.....	90
Configuring Mirror Replication Agent	96
Testing replication	111
Resetting the primary Oracle database for replication	112
 Glossary	 115
 Index	 123

About This Book

Sybase® Mirror Replication Agent™ version 15.1 extends the capabilities of Replication Server® to support the following primary data servers in a Sybase replication system:

- Adaptive Server® Enterprise (ASE)
- Microsoft SQL Server
- Oracle Database Server

Audience

This book is for anyone who needs to manage or administer a Sybase replication system with non-Sybase primary databases, or administer the non-Sybase primary databases in a Sybase replication system. This may include:

- Database Administrators
- Network Administrators
- System Administrators

How to use this book

Use the Mirror Replication Agent *Reference Manual* to look up detailed information about Mirror Replication Agent commands and configuration parameters.

This book is organized as follows:

Chapter 1, “Setting Up Replication for ASE,” describes the quick-start procedure to install and configure a sample replication environment for ASE.

Chapter 2, “Setting Up Replication for Microsoft SQL Server,” describes the quick-start procedure to install and configure a sample replication environment for Microsoft SQL Server.

Chapter 3, “Setting Up Replication for Oracle,” describes the quick-start procedure to install and configure a sample replication environment for Oracle.

Related documents

Mirror Replication Agent Refer to these documents to learn more about Mirror Replication Agent:

-
- Mirror Activator *Administration Guide* – for an overview of Mirror Replication Agent, information about configuring and administering Mirror Replication Agent instances, and information about configuring the other components in a Mirror Replication Agent system
 - Mirror Replication Agent *Primary Database Guide* – for detailed, database-specific information about each database that is supported by Mirror Replication Agent
 - Mirror Replication Agent *Installation Guide* – for information about installing Mirror Replication Agent software
 - Mirror Replication Agent *Reference Manual* – for information about all Mirror Replication Agent commands and configuration parameters, including syntax, examples, and detailed command usage notes
 - Mirror Replication Agent *Release Bulletin* – for last-minute information that was too late to be included in the documents

Note A more recent version of the Mirror Replication Agent release bulletin may be available on the World Wide Web. To check for critical product or document information that was added after the release of the product CD, use the Sybase Product Manuals Web site.

Replication Server Refer to these documents for more information about transaction replication systems and the Replication Server software:

- Replication Server *Design Guide* – for an introduction to basic transaction replication concepts and Sybase replication technology
- Replication Server *Administration Guide* – for guidelines for creating and managing a replication system
- Replication Server *Installation Guide* – for installation and upgrade procedures for all Replication Server and related products
- Replication Server *Reference Manual* – for the syntax and detailed descriptions of Replication Server commands
- Replication Server *Heterogeneous Replication Guide* – for detailed information about configuring Replication Server and implementing a Sybase replication system with Oracle databases

Enterprise Connect Data Access (ECDA) Refer to these documents for more information about ECDA:

- Enterprise Connect Data Access Option for ODBC *Users Guide for Access Services* – for information about the features and functionality of the ECDA Option for ODBC
- Enterprise Connect Data Access Option for Oracle *Server Administration and Users Guide* – for information about the features and functionality of the ECDA Option for Oracle
- Enterprise Connect Data Access *Installation Guide* – for installation and upgrade procedures for all ECDA and related products

Primary data server Make sure that you have appropriate documentation for the primary data server that you use with the Sybase replication system.

Java environment Mirror Replication Agent requires a Java Runtime Environment (JRE) on the Mirror Replication Agent host machine.

- The Mirror Replication Agent *Release Bulletin* contains the most up-to-date information about Java and JRE requirements.
- Java documentation available from your operating system vendor describes how to set up and manage the Java environment on your platform.

Other sources of information

Use the Sybase Getting Started CD, the SyBooks™ CD, and the Sybase Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the SyBooks CD. It is included with your software. To read or print documents on the Getting Started CD, you need Adobe Acrobat Reader, which you can download at no charge from the Adobe Web site using a link provided on the CD.
- The SyBooks CD contains product manuals and is included with your software. The Eclipse-based SyBooks browser allows you to access the manuals in an easy-to-use, HTML-based format.

Some documentation may be provided in PDF format, which you can access through the PDF directory on the SyBooks CD. To read or print the PDF files, you need Adobe Acrobat Reader.

Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.

-
- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Sybase Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click Certification Report.
- 3 In the Certification Report filter select a product, platform, and timeframe and then click Go.
- 4 Click a Certification Report title to display the report.

❖ Finding the latest information on component certifications

- 1 Point your Web browser to Availability and Certification Reports at <http://certification.sybase.com/>.
- 2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.
- 3 Select Search to display the availability and certification report for the selection.

❖ Creating a personalized view of the Sybase Web site (including support pages)

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

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

Sybase EBFs and software maintenance

❖ Finding the latest information on EBFs and software maintenance

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.
- 3 Select a product.
- 4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

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

- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Style conventions

The following style conventions are used in this book:

- In a sample screen display, commands that you should enter exactly as shown appear like this:

```
ra_init
```

- In the regular text of this document, variables or user-supplied words appear like this:

Specify the *value* option to change the setting of the configuration parameter.

- In a sample screen display, variables or words that you should replace with the appropriate value for your site appear like this:

```
resume connection to pds.pdb
```

Here, *pds* and *pdb* are the variables you should replace.

- In the regular text of this document, names of programs, utilities, procedures, and commands appear like this:

Use the `pdb_init` command to initialize the primary database.

- In the regular text of this document, names of database objects (tables, columns, stored procedures, etc.) appear like this:
Check the price column in the widgets table.
- In the regular text of this document, names of datatypes appear like this:
Use the date or datetime datatype.
- In the regular text of this document, names of files and directories appear like this:
Log files are located in the `$SYBASE/MA-15_1/inst_name/log` directory.

Syntax conventions

The following syntax conventions are used in this book:

Table 1: 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.

In reference sections of this document, statements that show the syntax of commands appear like this:

```
ra_config [param [, value]]
```

The words *param* and *value* in the syntax are variables or user-supplied words.

Character case conventions

The following character case conventions are used in this book:

- All command syntax and command examples are shown in lowercase. However, Mirror Replication Agent 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 Mirror Replication Agent commands. However, if you need to use a mixed-case object name in a Mirror Replication Agent command (to match a mixed-case object name in the primary database), you must delimit the object name with quote characters. For example:

```
pdb_get_tables "TableName"
```

Accessibility features

This document is available in an HTML version that is specialized for accessibility. You can navigate the HTML with an adaptive technology such as a screen reader, or view it with a screen enlarger.

Mirror Replication Agent and the HTML documentation have been tested for compliance with U.S. government Section 508 Accessibility requirements. 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.

The online help for this product is also provided in HTML, which you can navigate using a screen reader.

Note You might 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 Sybase Accessibility at <http://www.sybase.com/accessibility>. The Sybase Accessibility site includes links to information on Section 508 and W3C standards.

For a Section 508 compliance statement for Mirror Activator, see Sybase Accessibility at http://www.sybase.com/detail_list?id=52484.

If you need help

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.

Setting Up Replication for ASE

Topic	Page
System configuration	1
Prerequisites	2
Installing Mirror Activator components	3
Configuring Mirror Activator components for ASE	11

Follow the instructions in this chapter to install and configure a *sample* replication environment for Adaptive Server Enterprise (ASE).

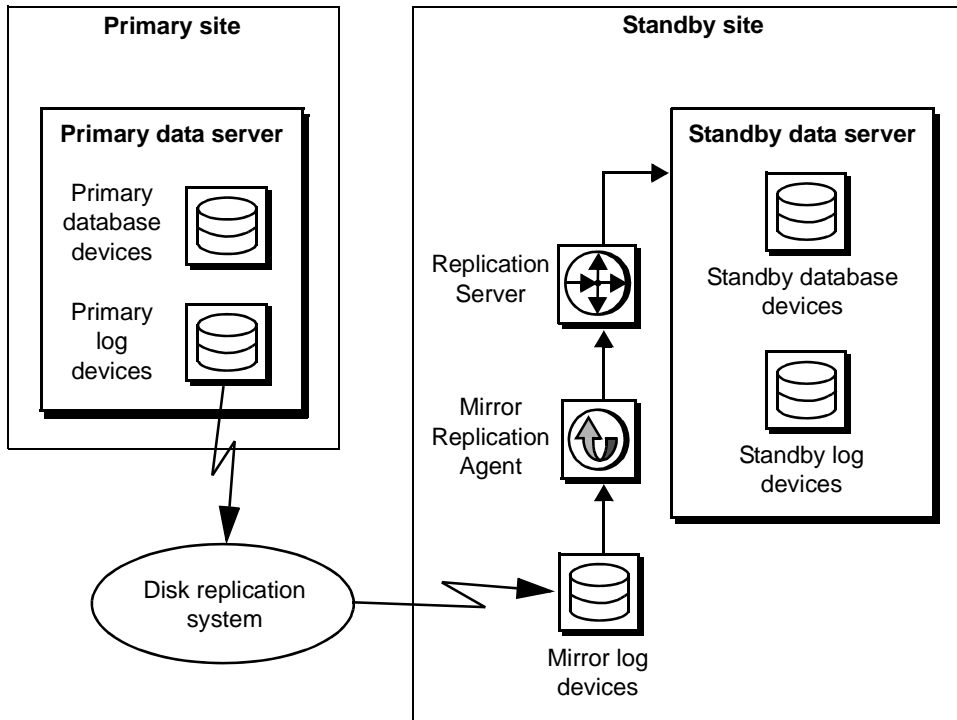
The procedure described in this chapter does **not** describe the steps for creating a mirrored copy.

Warning! This procedure is for Proof of Concept (POC) or test—*not* for production.

System configuration

Figure 1-1 shows a typical Mirror Activator system configuration for ASE. Arrows illustrate the flow of mirrored log device data and replicated transactions during normal Mirror Activator system operation, that is, while replicating transactions from the primary database to the standby database.

Figure 1-1: Sybase replication system with ASE



Prerequisites

Before using this procedure, you should be familiar with ASE and have an understanding of Sybase replication. In addition, these conditions must be met:

- The ASE primary and standby servers are installed and running.
- Mirror Replication Agent has direct access to the ASE *transaction* logs or a mirrored copy of the *transaction* logs.

During this procedure:

- Sybase recommends that you do not use Replication Server reserved words for object names and connection names. A complete list of reserved words are documented in Chapter 2 of the Replication Server *Reference Manual*.

- Because some directories, files, executable commands, and examples are shown only for Sun Solaris, you must adjust accordingly for Microsoft Windows and for other UNIX platforms.

System requirements

The target Windows computer on which you are installing the Mirror Activator components must meet the following memory and disk space requirements:

Table 1-1: Memory and disk requirements

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

Installing Mirror Activator components

This section explains pre-installation tasks and describes how to install and verify the installation of these components of Mirror Activator for ASE:

- Mirror Replication Agent
- Replication Server

Before you start

Before installing Mirror Activator components, identify the Sybase installation directories and obtain the installation software.

Identifying the Sybase installation directory

Mirror Replication Agent and Replication Server are installed in the same base directory (*\$SYBASE*), 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. For example:

```
/software/sybase
```

In addition, make sure you have a temporary installation directory with at least 100MB of space for the installation program to use.

Accessing products from the SPDC Web site

If you are not using installation CDs, you must go to the Sybase product Download Center (SPCD) Web site to obtain replication software.

❖ To download the replication software

Note This procedure requires you to repeat steps for downloading two products: Replication Server and Mirror Activator.

You must have a Login ID and Password to download software from the SPDC Web site.

- 1 Go to SPDC at <https://sybase.subscribenet.com/control/sybs/login> and enter your Login ID and Password.

The Product List page displays.

- 2 Select Mirror Activator.

The detailed Product Information page displays.

- 3 Select the Mirror Activator software you need by primary database, version, and platform.

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

- 4 Either:

- Select I Agree to accept the terms without reading them, 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 displays. 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 go back to the high-level Product Information page.

- 5 The Product Download page lists the software and documentation for the product. Select FTP Download for the Mirror Activator 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 SyBooks download contains the rest of the documentation set.

- 6 Uncompress and extract all the installation images to your local drive.

Installing Mirror Replication Agent

This section describes how to install Mirror Replication Agent. It also has a procedure to verify the installation.

The Mirror Replication Agent installation process creates a directory structure containing the Mirror Replication Agent binaries, scripts, Java Runtime Environment (JRE), and associated files. There are no configuration steps during installation.

Before you install Mirror Replication Agent, see the Mirror Replication Agent *Installation Guide* and complete the installation worksheet.

❖ To install Mirror Replication Agent software

- 1 Verify that the operating system (OS) patch levels are at the current level required to support Java 5.0. This is needed to support the JRE that is installed with Mirror Replication Agent.

For more information about this requirement, see the Mirror Replication Agent *Installation Guide*.

- 2 Use **one** of the following methods to install the software:
 - From the Sybase Software Product Download Center (SPDC), download and extract the Mirror Activator 15.1 CD install image for Sun Solaris (see “Accessing products from the SPDC Web site” on page 4 for details), or
 - Insert the Mirror Activator CD into your CD-ROM drive.
- 3 Launch the installation program by executing the setup routine for Solaris:

```
./setup
```

- 4 If you see an error message saying that there is not enough temporary disk space, add this to your `setup` command:

```
./setup -is:tempdir someDirectory
```

where *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.

Note If you are installing from a remote machine and do not have an X-windows client, execute `./setup --console` to run in console mode.

- 5 On the Welcome window, click Next.
- 6 Select a geographic location that displays the license agreement. Select I Agree and click Next.
- 7 Enter the directory where you want to install Mirror Replication Agent:
- If the directory does not exist, the installation program prompts you to create it. Click Yes to create it.
 - If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 8 Select Typical for the installation setup type and click Next.
- 9 The installation program displays a summary of what is to be installed. Click Next.
- 10 The installation program starts to install Mirror Replication Agent, displaying the progress status. You might see this message:

```
There are newer files already installed and do you
want to replace them?
```

Click No to All.

- 11 When you see a message that indicates that the installation was successful, click Next.
- 12 The SySAM License Server window opens and displays this prompt:
- ```
Will licenses be obtained from the License Server?
```
- If you answer Yes, you can leave the *license-mrx.lic* license file where it was installed (`$SYBASE/SYSAM-2_0/licenses`), or you can move the file to the *SYSAM-2\_0/licenses* directory of your SYSAM license server. In either case, the Mirror Replication Agent will verify the necessary license.

- If you answer No, no further action is needed. The Mirror Replication Agent will use the installed license (*license-mrx.lic*) as it was installed in *\$SYBASE/SYSAM-2\_0/licenses*.

Click Next.

- 13 The SySAM notification window prompts you to 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 these values from the installation worksheet you completed in the Mirror Replication Agent *Installation Guide*:

- SMTP server host name
- SMTP server port number
- E-mail return address
- Recipients of the notification
- Message severity level of an event that will trigger e-mail. Your choices are:
  - Informational
  - Warning
  - Error

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

Click Next.

- 14 When you receive a message indicating that Mirror Replication Agent was successfully installed, click Finish.

For more information about the installation process and the installation program, see the Mirror Replication Agent *Installation Guide*.

❖ **To verify the installation**

- 1 Change to the directory (*\$SYBASE*) where you installed Mirror Replication Agent.
- 2 Set the environment variable by sourcing *SYBASE.csh* or *SYBASE.sh*, as appropriate for your shell.
- 3 Change the directory to *\$SYBASE/MA-15\_1/bin*.
- 4 Obtain the Mirror Replication Agent version string:

```
ma.sh -v
```

You should see the Sybase copyright and the Mirror Replication Agent version string.

## Installing Replication Server

This section describes how to install Replication Server. It also has a procedure to verify the installation.

---

**Note** If Replication Server is already installed, go to the next section called “Configuring Replication Server for replication to standby.”

---

### ❖ To install Replication Server

- 1 Use *one* of the following methods to download Replication Server software.
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 4, go to the location where you extracted the Replication Server 15.1 for Windows installation image (see “Accessing products from the SPDC Web site” on page 4 for details), or
  - Insert the Replication Server CD in your CD-ROM drive.
- 2 Launch the installation program by executing the setup routine for Solaris:  

```
./setup
```
- 3 If you see an error message saying that there is not enough temporary disk space, add this line to your setup command:  

```
./setup -is:tempdir someDirectory
```

where *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.
- 4 On the Welcome window, click Next.
- 5 Select a geographic location that displays the license agreement. Select I Agree and click Next.
- 6 Enter the directory where you want to install Replication Server:
  - If the directory does not exist, the installation program prompts you to create it. Click Yes to create it.

- If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 7 Select Typical for the installation setup type and click Next.
  - 8 The installation program displays a summary of what is to be installed. Click Next.
  - 9 The installation program starts to install Replication Server, displaying the progress status. You may see the following message:

```
There are newer files already installed and do you
want to replace them?
```

Click No to All.
  - 10 When you receive a message that Replication Server software has successfully installed, click Next.

---

**Note** This procedure uses the SAMPLE\_RS Replication Server.

---

- 11 The SySAM License Server window opens and displays this prompt:

```
Will licenses be obtained from the License Server?
```

Respond to how you set up your environment and use the license obtained from the software product download site (SPDC) for Mirror Activator. Click Next.
- 12 The SySAM notification window prompts you to 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 your own values for the following fields:
  - SMTP server host name
  - SMTP server port number
  - Sender e-mail
  - Recipient e-mails
  - Message severity level of an event that will trigger e-mail. Your choices are:
    - Informational
    - Warning

- Error

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

Click Next.

13 To start a sample Replication Server, select Yes and click Next.

14 The following information is displayed:

This page contains detailed information regarding the sample Replication Server. Please record this information.

The sample Replication Server will be called SAMPLE\_RS and will run on port 11752. It will be configured with a user of *sa* and no password.

The sample Replication Server will use an embedded RSSD called SAMPLE\_RS\_ERSSD that runs on port 11751. It will be configured with a user of SAMPLE\_RS\_RSSD\_prim and a password of SAMPLE\_RS\_RSSD\_prim\_ps.

The installer has updated the appropriate *interfaces* file or *sql.ini* file.

All files and logs associated with the sample Replication Server will be located in the directory *\$SYBASE/REP-15\_1/samp\_repserver*.

The sample Replication Server will be configured using the file *\$SYBASE/REP-15\_1/samp\_repserver/SAMPLE\_RS.res*.

Keep this information for future reference. Click Next.

15 The SAMPLE\_RS is now running and a “successful installation” message appears. Click Finish.

❖ **To verify the installation**

- 1 Change to the directory (*\$SYBASE*) where you installed Replication Server.
- 2 Set the environment variables by sourcing the *SYBASE.csh* file.
- 3 Use *isql* to log in to Replication Server:

```
isql -Usa -P -SSAMPLE_RS
```

You should successfully log in to the Replication Server.

- 4 Exit isql.

## Configuring Mirror Activator components for ASE

This section describes the steps, in order, for configuring the components for Mirror Replication Agent for ASE:

- 1 Configuring Replication Server
- 2 Configuring Mirror Replication Agent

### Configuring Replication Server

This section describes the required steps—shown in order—to configure Replication Server for ASE:

- 1 Configure the Replication Server for remote access.
- 2 Configure Replication Server for replication to standby.
- 3 Prepare the primary ASE for Replication Server.
- 4 Configure Replication Server for replication from the primary database.

For more information, see the Mirror Activator *Administration Guide* and the Mirror Replication Agent *Primary Database Guide*.

### Configure the Replication Server for remote access

This section describes the required configurations to allow replication to occur when Replication Server is installed on a different (remote) machine than the primary or standby ASE servers.

The *interfaces* file entry for SAMPLE\_RS and SAMPLE\_RS\_ERSSD uses “localhost” as the name of the machine where SAMPLE\_RS and SAMPLE\_RS\_ERSSD reside. If the primary ASE database is not on the same machine as the SAMPLE\_RS, the *interfaces* file entry must be changed to use either the machine name or the IP address of the machine.

---

**Note** This procedure is not required if the primary and standby ASE servers are on the same machine as Replication Server.

---

❖ **To configure the SAMPLE\_RS for remote access**

- 1 Shut down Replication Server:

```
isql -Usa -P -SSAMPLE_RS
shutdown
go
```

A message displays indicating that Replication Server is shut down.

- 2 Edit the *interfaces* file.

---

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

---

- 3 Change the entry for the SAMPLE\_RS and SAMPLE\_RS\_ERSSD from “localhost” to either a machine name or an IP address.
- 4 Add the primary ASE server and the standby ASE servers to the *interfaces* file, if they do not exist.
- 5 Go to the `$SYBASE/$SYBASE_REP/samp_repserver` directory:

```
cd $SYBASE/$SYBASE_REP/samp_repserver/
```

- 6 Run the start-up script:

```
./RUN_SAMPLE_RS &
```

## Configure Replication Server for replication to standby

This section describes the required configurations for Replication Server for replication to the standby ASE server.

❖ **To create a Replication Server connection to standby**

- 1 If needed, add the standby server to the *interfaces* file of the Replication Server and reboot the Replication Server.
- 2 Go to the `$SYBASE/MA-15_1/scripts/ase` directory.
- 3 Make a copy of the *ase\_create\_rs\_standby\_connection.rs* file:

```
cp ase_create_rs_standby_connection.rs
myma_ase_create_rs_standby_connection.rs
```

- 4 Using a text editor, change these values in the newly-created standby script and save the script:



- `$SYBASE` – enter the path name of the `$SYBASE` where the Replication Server resides, for example. `/software/sybase`. (`$SYBASE` occurs three times in the script.)
- `$SYBASE_REP` – enter `REP-15_1`.
- `rds` – name of the data server where the standby database resides.
- `rdb` – name of the standby database.

---

**Note**

- If you are not using the `SAMPLE_RS`, change the `SAMPLE_RS` to the proper Replication Server name.
- The value entered for `rs.rs_ds_name` and `rs.rs_name` must exist in the *interfaces* file used by the Replication Server. The names are case sensitive.

- 
- 5 Using the `rs_init` utility in Replication Server, create the connection from the Replication Server to the standby database:

```
$SYBASE/$SYBASE_REP/install/rs_init -r
myma_ase_create_rs_standby_connection.rs
```

The connection to the standby database is created.

## Prepare the primary ASE for Replication Server

This section describes how to prepare the primary ASE database for Replication Server.

### ❖ To prepare the primary ASE database for Replication Server

- Add an entry for the Replication Server in the *interfaces* file of the primary ASE.

---

**Note** It is possible that a reboot will be required for the primary ASE. Sybase recommends that you reboot the ASE server to get a fresh copy of the *interfaces* file.

---

## Configure Replication Server for replication from the primary database

This section describes these configuration steps required for replication from the primary database:

- 1 Configure Replication Server for replication from the primary database.
- 2 Grant create object permission to the Mirror Replication Agent user.
- 3 Create the database replication definition.
- 4 Create the database replication subscription.

❖ **To configure Replication Server for replication from the primary database**

- 1 If needed, add the primary ASE server to the *interfaces* file of the Replication Server and reboot the Replication Server.
- 2 Change the directory to `$SYBASE/MA-15_1/scripts/ase`.
- 3 Make a copy of `ase_create_rs_primary_connection.rs`:

```
cp ase_create_rs_primary_connection.rs
 myma_ase_create_rs_primary_connection.rs
```
- 4 Using a text editor, change these values in the newly created script and save the script:
  - `$SYBASE` – enter the path name of the `$SYBASE` where the Replication Server resides, for example, `/software/sybase`. (`$SYBASE` occurs three times in the script).
  - `$SYBASE_REP` – enter `REP-15_1`.
  - `pds` – name of the ASE data server where the primary database resides.
  - `pdb` – name of the primary database.

---

**Note**

- If you are not using the `SAMPLE_RS`, change the `SAMPLE_RS` to the proper Replication Server name.
  - The value entered for `rs.rs_ds_name` and `rs.rs_name` must exist in the *interfaces* file used by the Replication Server, and the names are case sensitive.
- 

- 5 Using the `rs_init` utility in Replication Server, create the connection from the Replication Server to the primary database:

```
$SYBASE/$SYBASE_REP/install/rs_init -r
myma_ase_create_rs_primary_connection.rs
```

The connection to the primary database is created.

- 6 Verify that the connections to the primary and standby servers have been made:

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

| Spid | Name      | State            | Info                                   |
|------|-----------|------------------|----------------------------------------|
| 23   | DSI EXEC  | Awaiting Command | 101(1) SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD |
| 19   | DSI       | Awaiting Message | 101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD    |
| 18   | SQM       | Awaiting Message | 101:0 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD  |
| 28   | DSI EXEC  | Awaiting Command | 102(1) repit151i2.foo2                 |
| 27   | DSI       | Awaiting Message | 102 repit151i2.foo2                    |
| 26   | SQM       | Awaiting Message | 102:0 repit151i2.foo2                  |
| 38   | DSI EXEC  | Awaiting Command | 103(1) repit151i.foo                   |
| 33   | DSI       | Awaiting Message | 103 repit151i.foo                      |
| 35   | DIST      | Awaiting Wakeup  | 103 repit151i.foo                      |
| 36   | SQT       | Awaiting Wakeup  | 103:1 DIST repit151i.foo               |
| 34   | SQM       | Awaiting Wakeup  | 103:1 repit151i.foo                    |
| 32   | SQL       | Awaiting Message | 103:0 repit151i.foo                    |
| 37   | REP AGENT | Awaiting Command | repit151i.foo                          |
| 20   | dSUB      | Sleeping         |                                        |
| 15   | dCM       | Awaiting Message |                                        |
| 17   | dAIO      | Awaiting Message |                                        |
| 21   | dREC      | Sleeping         | dREC                                   |
| 39   | USER      | Active           | sa                                     |
| 14   | dALARM    | Awaiting Wakeup  |                                        |
| 22   | dSYSAM    | Sleeping         |                                        |

In this example, *repit151i.foo* is the primary database, it has a Mirror Replication Agent thread, and *repit151i2.foo2* is the standby database. All threads for the primary and standby databases should have the *Awaiting Message* or *Awaiting Command* state.

❖ **To grant *create object* permission to the Mirror Replication Agent user**

- 1 Allow the default Mirror Replication Agent user to create replication definitions:

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

Permission is granted to user “SAMPLE\_RS\_ra.”

---

**Note** Create object permissions must be granted to the value (located in *\$SYBASE/MA-15\_1/init/ase.rs*) that you use for the *rs\_username* configuration parameter.

---

2 Exit isql.

❖ **To create the database replication definition**

1 Go to the *\$SYBASE/MA-15\_1/scripts/sybase* directory:

```
cd $SYBASE/MA-15_1/scripts/sybase
```

2 Make a copy of the *create\_rs\_db\_repdef.sql* directory:

```
cp create_rs_db_repdef.sql
myma_create_rs_db_repdef.sql
```

3 Before executing the *myma\_create\_rs\_db\_repdef.sql* script, change the value of “{pds}.{pdb}” to the name of the connection string you defined for the primary database, where:

- *pds* is the name of the data server where the primary database resides.
- *pdb* is the name of the primary database.

For example:

```
replit151i.foo
```

Here is a database replication definition example:

```
create database replication definition replit151i_repdef1
with primary at replit151i.foo
replicate DDL
go
```

4 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imyma_create_rs_db_repdef.sql
```

A message appears indicating that the database replication definition is created.

❖ **To create the database replication subscription**

1 Make a copy of *create\_rs\_db\_sub.sql*:

```
cp create_rs_db_sub.sql myma_create_rs_db_sub.sql
```

- 2 Using a text editor, change these values in the copied script and save the changes:

`{pds}` – name of the data server where the primary database resides  
`{pdb}` – name of the primary database  
`{rds}` – name of the data server where the standby database resides  
`{rdb}` – name of the standby database

Here is a create database replication subscription example:

```
create subscription repit151i_sub1
for database replication definition
repit151i_repdef1
with primary at repit151i.foo
with replicate at repit151i2.foo2
without materialization
```

- 3 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imymra_create_rs_db_sub.sql
```

A message appears indicating that the subscription is being created.

## Configuring Mirror Replication Agent

This section describes the required steps—shown in order—to configure Mirror Replication Agent for ASE:

- 1 Prepare and create the Mirror Replication Agent instance.
- 2 Verify the Mirror Replication Agent instance installation.
- 3 Resume the Mirror Replication Agent instance.
- 4 Test replication.
- 5 Reset the primary ASE database for replication.

### Prepare and create the Mirror Replication Agent instance

A single installation of the Mirror Replication Agent can support replication from multiple databases; however, one Mirror Replication Agent instance is needed for each ASE database to be replicated. This section describes how to prepare and create a Mirror Replication Agent for ASE instance:

#### ❖ To prepare for a Mirror Replication Agent instance

- 1 Determine the default settings you want.

By default, Mirror Replication Agent is configured to mark every user table in the primary database for replication, and to create a replication definition for every table marked. Although these settings may be reasonable for production environments, they may not be useful in a test or proof-of-concept environment, where only a subset of tables are being replicated. If you do not want automatic the table marking or replication definition generation options, change the settings of two Mirror Replication Agent configuration parameters: `pdb_automark_tables` and `pdb_auto_create_repdefs`.

For more information about the configuration parameters, refer to the Mirror Replication Agent *Reference Manual*.

- 2 Locate the sample *resource* file.

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

A sample resource file is located in the Mirror Replication Agent `$SYBASE/MA-15_1/init/ase.rs` directory.

❖ **To create a Mirror Replication Agent instance**

- 1 Go to the `$SYBASE/MA-15_1/init` directory.
- 2 Create an instance *resource* file by copying the resource file template `$SYBASE/MA-15_1/init/ase.rs` file to another file that contains the configuration information for a new instance:

```
cp ase.rs myma.rs
```

- 3 Using a text editor, change the values assigned to the properties listed in Table 1-2:

**Table 1-2: Resource file parameters**

| Parameter                      | Description                                               | Example values     |
|--------------------------------|-----------------------------------------------------------|--------------------|
| <code>instance_name</code>     | Any valid name.                                           | <code>myma</code>  |
| <code>admin_port</code>        | Port number that Mirror Replication Agent will use.       | <code>20000</code> |
| <code>pds_host_name</code>     | Machine (host) where ASE is installed.                    | <code>repit</code> |
| <code>pds_port_number</code>   | Port number for ASE.                                      | <code>5600</code>  |
| <code>pds_database_name</code> | Name of database replicated from the primary data server. | <code>foo</code>   |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                       | Example values                                            |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------|
| pds_username        | A valid user name that has “sa” and replicate roles.                                                                                                                                                                                                                                                                              | sa                                                        |
| pds_password        | Password for pds_username.                                                                                                                                                                                                                                                                                                        | sa_ps<br><br><b>Note</b> For null passwords, leave blank. |
| create_pds_username | Specifies whether or not the ma_admin utility, invoked without the -r option, should create the Mirror Replication Agent primary database user login if the login does not already exist in the primary data server. For information about the ma_admin utility -r option, see the Mirror Activator <i>Administration Guide</i> . | yes                                                       |
| pds_sa_username     | System admin name.                                                                                                                                                                                                                                                                                                                | sa                                                        |
| pds_sa_password     | System admin name password.                                                                                                                                                                                                                                                                                                       | sa_ps<br><br><b>Note</b> For null passwords, leave blank. |
| rs_host_name        | Machine where Replication Server is installed.                                                                                                                                                                                                                                                                                    | repit                                                     |
| rs_port_number      | Port where Replication Server is installed.                                                                                                                                                                                                                                                                                       | 11752                                                     |
| rs_username         | User ID that Mirror Replication Agent uses to access primary Replication Server; it must have CONNECT SOURCE and CREATE OBJECT capabilities.                                                                                                                                                                                      | SAMPLE_RS_ra                                              |
| rs_password         | Password for rs_username.                                                                                                                                                                                                                                                                                                         | SAMPLE_RS_ra_ps                                           |
| rs_charset          | Character set that Replication Server is using.<br><br><b>Note</b> The rs_charset value must be set to match the Replication Server character set.                                                                                                                                                                                | iso_1                                                     |
| rs_source_ds        | ASE server of the primary database.                                                                                                                                                                                                                                                                                               | repit151i                                                 |
| rs_source_db        | Primary database.                                                                                                                                                                                                                                                                                                                 | foo                                                       |
| rssd_host_name      | Machine where RSSD resides.                                                                                                                                                                                                                                                                                                       | repit                                                     |
| rssd_port_number    | Port number where RSSD resides.                                                                                                                                                                                                                                                                                                   | 11751                                                     |
| rssd_database_name  | Database name for RSSD.                                                                                                                                                                                                                                                                                                           | SAMPLE_RS_ERSSD                                           |
| rssd_username       | Valid user for RSSD.                                                                                                                                                                                                                                                                                                              | SAMPLE_RS_RSSD_maint                                      |
| rssd_password       | Password for rssd_username.                                                                                                                                                                                                                                                                                                       | SAMPLE_RS_RSSD_maint_ps                                   |
| disable_rat         | Disable Mirror Replication Agent thread.                                                                                                                                                                                                                                                                                          | yes                                                       |
| start_instance      | Start the instance that was created.                                                                                                                                                                                                                                                                                              | yes                                                       |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                 | Example values                                        |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|
| initialize_instance | <p>Initialize the Mirror Replication Agent instance.</p> <hr/> <p><b>Note</b> Set the initialize_instance parameter to yes (the default is no) to initialize the Mirror Replication Agent instance. If you retain the default value of no, you must manually issue the ra_devicepath command after initialization to set the correct mirror path.</p> <hr/> | yes                                                   |
| ra_devicepath_x     | <p>Primary database log mirror location information. These parameter names must be unique in the format ra_devicepath_x, where x is a unique sequential number and the parameter value contains both the device name and the device path separated by a comma.</p>                                                                                          | <p>ra_devicepath_1= logA,<br/>/dev/pubs2/logA.dat</p> |

- 4 Validate and execute the new instance resource file:

```
$SYBASE/MA-15_1/bin/ma_admin.sh -vr myma.rs
```

where *myma.rs* is the path 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 *ma\_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 because of this process.

- 5 After the *resource* file has been validated, allow the *ma\_admin* utility to create the MRA instance:

```
$SYBASE/MA-15_1/bin/ma_admin.sh -r myma.rs
```

A message appears, indicating that response file processing is completed.

- 6 **Note** If in your resource file you set start\_instance to yes and the log indicates that the Mirror Replication Agent instance was started, skip this step, and continue with the next section.

Change to the instance directory and run Mirror Replication Agent in the background:



```
cd $SYBASE/MA-15_1/myma
```

Execute the *RUN* file in the background, for example:

```
./RUN_myma &
```

Your Mirror Replication Agent for ASE is now running.

Using the example  
settings in the  
*resource* file

Using the example settings specified in the *resource* file, the Mirror Replication Agent will:

- Stop, disable, and preserve the secondary truncation point of the current ASE Mirror Replication Agent thread
- Mark the primary database with `sp_reptostandby`

---

**Warning!** If you have a large amount of text or image data in existing tables, the `sp_reptostandby` might take a long time to complete.

---

- Read all the tables in the primary database and prepare them for replication
- Initialize the Mirror Replication Agent

## Verify the Mirror Replication Agent instance installation

This section describes the required steps to verify an Mirror Replication Agent instance:

### ❖ To verify the Mirror Replication Agent installation

- 1 Update the Replication Server *interfaces* file to include an entry for the Mirror Replication Agent location. For details on how to do this, see the Replication Server *Administration Guide*.
- 2 Verify the connection to the Mirror Replication Agent:
  - a Open a command window in the `$SYBASE` directory of your Replication Server installation.
  - b Set the environment variables by sourcing `SYBASE.csh` or `SYBASE.sh`, as appropriate for your shell.
  - c Use `isql` to log in to Mirror Replication Agent:

```
isql -Usa -P -Smyma
```

- 3 Verify the Mirror Replication Agent connection to Replication Server:
  - a Enter this command:

```
test_connection RS
```

go

- b This output appears:

```
Type Connection

RS succeeded
(1 row affected)
```

- c If the result indicates a failure, the Replication Server or the RSSD 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 and RSSD as the configured user to determine which property is incorrectly defined.

For additional information, refer to the Mirror Activator *Administration Guide*, Chapter 2, “Setup and Configuration,” in the section titled, “Testing network connectivity.”

- 4 Verify the Mirror Replication Agent connection to the primary ASE database:

- a Enter this command:

```
test_connection PDS
go
```

This output appears:

```
Type Connection

PDS succeeded
(1 row affected)
```

- b If the result indicates a failure:

- The server is not responding, or
- The connection properties host, port, pds\_database\_name, user, or password are incorrect.

Check the host and port configuration values, and manually log in to the primary ASE database as the configured user (pds\_username in the \$SYBASE/MA-15\_1/init/myma.rs) to find which property is incorrectly defined.

For more information about connection failures, see the Mirror Activator *Administration Guide* section titled, “Testing network connectivity.”

## Resume the Mirror Replication Agent instance

This section describes the commands you must issue to resume a Mirror Replication Agent instance.

### ❖ To resume the Mirror Replication Agent instance

- 1 The resume command puts the Mirror Replication Agent in a *REPLICATING* state, reading the ASE *active* log and sending commands to Replication Server.

To place the Mirror Replication Agent in a *REPLICATING* state, enter the resume command:

```
resume
go
```

If the Mirror Replication Agent successfully transfers to a *REPLICATING* state, this result is returned:

| State                               | Action                   |
|-------------------------------------|--------------------------|
| -----                               | -----                    |
| REPLICATING (WAITING AT END OF LOG) | Ready to replicate data. |

If the state returned is *ADMIN*, an error prevented the *REPLICATING* state from being achieved. To determine the error, review the contents of the Mirror Replication Agent instance log. Also, see the Mirror Activator *Administration Guide*, Chapter 2, “Troubleshooting Mirror Replication Agent,” and the section titled, “Examining the Mirror Replication Agent when replication failure occurs.”

- 2 The `ra_status` command returns the state of the Mirror Replication Agent. It is good practice to verify that the Mirror Replication Agent remains in *REPLICATING* state, even after the resume command executes successfully.

To detect an error that occurred after replication start-up, execute the `ra_status` command:

```
ra_status
go
```

If the Mirror Replication Agent is in *REPLICATING* state, this result is returned:

| State                               | Action                   |
|-------------------------------------|--------------------------|
| -----                               | -----                    |
| REPLICATING (WAITING AT END OF LOG) | Ready to replicate data. |

If the state returned is *ADMIN*, an error prevented the *REPLICATING* state from being achieved. To determine the error, review the contents of the Mirror Replication Agent instance log.

Also, see the Mirror Activator *Administration Guide*, Chapter 2, “Troubleshooting Mirror Replication Agent,” and the section titled, “Examining the Mirror Replication Agent when replication failure occurs.”

- 3 To validate that both primary and standby connections are active, issue the `admin who` command:

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

You should not proceed until `admin who` has a status for threads similar to this:

| Spid | Name      | State            | Info                                   |
|------|-----------|------------------|----------------------------------------|
| 30   | DSI EXEC  | Awaiting Command | 101(1) SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD |
| 22   | DSI       | Awaiting Message | 101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD    |
| 18   | SQM       | Awaiting Message | 101:0 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD  |
| 47   | DSI EXEC  | Awaiting Command | 102(1) repit151i2.foo2                 |
| 23   | DSI       | Awaiting Message | 102 repit151i2.foo2                    |
| 19   | SQM       | Awaiting Message | 102:0 repit151i2.foo2                  |
| 46   | DSI EXEC  | Awaiting Command | 103(1) repit151i.foo                   |
| 24   | DSI       | Awaiting Message | 103 repit151i.foo                      |
| 26   | DIST      | Awaiting Wakeup  | 103 repit151i.foo                      |
| 29   | SQT       | Awaiting Wakeup  | 103:1 DIST repit151i.foo               |
| 21   | SQM       | Awaiting Message | 103:1 repit151i.foo                    |
| 20   | SQM       | Awaiting Message | 103:0 repit151i.foo                    |
| 53   | REP AGENT | Awaiting Command | repit151i.f00                          |
| 25   | dSUB      | Sleeping         |                                        |
| 15   | dCM       | Awaiting Message |                                        |
| 17   | dAIO      | Awaiting Message |                                        |
| 27   | dREC      | Sleeping         | dREC                                   |
| 54   | USER      | Active           | sa                                     |
| 14   | dALARM    | Awaiting Wakeup  |                                        |
| 28   | dSYSAM    | Sleeping         |                                        |

## Test replication

This section describes the steps required to test replication.

**❖ To test replication**

- 1 Connect to the primary ASE database as a user.

---

**Warning!** Sybase recommends that you do *not* use the *maint user* because *maint user* filtering is enabled by default and transactions by this user will not be replicated.

---

- 2 Create a test table to replicate (unless it already exists):

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

After the elapsed time referenced by the Mirror Replication Agent setting in `scan_sleep_max`, the test table named “T1” should be replicated to the target ASE instance.

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

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

- 4 After the table is marked and the Mirror Replication Agent is in a *REPLICATING* state, insert test data into the test table and commit it:

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

go
```

- 5 Using your preferred ASE query tool, examine the standby site for results and compare the contents of your test table from both the primary and standby sites.

What to do if  
replication fails

If replication fails, refer to the Mirror Activator *Administration Guide*, in Chapter 2, “Troubleshooting Mirror Replication Agent,” the section titled, “Examining the Mirror Replication Agent when replication failure occurs.”

**Reset the primary ASE database for replication**

In a test environment, there may be times when the replication environment should be *reset*. Instead of deleting and re-creating a new Mirror Replication Agent instance, use this procedure to easily reset the environment.

**❖ To reset the primary ASE database for replication**

- 1 Protect your new environment from old log information by using this command in the Mirror Replication Agent to move the transaction log truncation point to the end of the log:

```
pdb_init move_truncpt
```

- 2 You can retain marking and log device information by re-initializing the Mirror Replication Agent using `ra_init` with the `force` option. This forces the Mirror Replication Agent repository to be refreshed instead of overwritten. Before and after invoking `ra_init`, you must quiesce and resume the primary database with the `pdb_quiesce` command:

```
pdb_quiesce hold
ra_init force
pdb_quiesce release
```

---

**Note** If you prefer to delete and replace all the information in the Mirror Replication Agent repository, issue the `ra_deinit` command followed by a normal `ra_init` command (without the `force` option):

---

For more information about the `ra_init` command, see the Mirror Replication Agent *Reference Manual*.

- 3 Zero-out the locator stored in Replication Server using the `rs_zeroltm` command, for example:

```
isql -USAMPLE_RS_RSSD_prim -PSAMPLE_RS_RSSD_prim_ps
-SSAMPLE_RS_ERSSD -DSAMPLE_RS_ERSSD

rs_zeroltm repit151i,foo

go
```

This message appears:

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

# Setting Up Replication for Microsoft SQL Server

| Topic                                   | Page |
|-----------------------------------------|------|
| System configuration                    | 27   |
| Prerequisites                           | 28   |
| Installing Mirror Activator components  | 30   |
| Configuring Mirror Activator components | 40   |
| Testing replication                     | 70   |

Follow the instructions in this chapter to install and configure a *sample* replication environment for Microsoft SQL Server.

The procedure described in this chapter does **not** describe the steps for creating a mirrored copy.

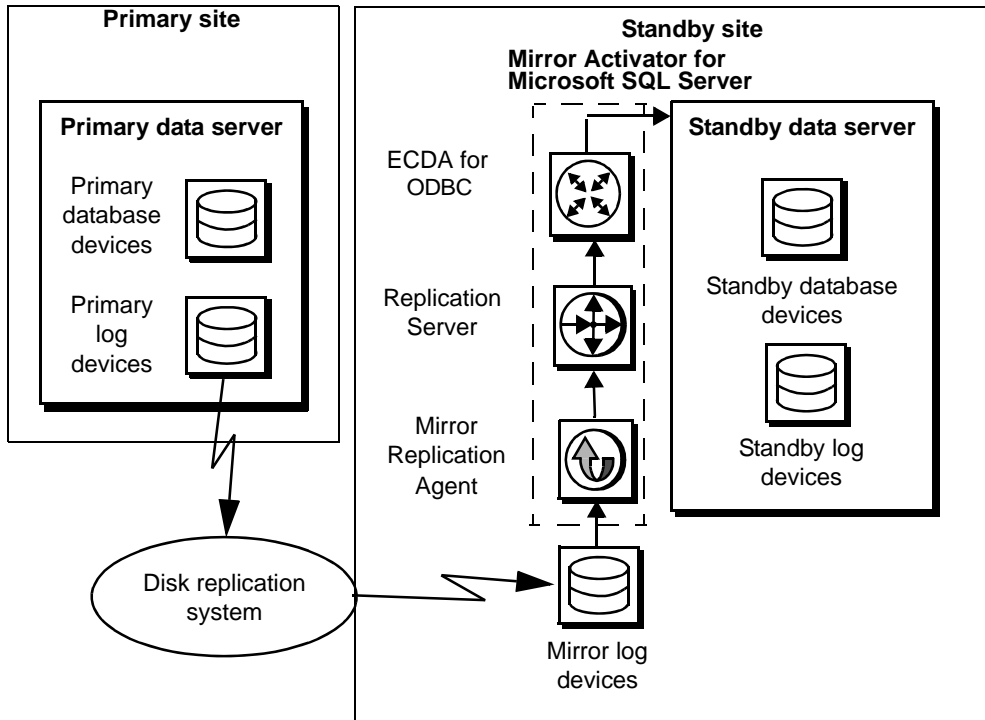
---

**Note** This procedure is only for proof-of-concept (POC) or testing, not for production.

---

## System configuration

Figure 2-1 shows a typical Mirror Replication Agent system configuration for Microsoft SQL Server. Arrows illustrate the flow of mirrored log device data and replicated transactions during normal Mirror Replication Agent system operation, that is, while replicating transactions from the primary database to the standby database.

**Figure 2-1: Sybase replication system for Microsoft SQL Server**

## Prerequisites

Before you begin setting up the replication system, you should be familiar with the Microsoft SQL Server data server and have an understanding of Sybase replication. In addition, these conditions must be met:

- ECDA for ODBC is installed on the same host as your standby Microsoft SQL Server database.
- You have obtained the licenses for the Replication Server and Mirror Activator.



- Microsoft SQL Server 2005 SP2, with database compatibility set to 2005(90), is installed and configured to serve as the target (standby) database and as the source (primary) database, from which Mirror Replication Agent replicates transactions. There must also be a user with permissions to log in to the standby database.
- Microsoft SQL Server JDBC driver 1.2 is installed.
- TCP/IP connectivity is available.
- Mirror Replication Agent has direct access to the Microsoft SQL Server online transaction logs or a mirrored copy of those logs.

During this procedure:

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

---

**Note** On Windows Vista, you must be logged in as an Administrator.

---

- Sybase recommends that you do not use Replication Server reserved words for object names and connection names. A complete list of reserved words are documented in Chapter 2 of the Replication Server *Reference Manual*.

## System requirements

The target Windows computer on which you are installing the Mirror Activator components must meet the following memory and disk space requirements:

**Table 2-1: Memory and disk requirements**

| Components               | Memory (RAM) | Disk space |
|--------------------------|--------------|------------|
| Replication Server       | 512MB        | 380MB      |
| ECDA                     | 512MB        | 440MB      |
| Mirror Replication Agent | 512MB        | 300MB      |

## Installing Mirror Activator components

This section explains pre-installation tasks and describes how to install and verify the installation of these components of Mirror Activator for Microsoft SQL Server:

- Mirror Replication Agent
- Replication Server
- ECDA Option for ODBC

---

**Note** Sybase recommends that you install all Mirror Activator components on the same host where a Microsoft SQL Server has already been installed and is running with both the primary and standby databases. This simplifies the quick-start process and avoids the step of separately installing a Microsoft SQL Server ODBC driver for ECDA Option for ODBC.

---

### Before you start

Before installing Mirror Activator components, identify the Sybase installation directories and obtain the installation software.

### Identifying the Sybase installation directory

Mirror Replication Agent, Replication Server, and ECDA are installed in the same base directory, which is identified by the SYBASE environment variable (%SYBASE%). Select a path on a host drive to be the recipient of the Sybase installation and configuration activities. For example:

```
c:\software\sybase
```

In addition, make sure you have a temporary installation directory with at least 100MB of space for the installation program to use.

### Accessing products from the SPDC Web site

If you are not using installation CDs, you must go to the Sybase product Download Center (SPCD) Web site to obtain replication software.

❖ **To download the replication software**

---

**Note** This procedure requires you repeat steps to download three products: Replication Server, Mirror Activator, and Enterprise Connect™ Data Access.

---

You must have a Login ID and Password to download software from the SPDC Web site.

- 1 Go to SPDC at <https://sybase.subscribenet.com/control/sybs/login> and enter your Login ID and Password.

The Product List page displays.

- 2 Select Mirror Activator.

The detailed Product Information page displays.

- 3 Select the Mirror Activator software you need by primary database, version, and platform.

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

- 4 Either:

- Select I Agree to accept the terms without reading them, 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 displays. 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 go back to the high-level Product Information page.

- 5 The Product Download Page lists the software and documentation for the product. Select FTP Download for the Mirror Activator 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 SyBooks download contains the rest of the documentation set.

---

- 6 Uncompress and extract all the installation images to your local drive.

## Installing Mirror Replication Agent

This section describes how to install Mirror Replication Agent. It also has a procedure to verify the installation.

The Mirror Replication Agent installation process creates a directory structure containing the Mirror Replication Agent binaries, scripts, Java Runtime Environment (JRE), and associated files. There are no configuration steps during installation.

Before you install Mirror Replication Agent, see the Mirror Replication Agent *Installation Guide* and complete the installation worksheet.

### Installation guidelines

While installing Mirror Replication Agent, use the following guidelines:

- Install Mirror Replication Agent on a server where it can directly access the Microsoft SQL Server mirror transaction logs.
- Create a Mirror Replication Agent instance and configure it for the Microsoft SQL Server database from which you want to replicate.

### ❖ To install Mirror Replication Agent software

- 1 Use *one* of the following methods to install the software.
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 30, go to the location where you extracted the Mirror Replication Agent 15.1 for Windows installation image, or
  - Insert the Mirror Replication Agent CD into your CD-ROM drive.
- 2 Launch the installation program by executing the setup routine for Windows:  

```
setup
```
- 3 If you see an error message saying that there is not enough temporary disk space, add the following to your setup command:  

```
setup -is:tempdir someDirectory
```

Here, *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.
- 4 On the Welcome window, click Next.
- 5 Select a geographic location that displays the license agreement. Select I Agree and click Next.
- 6 Enter the directory where you want to install Mirror Replication Agent:

- If the directory does not exist, the installation program prompts you to create it. Click Yes.
  - If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 7 Select Typical for the installation setup type and click Next.
  - 8 The installation program displays a summary of what is to be installed. Click Next.
  - 9 The installation program starts to install Mirror Replication Agent, displaying the progress status. If you see the following message;
 

There are newer files already installed and do you want to replace them?

Click No to All.
  - 10 When you receive a message that the installation was successful, click Next.
  - 11 The SySAM License Server window opens and displays this prompt:
 

Will licenses be obtained from the License Server?

    - If you answer Yes, you can leave the *license-mrx.lic* license file where it was installed (%SYBASE%\SYSAM-2\_0\licenses), or you can move the file to the *SYSAM-2\_0\licenses* directory of your SYSAM license server. In either case, the Mirror Replication Agent will verify the necessary license.
    - If you answer No, no further action is needed. The Mirror Replication Agent will use the installed license (*license-mrx.lic*) as it was installed in %SYBASE%\SYSAM-2\_0\licenses.

Click Next.
  - 12 The SySAM notification window prompts you to 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 these values from the installation worksheet you completed in the Mirror Replication Agent *Installation Guide*:
    - SMTP server host name
    - SMTP server port number
    - E-mail return address

- Recipients of the notification
- Message severity level of an event that will trigger e-mail. Your choices are:
  - Informational
  - Warning
  - Error

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

Click Next.

- 13 When you receive a message indicating that Mirror Replication Agent was successfully installed, click Finish.

---

**Note** Mirror Replication Agent should usually not be installed on the same host as the primary database for Mirror Activator.

---

For more information for the replication installation process, see the Mirror Replication Agent *Installation Guide*.

❖ **To verify the installation**

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

```
ma -v
```

If the Mirror Replication Agent installation completed successfully, you can see the Mirror Replication Agent version string and the Sybase copyright.

## Installing Replication Server

This section describes the steps for a Replication Server installation and a procedure to verify the installation.

---

**Note** If Replication Server is already installed, go to the next section called “Installing the ECDA Option for ODBC” on page 38.

---

### Installation guidelines

While installing Replication Server, use the following guidelines:

- 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.

### ❖ To install Replication Server

- 1 Use *one* of the following methods to download Replication Server software.
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 30, go to the location where you extracted the Replication Server 15.1 for Windows installation image (see “Accessing products from the SPDC Web site” on page 30 for details), or
  - Insert the Replication Server CD in your CD-ROM drive.
- 2 Launch the installation program by executing the setup routine for Solaris:  

```
setup
```
- 3 If you see an error message saying that there is not enough temporary disk space, add the following to your setup command:  

```
setup -is:tempdir someDirectory
```

Here, *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.
- 4 On the Welcome window, click Next.
- 5 Select a geographic location that displays the license agreement. Select I Agree and click Next.
- 6 Enter the directory where you want to install Replication Server:

- If the directory does not exist, the installation program prompts you to create it. Click Yes.
  - If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 7 Select Typical for the installation setup type and click Next.
  - 8 The installation program displays a summary of what is to be installed. Click Next.
  - 9 The installation program starts to install Replication Server, displaying the progress status. You may see the following message:  

```
There are newer files already installed and do you
want to replace them?
```

Click No to All.
  - 10 When you receive a message indicating that Replication Server software has successfully installed, click Next.

---

**Note** This procedure uses the SAMPLE\_RS Replication Server.

---

- 11 The SySAM License Server window opens and displays this prompt:  

```
Will licenses be obtained from the License Server?
```

Respond based on how you set up your environment and use the license obtained from the Software Product Download (SPDC) Web site for Replication Server.

Click Next.
- 12 The SySAM notification window prompts you to 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 your own values for the following fields:
  - SMTP server host name
  - SMTP server port number
  - Sender e-mail
  - Recipient e-mail
  - Message severity level of an event that triggers e-mail. Your choices are:



- Informational
- Warning
- Error

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

Click Next.

13 To start a sample Replication Server, select Yes and click Next.

14 The following information is displayed:

This page contains detailed information regarding the sample Replication Server. Please record this information.

The sample Replication Server will be called SAMPLE\_RS and will run on port 11752. It will be configured with a user of sa and no password.

The sample Replication Server will use an embedded RSSD called SAMPLE\_RS\_ERSSD that runs on port 11751. It will be configured with a user of SAMPLE\_RS\_RSSD\_prim and a password of SAMPLE\_RS\_RSSD\_prim\_ps.

The installer has updated the appropriate *interfaces* file or *sql.ini* file.

All files and logs associated with the sample Replication Server will be located in the directory %SYBASE%\REP-15\_1\samp\_repserver.

The sample Replication Server will be configured using the file %SYBASE%\REP-15\_1\samp\_repserver\SAMPLE\_RS.res.

Keep this information for future reference. Click Next.

15 When the SAMPLE\_RS is running and an “Installation was successful” message appears, click Finish.

#### ❖ To verify the installation

- 1 In a command window, change to the directory where you installed Replication Server.
- 2 Use isql to log in to Replication Server:

```
isql -Usa -P -SSAMPLE_RS
```

You should successfully log in to the Replication Server.

- 3 Exit isql.

## Installing the ECDA Option for ODBC

This section describes how to install the ECDA Option for ODBC.

---

**Note** If ECDA Option for ODBC is already installed, go to the next section called “Configuring Mirror Activator components” on page 40.

---

When installing ECDA for ODBC, be sure to install ECDA on the same server as the standby database. This eliminates a network hop and improves performance.

### ❖ To install ECDA for ODBC in GUI mode

---

**Note** If you want to install using the console mode or a response file, refer to the *ECDA Installation Guide* for Microsoft Windows.

---

- 1 Use *one* of the following methods to download the software:
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 30, go to the location where you extracted the ECDA Option for ODBC 15.0 installation image.
  - Insert the CD labeled “ECDA 15.0 Option for ODBC for Windows.”
- 2 Launch InstallShield by executing the Windows setup routine:

```
setup
```

- 3 If you see an error message saying that there is not enough temporary disk space, add the following to your setup command:

```
setup -is:tempdir someDirectory
```

Here, *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.

- 4 On the Welcome window, click Next.
- 5 Select a geographic location that displays the license agreement. Select I Agree and click Next.

- 6 Enter the directory where you want to install ECDA:
  - In the destination directory, enter the directory where the Mirror Replication Agent software is installed. Click Yes to install into an existing directory.
  - If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 7 Select Custom setup, and click Next.
- 8 Deselect ECDA Option for Oracle, and click Next.
- 9 Verify the modules to install. Click Next.

If you receive a message stating that your current installation contains newer files than those that are about to be installed and asking you if you want to replace these files, click No.
- 10 When you receive a message that the installation was successful, click Next.
- 11 The SySAM License Server window opens and displays this prompt:

Will licenses be obtained from the License Server?

Respond based on how you set up your environment and use the license obtained from the Software Product Download Center (SPDC) Web site for Mirror Activator.

Click Next.
- 12 The SySAM notification window prompts you to 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 your own values for the following fields:
  - SMTP server host name
  - SMTP server port number
  - Sender e-mail
  - Recipient e-mail
  - Message severity level of an event that triggers e-mail. Your choices are:
    - Informational
    - Warning

- Error

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

- 13 When you receive a message indicating that ECDA was successfully installed, click Finish.

ECDA for ODBC, including the DirectConnect™ server, is now installed.

## Configuring Mirror Activator components

This section describes the steps, in order, for configuring the components for Mirror Replication Agent for Microsoft SQL Server:

- 1 Configuring ECDA Option for ODBC
- 2 Configuring Replication Server
- 3 Configuring Mirror Replication Agent

### Configuring ECDA Option for ODBC

Configuring the ECDA Option for ODBC involves the following sequence of tasks:

- 1 Configuring the Microsoft SQL Server ODBC driver
- 2 Configuring and starting the DirectConnect server
- 3 Creating a maintenance user for replication
- 4 Creating the replicate objects
- 5 Verifying the ECDA installed objects

#### Configuring the Microsoft SQL Server ODBC driver

If you already have a data source configured for your standby database, skip this section and continue with “Configuring and starting the DirectConnect server” on page 41.

When you installed Microsoft SQL Server, an ODBC driver was also installed. Configure this driver to connect to the standby database.

**❖ To configure the Microsoft SQL Server ODBC driver**

- 1 Open the ODBC Driver Manager by navigating to Control Panel | Administrative Tools and double-clicking the Data Sources (ODBC) icon.
- 2 Click on the Drivers tab, and verify that the Microsoft SQL Server driver, which is named SQL Server, is version 2000.86.1830.00 or later. If it is an earlier version, go to the Microsoft Web site to download and install the correct version.
- 3 Click on the System DSN tab, click Add, select the Microsoft SQL Server driver. Click Finish.

---

**Note** Due to permissions, you may have to set up the ODBC data source as a user data source under the User DSN tab.

---

- 4 In the Name field, enter a name that identifies your standby Microsoft SQL Server, for example, *my\_standby\_mssql\_datasource*. Use this name later, when you configure the ECDA access service.
- 5 In the Server drop-down list, select the standby Microsoft SQL Server database, and click Next.
- 6 Be sure that “with SQL Server authentication using a login ID & password entered by the user” is selected, and the “Connect to Microsoft SQL Server to obtain default settings” check box is checked. Enter a valid Microsoft SQL Server Login ID and Password, and click Next.
- 7 Check the “Change the default database to” check box, select the standby database from the drop-down list, and click Next.
- 8 Click Finish.
- 9 Verify ODBC connectivity to the standby database by clicking Test Data Source. If connectivity is not successful, verify that you have configured the ODBC driver with the correct login name, database password, and permissions.
- 10 Click OK.
- 11 You should now see your Microsoft SQL Server data source in the list of System Data Sources. Click OK.

## Configuring and starting the DirectConnect server

Before you begin to create and configure the DirectConnect server, make sure that:

- The name for a valid Microsoft SQL Server account is available. This name is used as the administrator for the DirectConnect server.
- The number of the unused port to be used by the DirectConnect server is available.

❖ **Creating and configuring a DirectConnect server**

- 1 Start the Create Server wizard by navigating to the `%SYBASE%\DC-15_0\DCWizard` directory and executing the `DCWizard` script. The DirectConnect server creation wizard starts. Click Next.
- 2 Select ECDA Option for ODBC. Click Next.
- 3 Enter the server name for the new DirectConnect server (for example, “DCServer”) and the port number that you want the server to listen on. Click Next.
- 4 Enter the ECDA access service name that you want to use for the standby Microsoft SQL Server database, for example, `my_standby_mssql_service`. Click Next.
- 5 Verify the DirectConnect server information. If correct, select Create Server. Otherwise, click Back to return to the previous window and provide the correct information.
- 6 Click Next when the DirectConnect server prompts you.
- 7 When the Start DirectConnect Server (Optional) screen appears, click Finish.

If you have accidentally started the DirectConnect server, shut it down before proceeding:

```
%SYBASE%\DC-15_0\bin\stopsrvr -SDCServer
```

where `DCServer` is the name of the DirectConnect server.

- 8 Using a text editor, open the `%SYBASE%\DC-15_0\servers\<servername>\cfg\dcany.cfg` file, where `<servername>` is the name of the server you created in step 3 using the DCWizard.
- 9 In the file, locate the name of the Microsoft SQL Server access service that you created in step 4 using the DCWizard. Under that, set the `ConnectionSpec1` parameter to the name of the Microsoft SQL Server ODBC data source you created. Make sure that the file has at minimum the following configuration lines and that the options are set as follows:

```
[Service Library]
{Client Interaction}
```

```
[my_standby_mssql_service]
{ACS Required}
ConnectionSpec1=my_standby_mssql_datasource
{Client Interaction}
EnableAtStartup=yes
TransactionMode=long
```

---

**Note** TransactionMode=long indicates that the client session will manually issue the commit command for transactions.

---

Save the file.

10 Start the DirectConnect server:

- a Open a command window, and navigate to the DirectConnect installation directory. For example:

```
cd %SYBASE%\DC-15_0
```

- b Execute the *DC\_SYBASE.bat* script to set the DirectConnect environment variables.

- c Navigate to the DirectConnect *bin* directory:

```
cd bin
```

- d Start the DirectConnect server. This also starts all of the configured access services:

```
DCStart -SDCServer
```

where *DCServer* is the name of the DirectConnect server.

Wait until the DirectConnect server indicates that it is ready for connections before proceeding to the next step.

❖ **To verify that you can connect to Microsoft SQL Server**

- 1 Open a command window in the *%SYBASE%* directory of your ECDA installation.
- 2 Change to the *%SYBASE%\DC-15\_0* directory.
- 3 Set the environment variables by executing the *%SYBASE%\DC\_SYBASE.bat* file.
- 4 Use the Sybase isql utility to log in to Microsoft SQL Server through the DirectConnect service:

```
isql -Uinvaliduser -Ppassword
-Smy_standby_mssql_service
```

Here, *validuser* and *password* are any valid user and password, and *my\_standby\_mssql\_service* is the ECDA access service name identified in the previous procedure.

---

**Note** If the message `Data source not found` is returned and the data source has been entered correctly in the *dcany.cfg* file, you may need to set up the data source as a user data source under the User DSN tab of the ODBC Data Source Administrator instead of the System DSN tab. To do this, stop the DirectConnect server and return to “Configuring the Microsoft SQL Server ODBC driver” on page 40.

---

- 5 Verify the connection to the standby Microsoft SQL Server database by obtaining the DBMS name and version number:

```
select @@sqldbmsname
go
```

```
SQLDbmsName

Microsoft SQL Server
```

```
select @@sqlbmsver
go
```

```
SQLDbmsVer

09.00.1399
```

- 6 Exit the isql session with the quit command.

## Creating a maintenance user for replication

This section explains how to create a maintenance user in the standby Microsoft SQL Server database.

### Description of a maintenance user

The maintenance user is a valid Microsoft SQL Server user that the Replication Server uses to apply commands to the standby Microsoft SQL Server database. Replication Server requires one maintenance user to be defined for each connection.



**❖ To create a maintenance user in Microsoft SQL Server**

- 1 Using the Microsoft SQL Server Management Studio or the sqlcmd utility, connect to the standby database, and enter the following command to create maintenance user login:

```
CREATE LOGIN maintuser WITH PASSWORD='password',
DEFAULT_DATABASE=standby_db;
```

Here, *password* is the maintenance user login password, and *standby\_db* is the standby database name.

- 2 At the standby database, enter the following command to create a maintenance user:

```
CREATE USER maintuser FOR LOGIN maintuser;
```

- 3 Grant database owner permission to the maintenance user:

```
EXEC sp_addrolemember 'db_owner', 'maintuser';
```

## Setting up DDL replication

To replicate DDL commands for Microsoft SQL Server, you must create a unique DDL user that is granted privileges to execute all replicated DDL commands at the standby database. The DDL user is the standby database user name included in log transfer language (LTL) for replicating DDL commands to the standby database.

**❖ To create a DDL user in Microsoft SQL Server**

- 1 Using the Microsoft SQL Server Management Studio or the sqlcmd utility, connect to the standby database, and enter the following command to create a DDL user login:

```
CREATE LOGIN ddl_user WITH PASSWORD='password',
DEFAULT_DATABASE=standby_db;
```

Here, *password* is the DDL user login password, and *standby\_db* is the standby database name.

- 2 At the standby database, enter the following command to create a DDL user:

```
CREATE USER ddl_user FOR LOGIN ddl_user;
```

### Permissions

In addition to the permission to execute all replicated DDL commands at the standby database, the DDL user should also have the impersonate permission granted for all users whose DDL commands may be replicated to the standby database. This impersonate permission is necessary to switch session context in the standby database when executing a DDL command. This user switches context to apply the DDL command using the same privileges and default schema settings as the user who executed the DDL command at the primary database. To provide this context switch, the DDL user must have permission to execute the `execute as user` Microsoft SQL Server command for any user who might execute DDL commands to be replicated from the primary database.

For example, `user1` with a default schema of `schema1` executes the following DDL at the primary database:

```
create table tab1 (id int)
```

This results in the creation of a table named `schema1.tab1` at the primary database. At the standby database, `user2` with a default schema of `schema2`, cannot immediately execute this DDL because it will generate a table named `schema2.tab1`. Therefore, `user2`, whose name is specified by the `ddl_username` configuration parameter, must first execute the following command at the standby database to impersonate `user1`:

```
execute as user = 'user1'
```

The DDL can then be executed with the correct schema by `user2` at the standby database, generating a table named `schema1.tab1`.

### Granting impersonate permission

There are two ways to grant impersonate permission to the DDL user:

- You can grant database owner permission to the DDL user. In doing this, you implicitly grant impersonate permission.
- Alternately, you can grant impersonate permission explicitly with the following Microsoft SQL Server command:

```
GRANT IMPERSONATE ON USER::pri_ddl_user TO ddl_user
```

Here, *pri\_ddl\_user* is a user whose DDL is expected to be replicated to the standby database, and *ddl\_user* is the DDL user you created.

---

**Note** This grant command must be executed in the *standby* database, where the DDL user executes the DDL commands.

---

When you replicate DDL in Microsoft SQL Server, you must use Microsoft SQL Server as the standby database. You *cannot* replicate DDL commands from Microsoft SQL Server to non-Microsoft SQL Server standby databases.

---

**Note** To replicate DDL, Replication Server must have a database-level replication definition with replicate DDL set in the definition.

---

## Creating the replicate objects

### ❖ To create the replicate objects

- Sign on to the Microsoft SQL Server database through ECDA Option for ODBC as the maintenance user (assuming that the maintenance user that you created has the resource role to use create table, and execute this script using the Sybase isql utility):

```
isql -Umaintuser -Ppassword -Smy_standby_mssql_service
-i %SYBASE%\REP-15_1\scripts\hds_msss_setup_for_replicate.sql
```

Here, *my\_standby\_mssql\_service* is the name of the DirectConnect access service.

## Verifying the ECDA installed objects

### ❖ To verify the ECDA installed objects

- 1 Open another command window and navigate to the DirectConnect installation. For example:  

```
cd c:\sybase\DC-15_0
```
- 2 Set the environment variables by executing the *DC\_SYBASE.bat* file.
- 3 Use the Sybase isql utility to log in to the standby Microsoft SQL Server database through ECDA as *maintuser*.
  - To verify whether the *rs\_info* table has been created successfully:

```
isql -Umaintuser -Ppassword -Smy_standby_mssql_service
select * from rs_info
go
```

This following result is returned:

| rskey        | rsva1 |
|--------------|-------|
| -----        |       |
| charset_name | iso_1 |

```
sortorder_name bin_iso_1
(2 rows affected)
```

- To verify whether the `rs_lastcommit` table has been created successfully:

```
isql -Umaintuser -Ppassword -Smy_standby_mssql_service
select * from rs_lastcommit
go
```

The following is result is returned:

```
origin origin_qid secondary_qid origin_time
dest_commit_time

(0 rows affected)
```

If the ECDA objects are not installed successfully, no results are displayed, and you encounter an error message.

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

## Configuring Replication Server

Configuring Replication Server includes the following tasks:

- Configuring Replication Server for Microsoft SQL Server replication
- Configuring Replication Server for replication to the target database
- Configuring Replication Server for replication from the primary database

### Configuring Replication Server for Microsoft SQL Server replication

This section describes the configuration tasks for configuring Replication Server for Microsoft SQL Server.

#### ❖ To apply heterogeneous datatype support (HDS) scripts to RSSD

Prepare the Replication Server RSSD for replication to a Microsoft SQL Server database. To apply HDS support scripts to RSSD:

- 1 Set the environment variables by executing the `SYBASE.bat` file.
- 2 Change to the `%SYBASE%\REP-15_1\scripts` directory:
- 3 Load the following scripts into the RSSD:

```
isql -U SAMPLE_RS_RSSD_prim -P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ihds_msss_udds.sql -DSAMPLE_RS_ERSSD
```

```
isql -U SAMPLE_RS_RSSD_prim -P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ihds_msss_funcstrings.sql -DSAMPLE_RS_ERSSD
```

---

**Note** The *hds\_clt\_ase\_to\_msss.sql* script, while present in the *scripts* directory, does nothing.

---

- 4 Shut down Replication Server:
 

```
isql -Usa -P -SSAMPLE_RS
shutdown
go
```
- 5 Navigate to the `%SYBASE%\REP-15_1\samp_repserver` directory:
- 6 Run the *RUN\_SAMPLE\_RS* script, to start the SAMPLE\_RS Replication Server.

## Configuring Replication Server for replication to the target database

This section describes the required configurations for Replication Server for replication to the target Microsoft SQL Server database.

### ❖ To create a Replication Server connection to the standby database

- 1 Navigate to the Replication Server *scripts* directory:  
`%SYBASE%\REP-15_1\scripts`.
- 2 Make a copy of the *hds\_msss\_connection\_sample.sql* script:
 

```
copy hds_msss_connection_sample.sql
my_hds_msss_connection_sample.sql
```
- 3 Before executing the *my\_hds\_msss\_connection\_sample.sql* script against your Replication Server, change `{msss_dco}`.`{dbname}` to the name of the connection that Replication Server uses to connect to the DirectConnect access service, where:
  - `{msss_dco}` is the DirectConnect access service name.
  - `{dbname}` is the Microsoft SQL Server standby database name.
  - `rs_maint_user` and `rs_maint_user_pwd` are the maintenance user and password created in step 1 of “Creating a maintenance user for replication” on page 44. For example:

```
create connection to
```

```
my_standby_mssql_service.standby_db
set error class rs_sqlserver_error_class
set function string class rs_msss_function_class
set username maintuser
set password "password"
set batch to "off"
go
```

---

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

---

- 4 Create the connection to the standby database using the Sybase isql utility:

```
isql -Usa -P -SSAMPLE_RS -i my_hds_msss_connection_sample.sql
```

- 5 To verify the Replication Server connection to the standby database, log in to SAMPLE\_RS and verify the Replication Server connection to the standby database:

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

You should see a message returned for the DSI connection  
*my\_standby\_mssql\_service.standby\_db*:

| Spid | Name     | State            | Info                                    |
|------|----------|------------------|-----------------------------------------|
| 14   | DSI EXEC | Awaiting Command | 101(1) SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD  |
| 10   | DSI      | Awaiting Message | 101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD     |
| 9    | SQM      | Awaiting Message | 101:0 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD   |
| 19   | DSI EXEC | Awaiting Command | 102(1) my_standby_mssql_service.standby |
| 18   | DSI      | Awaiting Message | 102 my_standby_mssql_service.standby_db |
| 17   | SQM      | Awaiting Message | 102:0 my_standby_mssql_service.standby_ |
| 11   | dSUB     | Sleeping         |                                         |
| 6    | dCM      | Awaiting Message |                                         |
| 7    | dAIO     | Awaiting Message |                                         |
| 12   | dREC     | Sleeping         | dREC                                    |
| 21   | USER     | Active           | sa                                      |
| 5    | dALARM   | Awaiting Wakeup  |                                         |
| 13   | dSYSAM   | Sleeping         |                                         |

Verify that the status is “Awaiting Message” or “Awaiting Command.”

- 6 Exit the isql session with the quit command.

## Configuring Replication Server for replication from the primary database

This section describes the configuration steps required for replication from the primary database. It includes:

- Creating a Replication Server connection to the primary database
- Creating a database replication definition
- Creating a database replication subscription

### ❖ To create a Replication Server connection to the primary database

- 1 Go to the Mirror Replication Agent scripts directory,  
`%SYBASE%\MA-15_I\scripts\mssql`
- 2 Make a copy of the `mssql_create_rs_primary_connection.sql` script:  

```
copy mssql_create_rs_primary_connection.sql
my_mssql_create_rs_primary_connection.sql
```
- 3 Before executing the `my_mssql_create_rs_primary_connection.sql` script against your Replication Server, change `{pds}`.`{pdb}` to the name of the connection between Mirror Replication Agent and Replication Server, where:
  - `{pds}` is `rs_source_ds` (located in `%SYBASE%\MA-15_I\init\mssql.rs`).
  - `{pdb}` is `rs_source_db` (located in `%SYBASE%\MA-15_I\init\mssql.rs`).

---

**Note** You must save the defined values to be used for the Mirror Replication Agent `mssql.rs` resource file later in the process.

---

For example, `mymrm.pubs`.

- 4 Change `sa` and `sa_pwd` to the user ID and password of the Microsoft SQL Server user who must have permission to apply DML operations against all user tables to be replicated, where:
  - `sa` is the Microsoft SQL Server user ID.
  - `sa_pwd` is the Microsoft SQL Server password.

---

**Note** This user must be a valid user in the Microsoft SQL Server database.

---

The following example command creates a Replication Server connection to the primary database:

```
create connection to mymr.pubs
set error class rs_sqlserver_error_class
set function string class rs_msss_function_class
set username newuser
set password newuser
with log transfer on, dsi_suspended
go
```

5 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imy_mssql_create_rs_primary_connection.sql
```

A message is displayed that indicates the Replication Server connection to the primary database is created.

❖ **To create a database replication definition**

1 Navigate to the `%SYBASE%\MA-15-1\scripts\sybase` directory.

2 Make a copy of the `create_rs_db_repdef.sql` script:

```
copy create_rs_db_repdef.sql
my_mssql_create_rs_db_repdef.sql
```

3 Before executing the `my_mssql_create_rs_db_repdef.sql` script, change the value of “`{pds}.pdb`” to the name of the connection string you defined for the primary database, where:

- `pds` is `rs_source_ds`, located in `%SYBASE%\MA-15-1\init\mssql.rs`, for example, `mymrm`.
- `pdb` is `rs_source_db`, located in `%SYBASE%\MA-15-1\init\mssql.rs`, for example, `pubs`.

This is a database replication definition example:

```
create database replication definition mymrm_repdef1
with primary at mymrm.pubs
replicate DDL
go
```

4 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imy_mssql_create_rs_db_repdef.sql
```

A message is displayed that indicates the database replication definition is created.

❖ **To create a database replication subscription**

1 Make a copy of the `create_rs_db_sub.sql` script:

```
copy create_rs_db_sub.sql
```



`my_mssql_create_rs_db_sub.sql`

- 2 Before executing the `my_mssql_create_rs_db_sub.sql` script, change “{pds}.{pdb}” and “{rds}.{rdb}” to the appropriate connection name, where:
  - `pds` is `rs_source_ds`, located in `%SYBASE%\MA-15_I\init\mssql.rs`.
  - `pdb` is `rs_source_db`, located in `%SYBASE%\MA-15_I\init\mssql.rs`.
  - `rds` is the DirectConnect access service name, for example, `my_standby_mssql_service`.
  - `rdb` is the name of the standby database, for example, `standby_db`.

Here is an example of create database replication subscription:

```
create subscription mymr_sub1
for database replication definition mymr_repdef1
with primary at mymr.pubs
with replicate at my_standby_mssql_service.standby_db
without materialization
go
```

- 3 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imy_mssql_create_rs_db_sub.sql
```

A message is displayed that indicates the subscription is in the process of being created.

- 4 Use `isql` to log in to the Replication Server:

```
isql -Usa -P -SSAMPLE_RS
```

- 5 Verify that the database replication definition database subscription exists:

```
check subscription mymr_sub1
for database replication definition mymr_repdef1
with primary at mymr.pubs
with replicate at my_standby_mssql_service.standby_db
go
```

A message is displayed that indicates the subscription is valid for the primary and replicate databases.

- 6 Replication Server comes with users who have the permissions necessary for Mirror Replication Agents. However, because Mirror Activator has the capability to create replication definitions, you must grant `CREATE OBJECT` permission to the Mirror Replication Agent user:

```
grant create object to SAMPLE_RS_ra
```

go

Permission is granted to user SAMPLE\_RS\_ra.

---

**Note** This user will be entered in the `rs_username` parameter when building the Mirror Replication Agent for Microsoft SQL Server instance. If you want a user other than the default replication agent user, make sure that the user is granted `CONNECT SOURCE` and `CREATE OBJECT` permission.

---

## Configuring Mirror Replication Agent

Configuring Mirror Replication Agent involves the following sequence of tasks:

---

**Note** On Windows Vista, you must be logged in as an Administrator.

---

- 1 Configuring and verifying the primary Microsoft SQL Server
- 2 Configuring and validating the Mirror Replication Agent resource file
- 3 First-time primary database initialization
- 4 Subsequent primary database initialization
- 5 Verifying the correct initialization for the primary database
- 6 Initializing a Mirror Replication Agent instance
- 7 Verifying the Mirror Replication Agent `REPLICATING` state

## Configuring and verifying the primary Microsoft SQL Server

Before you configure Mirror Replication Agent and configure the primary Microsoft SQL Server database for replication, complete these procedures, which are explained in the following subsections.

- Verify the compatibility level of Microsoft SQL Server.
- Disable Microsoft SQL Server replication.
- Create a Microsoft SQL Server user and grant Microsoft SQL Server permissions.
- Enable remote DAC.

- Obtain the Microsoft SQL Server DAC port number.
- Verify the supplemental logging of primary key data.
- Make the primary transaction log files readable for Mirror Replication Agent.
- Stop all Microsoft SQL Server services related to your Microsoft SQL Server instance, including the Analysis Service and Reporting Service.
- Enable TCP/IP.

❖ **To verify the compatibility level of Microsoft SQL Server**

Mirror Replication Agent supports only Microsoft SQL Server 2005 Service Pack 2 and later. Verify that the database compatibility level is set to Microsoft SQL Server 2005(90).

- 1 In the Microsoft SQL Server Management Studio, right-click on your primary database and select Properties. This brings up the Database Properties dialog box.
- 2 In Database Properties, click Options, and verify that the Compatibility level is “Microsoft SQL Server 2005(90).” If it is not, select that compatibility from the drop-down list and click OK.

❖ **To disable Microsoft SQL Server replication**

- A Microsoft SQL Server publication cannot be created in the primary database used by Mirror Replication Agent for Microsoft SQL Server, and you cannot simultaneously use Microsoft replication and Mirror Replication Agent on the same Microsoft SQL Server database. If a Microsoft SQL Server publication already exists, you must remove the publication before using Mirror Replication Agent for Microsoft SQL Server.

❖ **To create a Microsoft SQL Server user and grant permissions**

- 1 Using Microsoft SQL Server Management Studio or the sqlcmd utility, connect to the primary Microsoft SQL Server, run the following commands to create a Microsoft SQL Server user named “ra\_user” with the password “sybase,” and grant permissions to the user:

```
CREATE LOGIN ra_user WITH PASSWORD='sybase',
DEFAULT_DATABASE=primary_database;
```

- 2 At the primary database, enter the following command to allow the user to access the primary database:

```
CREATE USER ra_user FOR LOGIN ra_user;
```

- 3 Grant sysadmin permission to ra\_user:

```
EXEC sp_addrolemember 'sysadmin', 'ra_user';
```

❖ **To enable remote DAC from the GUI**

- 1 From the Windows Start menu, choose Microsoft SQL Server | Configuration Tools | Surface Area Configuration | Surface Area Configuration for Features.
- 2 In the Surface Area Configuration for Features window, choose DAC under the MSSQLSERVER/Database Engine tree, and make sure the Enable remote DAC check box is selected.

---

**Note** After you select Enable remote DAC, you should only click Apply and OK.

---

❖ **To enable remote DAC from the command line**

- 1 Log in to Microsoft SQL Server as a system administrator.
- 2 Invoke the sp\_configure stored procedure as follows:

```
sp_configure 'remote admin connections', 1
```
- 3 Issue the RECONFIGURE command to validate the configuration changes.

❖ **To obtain the Microsoft SQL Server DAC port number**

- 1 Open the *ERRORLOG* file in a text editor. This file is located in the *log* directory of your Microsoft SQL Server. For example:

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

- 2 Search for the string “Dedicated admin” to find an entry similar to the following:

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

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

❖ **To verify the version of the Microsoft Filer Manager Library**

To make the primary transaction log files readable for Mirror Replication Agent, the Microsoft Filter Manager Library must be version 5.1.2600.2978 or later.

- 1 To determine the version of the library, in Windows Explorer, right-click *c:\windows\system32\ftlib.dll*, select Properties, and click the Version tab in the Properties dialog box.
- 2 If the version is earlier than 5.1.2600.2978, go to the Microsoft Web site at <http://windowsupdate.microsoft.com>, and update your Windows system.

❖ **To make the primary transaction log files readable for Mirror Replication Agent**

Perform the following steps to install and set up the *sybfilter* driver so that Mirror Replication Agent can read the primary transaction log files.

- 1 In Windows Explorer, navigate to the *sybfilter* driver installation directory. On Windows, this directory is located at *%SYBASE%\MA-15\_1\system\<platform>*. Here, *<platform>* is *winx86*, *winx64*, or *winvistax64*.
- 2 Right-click the *sybfilter.inf* file, and click on Install to install the *sybfilter* driver. To reinstall the *sybfilter* driver, click Install on the correct filter.

---

**Note** There can be only one installation of the *sybfilter* driver on a Windows machine. After the driver is installed, it works for all Mirror Replication Agent for Microsoft SQL Server instances running on the same machine. To verify that the *sybfilter* driver has been installed and is running, enter the following command at the DOS prompt:

```
sc query sybfilter
```

If the *sybfilter* driver has been installed, the *sc* command displays details about the driver.

---

- 3 Under any directory, create a configuration file to store all log file paths for primary databases. The configuration file must have a *.cfg* suffix. For example, under the directory *%SYBASE%\MA-15\_1\system\<platform>*, create a file named *LogPath.cfg*.
- 4 Add a system environment variable named *RACFGFilePath*, and set its value to the path of the configuration file.
  - a Open the Control Panel, click System, click the Advanced tab, and choose Environment Variables.
  - b Click New to add a new system variable.
  - c Name the variable *RACFGFilePath*, and set its value to the location of the your configuration file.

- 5 Start the sybfilter driver.
  - a Set the environment variables at a DOS prompt by executing `%SYBASE%\SYBASE.bat`.
  - b Change to `%SYBASE%\MA-15_I\bin`, and execute `sybfiltermgr.exe`.
  - c Enter `start` at the console.
- 6 Add the log file path to the *sybfilter* driver with the user manager or by modifying the configuration file.

- *User manager* - Use the add command in the management console. The syntax for this command is as follows:

`add serverName dbName logFilePath`

For example, to add the log file named *pdb2\_log.ldf* at *C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\Data\* to the *dbName* database on the *serverName* data server, use the following:

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

---

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

---

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

```
[myserver, pdb1]
log_file_path=C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb1_log.ldf
log_file_path=C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb12_log.ldf
[myserver, pdb2]
log_file_path=C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Data\pdb2_log.ldf
```

---

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

---

- 7 If you added a log file for your primary database before adding the log file path to the *sybfilter* driver, you must restart Microsoft SQL Server to make the log file readable.

- 8 At the management console, enter `check` to verify that log files are readable. If some log files are unreadable, make sure the files have been created and that Microsoft SQL Server has been restarted, if necessary.

- 9 Exit the management console by entering `exit`.

For more information on the *sybfilter* driver, see the Mirror Replication Agent *Primary Database Guide*.

❖ **To stop the Microsoft SQL Server services related to your Microsoft SQL Server instance**

If you have not installed Microsoft SQL Server-related services, skip to the next procedure.

- 1 From the Windows Control Panel, select Administrative Tools | Services.
- 2 In the listed services, find the services named Microsoft SQL Server Analysis Service (*SERVER*) and Microsoft SQL Server Reporting Service (*SERVER*), where *SERVER* is the name of your Microsoft SQL Server data server.
- 3 Stop these services.

❖ **To enable TCP/IP from Windows**

- 1 From the Windows Start menu, choose Microsoft SQL Server | Configuration Tools | Surface Area Configuration | Surface Area Configuration for Services and Connections.
- 2 Under the Database Engine tree, click Remote Connections.
- 3 Select “Local and remote connections,” and “Using both TCP/IP and Named Pipes protocols.”

❖ **To enable TCP/IP from the command line**

- 1 Log in to Microsoft SQL Server as a system administrator.
- 2 Invoke the `sp_configure` stored procedure as follows:
 

```
sp_configure 'remote access', 1
```
- 3 Issue the `RECONFIGURE` command to validate the configuration changes.

❖ **To install the Microsoft SQL Server JDBC driver**

- 1 Go to the Microsoft download site (<http://www.microsoft.com/downloads>) and search for the following string:

Microsoft SQL Server 2005 JDBC driver

- 2 Click on the link for the 1.2 driver, and follow the instructions to download and install the driver.
- 3 Open Control Panel | System, click on the Advanced tab, and click Environment Variables.
- 4 In the System Variables window, if the CLASSPATH variable is not listed, click New. Otherwise, select it and click Edit.
- 5 Enter the full path of the JDBC driver using the semi-colon (;) to separate it from any other drivers, for example, *C:\path\sqljdbc.jar*.
- 6 Click OK three times.

## Configuring and validating the Mirror Replication Agent resource file

A single installation of the Mirror Replication Agent can support replication from multiple databases. However, one Mirror Replication Agent instance is needed for each Microsoft SQL Server database that is to be replicated. This section describes how to prepare for and create a Mirror Replication Agent instance.

### ❖ To configure and validate the Mirror Replication Agent parameters in a resource file

- 1 Locate the resource file template.

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

The Microsoft SQL Server resource file template can be found at *%SYBASE%\MA-15\_I\init\mssql.rs*.

- 2 Create an instance resource file.

Copy the resource file template *%SYBASE%\MA-15\_I\init\mssql.rs* to another file. For example:

```
copy mssql.rs myma.rs
```

Using a text editor, change the values assigned to properties in the resource file so they match the values required for your environment. The majority of these values are host, port, user, and password information for the Microsoft SQL Server, Replication Server, and RSSD connections.



The resource file is self-documenting, with each parameter described. For more information, refer to the Mirror Activator *Administration Guide*, in Chapter 2, “Setup and Configuration,” the section titled, “Creating a Mirror Replication Agent instance using resource files.”

**Table 2-2: Resource file parameters**

| Parameter           | Description                                                                                                                                                                                                                                                                                                                       | Example values                                    |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| instance_name       | Any valid name                                                                                                                                                                                                                                                                                                                    | myma                                              |
| admin_port          | Port number that Mirror Replication Agent uses.                                                                                                                                                                                                                                                                                   | 20500 (if in use, select a different port number) |
| pds_server_name     | Name of the primary Microsoft SQL Server.                                                                                                                                                                                                                                                                                         | TEAMSTER                                          |
| pds_port_number     | Port number for Microsoft SQL Server.                                                                                                                                                                                                                                                                                             | 1521                                              |
| pds_dac_port_number | Microsoft SQL Server Dedicated Administration connection port number.                                                                                                                                                                                                                                                             | 1348                                              |
| pds_database_name   | Name of the primary database.                                                                                                                                                                                                                                                                                                     | pubs                                              |
| pds_username        | User ID that Mirror Replication Agent uses to access the primary data server.                                                                                                                                                                                                                                                     | ra_user                                           |
| pds_password        | Password for pds_username.                                                                                                                                                                                                                                                                                                        | sybase                                            |
| create_pds_username | Specifies whether or not the ma_admin utility, invoked without the -r option, should create the Mirror Replication Agent primary database user login if the login does not already exist in the primary data server. For information about the ma_admin utility -r option, see the Mirror Activator <i>Administration Guide</i> . | yes                                               |
| rs_host_name        | Machine where Replication Server is installed.                                                                                                                                                                                                                                                                                    | teamster                                          |
| rs_port_number      | Port where Replication Server is installed.                                                                                                                                                                                                                                                                                       | 11752                                             |
| rs_username         | Replication Server user with CONNECT SOURCE and CREATE OBJECT capabilities.                                                                                                                                                                                                                                                       | SAMPLE_RS_ra                                      |
| rs_password         | Password for rs_username.                                                                                                                                                                                                                                                                                                         | SAMPLE_RS_ra_ps                                   |
| rs_source_ds        | Valid name representing data server of primary Microsoft SQL Server database.                                                                                                                                                                                                                                                     | mymrm                                             |
| rs_source_db        | Valid name representing primary Microsoft SQL Server database.                                                                                                                                                                                                                                                                    | pubs                                              |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                | Example values                                                                                       |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| rs_charset          | Character set that Replication Server is using.<br><br><b>Note</b> The value defined for the rs_charset configuration parameter MUST match the RS_charset value in Replication Server's configuration file,<br><i>%SYBASE%\REP-15_1\install\&lt;server&gt;.cfg</i> or<br><i>%SYBASE%\REP-15_1\samp_repserver\SAMPLE_RS.cfg</i> .                           | cp850                                                                                                |
| rssd_host_name      | Machine where RSSD resides.                                                                                                                                                                                                                                                                                                                                | teamster                                                                                             |
| rssd_port_number    | Port number where RSSD resides.                                                                                                                                                                                                                                                                                                                            | 11751                                                                                                |
| rssd_database_name  | Database name for RSSD.                                                                                                                                                                                                                                                                                                                                    | SAMPLE_RS_ERSSD                                                                                      |
| rssd_username       | Valid user for RSSD.                                                                                                                                                                                                                                                                                                                                       | SAMPLE_RS_RSSD_maint                                                                                 |
| rssd_password       | Password for rssd_username.                                                                                                                                                                                                                                                                                                                                | SAMPLE_RS_RSSD_maint_ps                                                                              |
| ddl_username        | Identifies the database user name included in LTL for replicating DDL commands to the standby database.                                                                                                                                                                                                                                                    | ddl_user                                                                                             |
| ddl_password        | Identifies the password for ddl_username.                                                                                                                                                                                                                                                                                                                  | The password you created previously                                                                  |
| start_instance      | Automatically start the instance that was created.                                                                                                                                                                                                                                                                                                         | yes                                                                                                  |
| initialize_instance | Automatically initialize the new Mirror Replication Agent instance.<br><br><b>Note</b> Set the initialize_instance parameter to yes (the default is no) to initialize the Mirror Replication Agent instance. If you retain the default value of no, you must manually issue the ra_devicepath command after initialization to set the correct mirror path. | yes                                                                                                  |
| ra_devicepath_x     | Primary database log mirror location information. These parameter names must be unique in the format ra_devicepath_x, where x is a unique sequential number and the parameter value contains both the device name and the device path separated by a comma.                                                                                                | <i>ra_devicepath_1= logA,<br/>/dev/pubs2/logA.dat</i> , or leave commented if not using mirrored log |

---

**Warning!** The combined values of rs\_source\_ds and rs\_source\_db must match the “pds.pdb” values of the Replication Server primary connection name.

---

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

```
%SYBASE%\MA-15_1\bin\ma_admin -vr myma.rs
```

Validation results are returned as one of the following:

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

If any validation fails, the `ma_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 because of this process.

## First-time primary database initialization

To initialize the primary data server, Mirror Replication Agent for Microsoft SQL Server installs objects at the data server level and at the database level. The data server-level modifications are required only once. Use this procedure if you are initializing the primary data server and a Mirror Replication Agent instance for the first time.

If you have previously initialized the primary data server and a Mirror Replication Agent instance and want to initialize another Mirror Replication Agent instance, skip this procedure, and follow the instructions in “Subsequent primary database initialization” on page 66.

### ❖ To configure the primary data server and a Mirror Replication Agent instance for the first time

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

```
SQL Server (TEAMSTER)
```
  - b Stop the service.

---

**Note** If you have a DirectConnect service attached to the primary data server, make sure that service is stopped.

---

- 2 Open a command window, and restart Microsoft SQL Server in single-user mode:

```
"C:\Program Files\Microsoft SQL
Server\MSSQL.1\MSSQL\Binn\sqlservr.exe" -m -s
instanceName
```

Here, *instanceName* is the name of the Microsoft SQL Server instance.

- 3 Execute the resource file to create the Mirror Replication Agent instance. In your resource file, if you set *start\_instance* to yes, this step also starts the Mirror Replication Agent instance. In your resource file, if you also configured *initialize\_instance* to yes, this step also initializes the Microsoft SQL Server and the Mirror Replication Agent instance.

Use the *ma\_admin* utility to create the Mirror Replication Agent instance using the *-r* parameter. For example:

```
%SYBASE%\MA-15_1\bin\ma_admin -r myma.rs
```

Execution results are returned as one of the following:

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

For more information, refer to the Mirror Activator *Administration Guide*, Chapter 2, “Setup and Configuration,” in the section titled, “Creating a Mirror Replication Agent instance using resource files.”

- 4 If your resource file set *start\_instance* to no, change to the *SYBASE* directory and start the Mirror Replication Agent instance:

```
cd %SYBASE%\MA-15_1\myma
```

Execute the *RUN* file, for example:

```
RUN_myma
```

- 5 Using the *dsedit* utility, update the Replication Server *sql.ini* file with the Mirror Replication Agent location, and verify the connection to the Mirror Replication Agent:

- a Open a command window in the *SYBASE* directory of your Mirror Replication Agent installation.
- b Set the environment variables by executing the *SYBASE.bat* file.
- c Use the Sybase *isql* utility to log in to Mirror Replication Agent:

```
isql -Usa -P -Smyma
```

Here, *myma* is the name of Mirror Replication Agent.

- 6 If your resource file set `initialize_instance` to `no`, initialize the primary data server and the Mirror Replication Agent instance:
  - a Open a command window, and log in to the Mirror Replication Agent instance.
  - b Issue the following command:
 

```
pdb_init
```
  - c Shut down the Mirror Replication Agent instance:
 

```
shutdown
```
- 7 Stop the Microsoft SQL Server in single-user mode (only one administrator can log in when Microsoft SQL Server is running in single-user mode):
  - a Use the `sqlcmd` utility to log in to the server:
 

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

Here, *username*, *password*, and *serverName* are your user ID, password, and Microsoft SQL Server name.
  - b Issue the shutdown command.
- 8 Restart Microsoft SQL Server in multi-user mode (normal start):
  - a In Control Panel | Administrative Tools | Services, find the service named SQL Server (*SERVER*) where *SERVER* is the name of your Microsoft SQL Server data server. For example:
 

```
SQL Server (TEAMSTER)
```
  - b Start the service.
- 9 Start the Mirror Replication Agent instance:
 

```
%SYBASE%\MA-15_1\instancename\RUN_instancename
```

where *instancename* is the name of the Mirror Replication Agent instance.

The Mirror Replication Agent instance is now running, and the primary data server and the Mirror Replication Agent instance have been initialized. Continue with “Verifying the Mirror Replication Agent REPLICATING state” on page 68.

## Subsequent primary database initialization

Use the following procedure only if you have already initialized the primary data server, as described in “First-time primary database initialization” on page 63, and you want to create another Mirror Replication Agent instance. This procedure creates and initializes only the Mirror Replication Agent instance, not the primary data server.

### ❖ To create and initialize a Mirror Replication Agent instance

- 1 Execute the resource file to create the Mirror Replication Agent instance.

- After the resource file has been validated as described in “Configuring and validating the Mirror Replication Agent resource file” on page 60, use the `ma_admin` utility to create the Mirror Replication Agent instance using the `-r` parameter, for example:

```
%SYBASE%\MA-15_1\bin\ma_admin -r myma.rs
```

Execution results are returned as one of the following:

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

For more information about the resource file or its use, refer to the *Mirror Activator Administration Guide*, Chapter 2, “Setup and Configuration,” in the section titled, “Creating a Mirror Replication Agent instance using resource files.”

- 2 If your resource file set `start_instance` to `no`, start the Mirror Replication Agent instance.
  - a Change to the directory containing your Mirror Replication Agent instance *RUN* file.
  - b To start the Mirror Replication Agent instance, execute the *RUN* file, for example:

```
RUN_myma
```

- 3 Using the `dsedit` utility, update the Replication Server *sql.ini* file with the Mirror Replication Agent location, and verify the connection to the Mirror Replication Agent.
  - a Open a command window in the *SYBASE* directory of your Mirror Replication Agent installation.
  - b Set the environment variables by executing the *SYBASE.bat* file.
  - c Use the Sybase `isql` utility to log in to Mirror Replication Agent:

```
isql -Usa -P -Smyma
```

Here, *myma* is the name of Mirror Replication Agent.

- 4 If your resource file set `initialize_instance` to `no`, open a new command window, log in to the Mirror Replication Agent instance, and initialize the Mirror Replication Agent instance and the primary Microsoft SQL Server database using the following Mirror Replication Agent command:

```
pdb_init
```

- 5 If your resource file completed successfully and the instance was initialized and started, proceed to “Verifying the Mirror Replication Agent REPLICATING state” on page 68.

## Verifying the correct initialization for the primary database

Verify the correct initialization for the primary database.

### ❖ To verify the correct initialization for the primary database

- 1 Use the Sybase `isql` utility to log in to Mirror Replication Agent:

```
isql -Usa -P -Smyma
```

Here, *myma* is the name of Mirror Replication Agent.

- 2 Invoke the `pdb_init` command with the `move_truncpt` keyword, which sets the truncation point to the end of the primary database transaction log:

```
pdb_init move_truncpt
go
```

A message should indicate that the procedure was successful. For more information, refer to the Mirror Activator *Administration Guide*, in Chapter 2, “Setup and Configuration.”

## Initializing a Mirror Replication Agent instance

To initialize a Mirror Replication Agent instance, use the `ra_init` command, which initializes the Mirror Replication Agent system database by reading schema information and transaction log location information from the primary Microsoft SQL Server database. The `ra_init` command also causes the `pdb_automark_tables` and `pdb_auto_create_repdefs` settings to take effect.

---

**Warning!** The default value for the `pdb_auto_create_repdefs` configuration property is true, which will cause a replication definition to be created for each table that is marked for replication during processing of the `ra_init` command. If you have thousands of tables, this may result in significant additional execution time. To avoid this additional execution time, set the `pdb_auto_create_repdefs` configuration property to false before invoking the `ra_init` command. When `ra_init` execution completes and before replicating, execute the `rs_create_repdef all` command to create your replication definitions.

---



---

**Note** For a production setup, initializing a Mirror Replication Agent instance should coincide with creating a backup copy of data used to materialize the standby database.

---

### ❖ To initialize a Mirror Replication Agent instance

- Invoke the `ra_init` command:

```
ra_init
go
```

A message should indicate that the procedure was successful. The Mirror Replication Agent instance is now running and has been initialized.

## Verifying the Mirror Replication Agent *REPLICATING* state

### ❖ To verify the Mirror Replication Agent *REPLICATING* state

- 1 Enter:

```
resume
go
```

If the Mirror Replication Agent successfully transfers to a *REPLICATING* state, the following result is returned:

```
State Action

```



REPLICATING Ready to replicate data.

A returned state of *ADMIN* indicates that an error prevented the *REPLICATING* state from being achieved. To determine the error, review the contents of the Mirror Replication Agent system log.

- 2 The `ra_status` command returns the state of the Mirror Replication Agent. It is good practice to verify that the Mirror Replication Agent remains in *REPLICATING* state, even after the resume command executes successfully.

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

```
ra_status
go
```

If the Mirror Replication Agent is in *REPLICATING* state, `ra_status` returns:

```
State Action

REPLICATING Ready to replicate data.
```

A returned state of *ADMIN* indicates that an error prevented the *REPLICATING* state from being achieved. To determine the error, review the contents of the Mirror Replication Agent system log.

- 3 Validate that both primary and standby connections are active:

```
isql -Usa -P -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 Mirror Replication Agent connection, if established for the standby database connection, is usually down, because you are not replicating data from the standby database.
- 

The following is displayed:

| Spid | Name     | State            | Info                                   |
|------|----------|------------------|----------------------------------------|
| 14   | DSI EXEC | Awaiting Command | 101(1) SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD |
| 10   | DSI      | Awaiting Message | 101 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD    |

|    |          |                                                          |
|----|----------|----------------------------------------------------------|
| 9  | SQM      | Awaiting Message 101:0 SAMPLE_RS_ERSSD.SAMPLE_RS_ERSSD   |
| 19 | DSI EXEC | Awaiting Command 102(1) my_standby_mssql_service.standby |
| 18 | DSI      | Awaiting Message 102 my_standby_mssql_service.standby_db |
| 17 | SQM      | Awaiting Message 102:0 my_standby_mssql_service.standby_ |
| 11 | dSUB     | Sleeping                                                 |
| 6  | dCM      | Awaiting Message                                         |
| 7  | dAIO     | Awaiting Message                                         |
| 12 | dREC     | Sleeping dREC                                            |
| 21 | USER     | Active sa                                                |
| 5  | dALARM   | Awaiting Wakeup                                          |
| 13 | dSYSAM   | Sleeping                                                 |

## Testing replication

After you finish setting up the replication system, you can test replication.

### ❖ To test replication

- 1 Connect to the primary Microsoft SQL Server database as a regular user, not the maintenance user. Make sure the regular user also exists in the standby database.

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

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

This SQL statement is replicated by the user defined in the `ddl_username` Mirror Replication Agent configuration parameter. After the time interval indicated in the `scan_sleep_max` Mirror Replication Agent configuration parameter, the T1 table is replicated to the target Microsoft SQL Server instance.

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

```
grant all on T1 to public
go
```

This SQL statement is also replicated to the standby database by the user defined in the `ddl_username` Mirror Replication Agent configuration parameter.

- 4 Connect to the standby Microsoft SQL Server database as a regular user, not the maintenance user.

- 5 After the table is marked and the Mirror 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. See “Configuring Replication Server for replication from the primary database” on page 51.

Mirror Replication Agent applies only committed transactions to a standby database.

```
insert into T1 values (42, 'foo')
go
```

- 6 Use your preferred Microsoft SQL Server query tool to examine the standby site for results, and compare the contents of your test table from both the primary and standby sites.

## Resetting the primary Microsoft SQL Server database for replication

In a test environment, you may occasionally need to reset the replication environment. Instead of deleting and re-creating a new Mirror Replication Agent instance, use this procedure to facilitate resetting the environment.

### ❖ To reset the primary database for replication

- 1 Protect your new environment from old log information by using the following command in the Mirror Replication Agent to:

```
pdb_init move_truncpt
```

- 2 To retain marking and log device information, reinitialize the Mirror Replication Agent using `ra_init` with the `force` option, which forces the Mirror Replication Agent repository to be refreshed instead of overwritten:

```
ra_init force
```

---

**Note** If you prefer to delete and replace all the information in the Mirror Replication Agent repository, issue the `ra_deinit` command followed by a normal `ra_init` command (without the `force` option).

---

- 3 Zero-out the locator stored in Replication Server using the `rs_zeroltm` command. For example:

```
isql -USAMPLE_RS_RSSD_prim -PSAMPLE_RS_RSSD_prim_ps
-SSAMPLE_RS_ERSSD -DSAMPLE_RS_ERSSD
```

```
rs_zeroltm mymr, pubs
go
```

The following is displayed:

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

| Topic                                   | Page |
|-----------------------------------------|------|
| System configuration                    | 73   |
| Prerequisites                           | 74   |
| Installing Mirror Activator components  | 76   |
| Configuring Mirror Activator components | 86   |
| Testing replication                     | 111  |

This chapter identifies the components that are required to implement a Sybase replication system for Oracle. This chapter contains an overview of the tasks and instructions for installing and configuring a sample replication environment for Oracle.

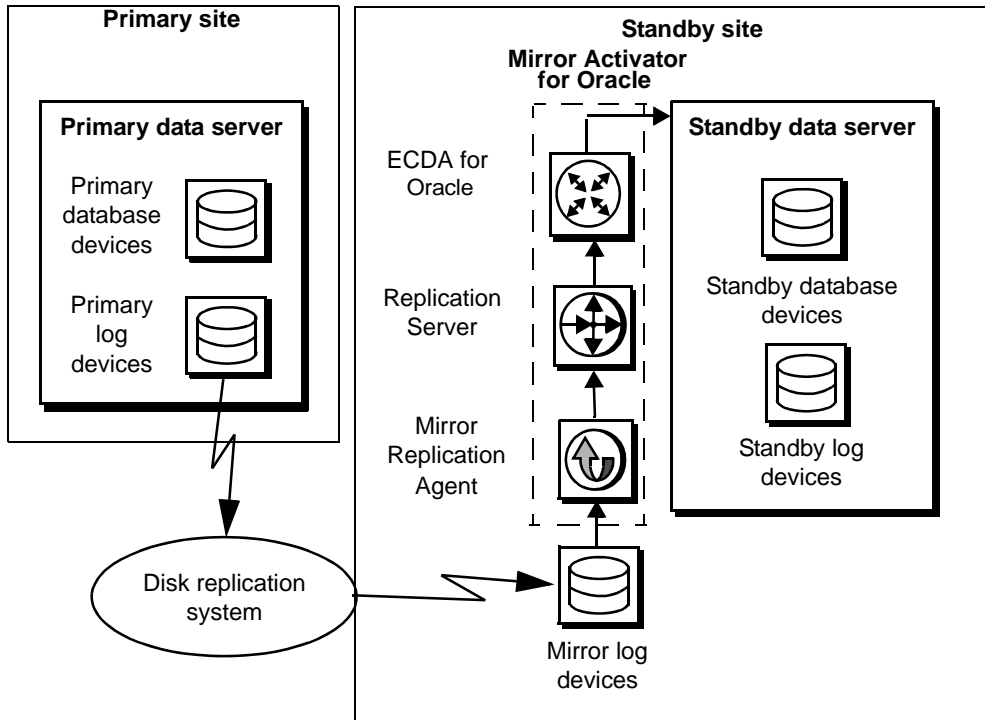
---

**Note** This procedure is only for proof-of-concept (POC) or testing, not for production.

---

## System configuration

Figure 3-1 shows a typical Mirror Replication Agent system configuration for Oracle. Arrows illustrate the flow of mirrored log device data and replicated transactions during normal Mirror Replication Agent system operation, that is, while replicating transactions from the primary database to the standby database.

**Figure 3-1: Sybase replication system with Oracle data servers**

## Prerequisites

Before you begin setting up the replication system, you should be familiar with Oracle data servers and have an understanding of Sybase replication. In addition, these conditions must be met:

- You must have a valid license for Replication Server and Mirror Activator.
- Oracle Enterprise Edition 9.2 or 10.2, is installed and configured, to serve as the source (primary) database and as the target (standby) database from which Mirror Replication Agent replicates transactions.
- Oracle JDBC thin driver for Oracle 10.2 and for JDK 1.4 and 1.5 is installed.
- TCP/IP connectivity is available.

- Any OS patches required for Java 5.0 have been installed.
- Mirror Replication Agent must have direct access to the Oracle online *redo* logs or a mirrored copy of those logs. Mirror Replication Agent must also have direct access to a copy of the Oracle archived *redo* logs or a mirrored copy of those logs.

During this procedure:

- Because some directories, files, executable commands, and examples are shown only for Sun Solaris, you must adjust accordingly for Windows and for other UNIX platforms.
- Sybase recommends that you do not use Replication Server reserved words for object names and connection names. A complete list of reserved words are documented in Chapter 2 of the Replication Server *Reference Manual*.

## System requirements

The target computer on which you are installing the Mirror Activator components must meet the following memory and disk space requirements:

**Table 3-1: Memory and disk requirements**

| Components               | Memory (RAM) | Disk space |
|--------------------------|--------------|------------|
| Replication Server       | 512MB        | 520MB      |
| ECDA                     | 256MB        | 700MB      |
| Mirror Replication Agent | 512MB        | 300MB      |

## Installing Mirror Activator components

This section explains pre-installation tasks and describes how to install and verify the installation of these components of Mirror Activator for Oracle:

- Mirror Replication Agent
- Replication Server
- ECDA Option for Oracle

---

**Note** The replication of sequences is not covered in this document. If you need to set up sequence replication, refer to the Mirror Replication Agent *Installation Guide* and the Mirror Replication Agent *Primary Database Guide*.

---

### Before you start

Before installing Mirror Activator components, identify the Sybase installation directories and obtain the installation software.

### Identifying the Sybase installation directory

Mirror Replication Agent, Replication Server, and ECDA are installed in the same base directory, which is identified by the SYBASE environment variable (\$SYBASE). Select a path on a host drive to be the recipient of the Sybase installation and configuration activities. For example:

`/software/sybase`

### Accessing products from the SPDC Web site

If you are not using installation CDs, you must go to the Sybase product Download Center (SPCD) Web site to obtain replication software.

❖ **To download the replication software**

---

**Note** This procedure requires you repeat steps to download three products: Replication Server, Mirror Activator, and Enterprise Connect Data Access.

---

You must have a Login ID and Password to download software from the SPDC Web site.



- 1 Go to SPDC at <https://sybase.subscribenet.com/control/sybs/login> and enter your Login ID and Password.  
The Product List page displays.
- 2 Select Mirror Activator.  
The detailed Product Information page displays.
- 3 Select the Mirror Activator software you need by primary database, version, and platform.  
The Software Terms and Conditions page displays, listing countries and regions.
- 4 Either:
  - Select I Agree to accept the terms without reading them, 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 displays. 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 go back to the high-level Product Information page.
- 5 The Product Download Page lists the software and documentation for the product. Select FTP Download for the Mirror Activator 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 SyBooks download contains the rest of the documentation set.

---

- 6 Uncompress and extract all the installation images to your local drive.

## Installing Mirror Replication Agent

This section describes how to install Mirror Replication Agent. It also has a procedure to verify the installation.

The Mirror Replication Agent installation process creates a directory structure containing the Mirror Replication Agent binaries, scripts, Java Runtime Environment (JRE), and associated files. There are no configuration steps during installation.

Before you install Mirror Replication Agent, see the Mirror Replication Agent *Installation Guide* and complete the installation worksheet.

### Installation guidelines

While installing Mirror Replication Agent, use the following guidelines:

- Verify that the operating system patch levels are at the current level required to support Java 5.0. This is necessary to support the JRE installed with Mirror Replication Agent.
- Install Mirror Replication Agent for Oracle on a server where it can directly access the Oracle online *redo* logs and the archived *redo* logs or access a mirrored copy of these logs.
- Create a Mirror Replication Agent instance and configure it for the Oracle database from which you want to replicate.

### ❖ To install Mirror Replication Agent software

- 1 Use *one* of the following methods to install the software.
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 76, go to the location where you extracted the Mirror Replication Agent 15.1 for Sun Solaris installation image, or
  - Insert the Mirror Replication Agent CD into your CD drive.
- 2 Launch the installation program by executing the setup routine for Solaris:
- 3 If you see an error message saying that there is not enough temporary disk space, add the following to your setup command:

```
./setup -is:tempdir someDirectory
```

Here, *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.

---

**Note** If you are installing from a remote machine and do not have an X-windows client, execute `./setup --console`, to run in console mode.

---

- 4 On the Welcome window, click Next.

- 5 Select a geographic location that displays the license agreement. Select I Agree and click Next.
- 6 Enter the directory where you want to install Mirror Replication Agent:
  - If the directory does not exist, the installation program prompts you to create it. Click Yes.
  - If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 7 Select Typical for the installation setup type and click Next.
- 8 The installation program displays a summary of what is to be installed. Click Next.
- 9 The installation program starts to install Mirror Replication Agent, displaying the progress status. If you see the following message;

```
There are newer files already installed and do you
want to replace them?
```

Click No to All.
- 10 When you receive a message that the installation was successful, click Next.
- 11 The SySAM License Server window opens and displays this prompt:

```
Will licenses be obtained from the License Server?
```

  - If you answer Yes, you can leave the *license-mrx.lic* license file where it was installed (*\$SYBASE/SYSAM-2\_0/licenses*), or you can move the file to the *SYSAM-2\_0/licenses* directory of your SYSAM license server. In either case, the Mirror Replication Agent will verify the necessary license.
  - If you answer No, no further action is needed. The Mirror Replication Agent will use the installed license (*license-mrx.lic*) as it was installed in *\$SYBASE/SYSAM-2\_0/licenses*.Click Next.
- 12 The SySAM notification window prompts you to 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 these values from the installation worksheet you completed in the Mirror Replication Agent *Installation Guide*:
  - SMTP server host name

- SMTP server port number
- E-mail return address
- Recipients of the notification
- Message severity level of an event that will trigger e-mail. Your choices are:
  - Informational
  - Warning
  - Error

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

Click Next.

- 13 When you receive a message indicating that Mirror Replication Agent was successfully installed, click Finish.

For more information for the replication installation process, see the Mirror Replication Agent *Installation Guide*.

❖ **To verify the installation**

- 1 Change to the directory where you installed Mirror Replication Agent.
- 2 Set the environment variable by sourcing the *SYBASE.csh* file.
- 3 Change the directory to *\$SYBASE/MA-15\_1/bin*.
- 4 Obtain the Mirror Replication Agent version string:

```
./ma.sh -v
```

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

## Installing Replication Server

This section describes the steps for a Replication Server installation and a procedure to verify the installation.

---

**Note** If Replication Server is already installed go to the next section called “Installing the ECDA Option for Oracle” on page 84.

---

**Installation guidelines**

While installing Replication Server, use the following guidelines:

- 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.

**❖ To install Replication Server**

- 1 Use *one* of the following methods to download Replication Server software.
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 76, go to the location where you extracted the Replication Server 15.1 for Sun Solaris installation image, or
  - Insert the Replication Server CD in your CD drive.
- 2 Launch the installation program by executing the setup routine for Solaris:  

```
./setup
```
- 3 If you see an error message saying that there is not enough temporary disk space, add the following to your setup command:  

```
./setup -is:tempdir someDirectory
```

Here, *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.
- 4 On the Welcome window, click Next.
- 5 Select a geographic location that displays the license agreement. Select I agree and click Next.
- 6 Enter the directory where you want to install Replication Server:
  - If the directory does not exist, the installation program prompts you to create it. Click Yes.
  - If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.
- 7 Select Typical for the installation setup type and click Next.
- 8 The installation program displays a summary of what is to be installed. Click Next.

- 9 The installation program starts to install Replication Server, displaying the progress status. You may see the following message,

There are newer files already installed and do you want to replace them?

Click No to All.

- 10 When you receive a message indicating that Replication Server software has successfully installed, click Next.

---

**Note** This procedure uses the SAMPLE\_RS Replication Server.

---

- 11 The SySAM License Server window opens and displays this prompt:

Will licenses be obtained from the License Server?

Respond based on how you set up your environment and use the license obtained from the Software Product Download (SPDC) Web site for Replication Server.

Click Next.

- 12 The SySAM notification window prompts you to 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 your own values for the following fields:

- SMTP server host name
- SMTP server port number
- Sender e-mail
- Recipient e-mail
- Message severity level of an event that triggers e-mail. Your choices are:
  - Informational
  - Warning
  - Error

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

Click Next.

13 To start a sample Replication Server, select Yes and click Next.

14 The following information is displayed:

This page contains detailed information regarding the sample Replication Server. Please record this information.

The sample Replication Server will be called `SAMPLE_RS` and will run on port 11752. It will be configured with a user of `sa` and no password.

The sample Replication Server will use an embedded RSSD called `SAMPLE_RS_ERSSD` that runs on port 11751. It will be configured with a user of `SAMPLE_RS_RSSD_prim` and a password of `SAMPLE_RS_RSSD_prim_ps`.

The installer has updated the appropriate *interfaces* file or *sql.ini* file.

All files and logs associated with the sample Replication Server will be located in the directory `$SYBASE/REP-15_1/samp_repserver`.

The sample Replication Server will be configured using the file `$SYBASE/REP-15_1/samp_repserver/SAMPLE_RS.res`.

Keep this information for future reference. Click Next.

15 When the `SAMPLE_RS` is running and an “Installation was successful” message appears, click Finish.

❖ **To verify the installation**

1 Change to the directory (`$SYBASE`) where you installed Replication Server.

2 Set the environment variables by sourcing the `SYBASE.csh` file.

3 Use `isql` to log in to Replication Server:

```
isql -Usa -P -SSAMPLE_RS
```

You should successfully log in to the Replication Server.

4 Exit `isql`.

## Installing the ECDA Option for Oracle

This section describes how to install the ECDA Option for Oracle.

---

**Note** If ECDA Option for Oracle was already installed, go to the next section called “Configuring Mirror Activator components.”

---

When installing ECDA for Oracle, be sure to install ECDA on the same server as the standby database. This eliminates a network hop and improves performance.

### ❖ To install ECDA for Oracle in GUI mode

---

**Note** If you want to install using the console mode or a response file, refer to the ECDA *Installation Guide* for Linux and UNIX.

---

- 1 Use *one* of the following methods to download the software:
  - After following the download instructions from the SPDC Web site in the section “Accessing products from the SPDC Web site” on page 76, go to the location where you extracted the ECDA Option for Oracle 15.0 installation image.
  - Insert the CD labeled ECDA 15.0 Option for Oracle for Sun Solaris.

- 2 Launch InstallShield by executing the Sun Solaris setup routine:

```
./setup
```

- 3 If you see an error message saying that there is not enough temporary disk space, add the following to your setup command:

```
./setup -is:tempdir someDirectory
```

Here, *someDirectory* is a directory with at least 100MB of disk space that the installation program can use.

- 4 On the Welcome window, click Next.
- 5 Select a geographic location that displays the license agreement. Select I agree and click Next.
- 6 Enter the directory where you want to install ECDA:
  - In the destination directory, enter the directory where the Mirror Replication Agent software is installed. Click Yes to install into an existing directory.



- If the destination directory exists, you receive a warning message that you are installing into an existing directory. Click Yes.

- 7 Select Custom setup, and click Next.
- 8 Deselect ECDA Option for ODBC, and click Next.
- 9 Verify the modules to install. Click Next.

If you receive a message stating that your current installation contains newer files than those that are about to be installed and asking you if you want to replace these files, click No.

- 10 When you receive a message that the installation was successful, click Next.

- 11 The SySAM License Server window opens and displays this prompt:

Will licenses be obtained from the License Server?

Respond based on how you set up your environment and use the license obtained from the Software Product Download Center (SPDC) Web site for Mirror Activator.

Click Next.

- 12 The SySAM notification window prompts you to 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 your own values for the following fields:

- SMTP server host name
- SMTP server port number
- Sender e-mail
- Recipient e-mail
- Message severity level of an event that triggers e-mail. Your choices are:
  - Informational
  - Warning
  - Error

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

- 13 When you receive a message indicating that ECDA was successfully installed, click Finish.

ECDA, including the DirectConnect server, is installed.

## Configuring Mirror Activator components

This section describes the steps, in order, for configuring the components for Mirror Replication Agent for Oracle:

- 1 Configuring ECDA Option for Oracle
- 2 Configuring Replication Server
- 3 Configuring Mirror Replication Agent

### Configuring ECDA Option for Oracle

Configuring ECDA Option for Oracle involves the following sequence of tasks:

- 1 Configuring and starting the DirectConnect server
- 2 Creating a maintenance user and DDL user for replication
- 3 Creating objects and public synonyms
- 4 Verifying the ECDA installed objects

### Configuring and starting the DirectConnect server

Before you begin to create and configure the DirectConnect server, make sure that:

- You have made a copy of the existing *tnsnames.ora* file, currently being used to connect to Oracle, and have placed it in a temporary file location.
- The Oracle connect string is available.
- The name for a valid Oracle account, which is used as the administrator for the DirectConnect server, is available.
- The number of the unused port to be used by the DirectConnect server is available.

**❖ Creating and configuring a DirectConnect server**

- 1 Start the Create Server wizard navigate to the `$SYBASE/DCO-15_0/DCWizard` directory and execute `DCWizard.sh` script  
  
The Welcome Create Server Wizard window appears. Click Next.
- 2 Select ECDA Option for Oracle. Click Next.
- 3 Enter the server name for the new DirectConnect server and the port number that you want the server to listen on. Click Next.
- 4 Enter a valid Admin Account name. Click Next.
- 5 Enter the Oracle connection string as previously defined in the `tnsnames.ora` file, and the path (including the file name) to the previously defined temporary location of the `tnsnames.ora` file. Click Next.
- 6 Verify the DirectConnect server information. If correct, select Create Server otherwise, click Back to return to the previous window and provide the correct information.
- 7 Verify if the DirectConnect configuration is successful. If yes, click Next otherwise, click Back to return to the previous window, provide the correct information, and recreate the server.
- 8 In the Start DirectConnect Server window, click Start `ecda_servername` to start the newly configured DirectConnect server, where `ecda_servername` is the server name you indicated for the new DirectConnect server.
- 9 Verify that the DirectConnect server startup is successful and click Finish.

**❖ To verify that you can connect to Oracle**

- 1 Open a command window in the `$SYBASE` directory of your ECDA installation.
- 2 Change to the `$SYBASE/DCO-15_0` directory.
- 3 Set the environment variables by sourcing the `$SYBASE/DCO_SYBASE.csh` file.
- 4 Use the `isql` utility to log in to Oracle through the DirectConnect server:  
  

```
isql -Uvaliduser -Ppassword -SDCOServer
```

  
Here, `validuser` and `password` are any valid user and password, and `DCOServer` is the DirectConnect for Oracle server name.
- 5 Verify the connection to the standby Oracle database:

```
select banner from v$version
go

BANNER

Oracle Database 10g Enterprise Edition Release
10.2.0.1.0 - Prod
PL/SQL Release 10.2.0.1.0 - Production
CORE 10.2.0.1.0 Production
TNS for Solaris Version 10.2.0.1.0 - Production
NLSRTL Version 10.2.0.1.0 - Production
(5 rows affected)
```

- 6 Exit the isql session with the quit command.

## Creating a maintenance user and DDL user for replication

Description of a  
maintenance user

This section explains how to create a maintenance user in Oracle.

The maintenance user is a valid Oracle user that the Replication Server uses to apply commands to the standby Oracle database. Replication Server requires one maintenance user to be defined for each connection. Do not use the same name as the DCO Admin Account Name used in configuring the DirectConnect server.

### ❖ To create a maintenance user in Oracle

- 1 Using SQLPLUS, connect to the standby database and enter:

```
create user maintuser identified by password;
```

- 2 Grant dba and create session permissions to the maintenance user, who must have permission to issue SQL commands against any table to be replicated:

```
grant dba to maintuser;
grant create session to maintuser;
```

### ❖ To create a DDL user in Oracle

- 1 Create a DDL user in the standby database:

```
create user ddl_user identified by password;
```

---

**Note** The value of the ddl\_username parameter must not be the same as the value of the maintenance user defined in Replication Server for the standby connection.

---

- 2 Grant permissions to the DDL user, which must have privileges to perform DDL activities and issue any DDL command on the correct database.

---

**Note** The user ID for the DDL user must have authority in the standby database to issue the

`ALTER SESSION SET CURRENT_SCHEMA=user`

command for any user who can create DDL in the primary database.

---

## Creating objects and public synonyms

### ❖ To create objects and public synonyms

- 1 Sign on to Oracle through ECDA Option for Oracle as the maintenance user (assuming that the maintenance user that you created has the resource role to create tables).
- 2 Create the `rs_info` and `rs_lastcommit` tables in the Oracle database by executing this script:

```
isql -Umaintuser -Ppassword -SDCOserver
-i $SYBASE/MA-15_1/scripts/oracle
/hds_oracle_new_setup_for_replicate.sql
```

Here, *DCOserver* is the name of the DirectConnect server.

---

**Note** You may receive a ORA-00942 message that a table or view does not exist. Ignore the message.

---

## Verifying the ECDA installed objects

### ❖ To verify the ECDA installed objects

- 1 Open another command window in the `$SYBASE` directory of your ECDA installation.
- 2 Set the environment variables by sourcing the `$SYBASE/DCO-15_0/DCO_SYBASE.csh` file.
- 3 Use the `isql` utility to log in to Oracle through ECDA as *maintuser*.
  - To verify whether the `rs_info` table has been created successfully:

```
isql -Umaintuser -Ppassword -SDCOserver
select * from rs_info
go
```

This is returned:

```

RSKEY RSVAL

charset_name iso_1
sortorder_name bin_iso_1
(2 rows affected)

```

- To verify whether the rs\_lastcommit table has been created successfully:

```

isql -Umaintuser -Ppassword -SDCOserver
select * from rs_lastcommit
go

```

The following is returned:

```

ORIGIN ORIGIN_QID SECONDARY_QID ORIGIN_TIME DEST_COMMIT_TIME

(0 rows affected)

```

If the ECDA objects are not installed successfully, no results are displayed, and you encounter an error message.

- 4 Exit the isql session with the quit command.

## Configuring Replication Server

Configuring Replication Server includes the following sequence of tasks:

- 1 Configuring Replication Server for Oracle replication
- 2 Configuring Replication Server for replication to the target database
- 3 Configuring Replication Server for replication from the primary database

### Configuring Replication Server for Oracle replication

This section describes the configuration tasks for configuring Replication Server for Oracle.

#### ❖ To apply heterogeneous datatype support (HDS) scripts to RSSD

Prepare the Replication Server RSSD for replication to an Oracle database. To apply HDS support scripts to RSSD:

- 1 Set the environment variables by sourcing *SYBASE.csh* file.
- 2 Change to the *SYBASE/REP-15\_1/scripts* directory:

- 3 Load the following scripts into the RSSD:

```
isql -U SAMPLE_RS_RSSD_prim -P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ihds_clt_ase_to_oracle.sql -DSAMPLE_RS_ERSSD
```

```
isql -U SAMPLE_RS_RSSD_prim -P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ihds_oracle_udds.sql -DSAMPLE_RS_ERSSD
```

```
isql -U SAMPLE_RS_RSSD_prim -P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ihds_oracle_funcstrings.sql -DSAMPLE_RS_ERSSD
```

---

**Note** The message “ASA -157 cannot convert to a timestamp” is displayed. You can ignore this message.

---

- 4 Go to the `$$SYBASE/MA-15_1/scripts/oracle` directory.
- 5 Load this script into the RSSD:

```
isql -U SAMPLE_RS_RSSD_prim
-P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ihds_oracle_new_udds.sql -DSAMPLE_RS_ERSSD
```
- 6 Shut down Replication Server:

```
isql -Usa -P -SSAMPLE_RS
shutdown
go
```
- 7 Navigate to the `$$SYBASE/REP-15_1/samp_repserver` directory:
- 8 Run the `RUN_SAMPLE_RS` script, to start the `SAMPLE_RS` Replication Server.

## Configuring Replication Server for replication to the target database

This section describes the required configurations for Replication Server for replication to the target Oracle database.

### ❖ To create a Replication Server connection to replicate

- 1 Navigate to the Mirror Replication Agent scripts `$$SYBASE/MA-15_1/scripts/oracle` directory.
- 2 Create the Oracle error class with the following scripts.

---

**Note** The `oracle_error_class_1_rs.sql` and `oracle_error_class_3_rs.sql` scripts are executed to the Replication Server. The `oracle_error_class_2_rssd.sql` script is executed in the RSSD.

---

```
isql -Usa -P -SSAMPLE_RS -ioracle_create_error_class_1_rs.sql
isql -U SAMPLE_RS_RSSD_prim -P SAMPLE_RS_RSSD_prim_ps -SSAMPLE_RS_ERSSD
-ioracle_create_error_class_2_rssd.sql -DSAMPLE_RS_ERSSD
isql -Usa -P -SSAMPLE_RS -ioracle_create_error_class_3_rs.sql
```

- 3 Make a copy of the *oracle\_create\_rs\_standby\_connection.sql* script:

```
cp oracle_create_rs_standby_connection.sql
my_oracle_create_rs_standby_connection.sql
```

- 4 Before executing the *my\_oracle\_create\_rs\_standby\_connection.sql* script against your Replication Server, change all occurrences of value *{rds}.**{rdb}* to the name of the connection that Replication Server uses to connect to DirectConnect server, where:

- *rds* is the DirectConnect server name.
- *rdb* is any valid identifier. Sybase recommends that you use the Oracle SID name.
- *rs\_maint\_user* and *rs\_maint\_user\_pwd* are the maintenance user and password created in ECDA Option for Oracle, in step 2 of “Configuring ECDA Option for Oracle” on page 86.
- *maintuser* and *password* are the user name and password created in ECDA Option for Oracle.

For example:

```
create connection to DCOServer.oratest2
set error class oracle_error_class
set function string class rs_oracle_function_class
set username maintuser
set password "password"
set batch to "off"
go
```

---

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

---

- 5 Create the connection to the standby database:

```
isql -Usa -P -SSAMPLE_RS -i my_oracle_create_rs_standby_connection.sql
```

- 6 To verify the Replication Server connection to the standby database, log in to *SAMPLE\_RS* and verify the Replication Server connection to the standby database:



```
isql -Usa -P -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.”

- 7 Exit the isql session with the quit command.

## Configuring Replication Server for replication from the primary database

This section describes the configuration steps required for replication from the primary database. It includes:

- Creating a Replication Server connection to the primary database
- Creating a database replication definition
- Creating the database replication subscription

### ❖ To create a Replication Server connection to the primary database

- 1 Go to the Mirror Replication Agent scripts, *\$SYBASE/MA-15\_1/scripts/oracle* directory:
- 2 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
```
- 3 Before executing the *my\_oracle\_create\_rs\_primary\_connection.sql* script against your Replication Server, change all occurrences of value *{pds}.{pdb}* to the name of the Replication Server connection used for the connection from Mirror Replication Agent for Oracle, where:
  - *pds* is *rs\_source\_ds* (located in *\$SYBASE/MA-15\_1/init/oracle.rs*).
  - *pdb* is *rs\_source\_db* (located in *\$SYBASE/MA-15\_1/init/oracle.rs*).

---

**Note** You must save the defined values to be used for the Mirror Replication Agent *oracle.rs* file later in the process.

---

For example, *NY.NYora92*.

- 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.

---

The following command creates a Replication Server connection to the primary database example:

```
create connection to NY.NYora92
set error class rs_sqlserver_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 -Usa -P -SSAMPLE_RS -i my_oracle_create_rs_primary_connection.sql
```

A message is displayed that indicates the Replication Server connection to the primary database is created.

❖ **To create the database replication definition**

- 1 Navigate to the *\$SYBASE/MA-15\_1/scripts/sybase* directory:
- 2 Make a copy of the *rs\_create\_test\_db\_repdef.sql* script:

```
cp rs_create_test_db_repdef.sql
my_rs_create_test_db_repdef.sql
```

- 3 Before executing the *my\_rs\_create\_test\_db\_repdef.sql* script, change the value of “*{pds}.{pdb}*” to the name of the connection string you defined for the primary database, where:

- *pds* is *rs\_source\_ds* (located in *\$SYBASE/RAX\_15\_1/init/oracle.rs*).
- *pdb* is *rs\_source\_db* (located in *\$SYBASE/RAX\_15\_1/init/oracle.rs*).

For example, *NY.NYora92*

This is a database replication definition example:

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

- 4 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imy_rs_create_test_db_repdef.sql
```

A message is displayed that indicates the database replication definition is created.

❖ **To create the database replication subscription**

- 1 Make a copy of the *rs\_create\_test\_db\_sub.sql* script:

```
cp rs_create_test_db_sub.sql
my_rs_create_test_db_sub.sql
```

- 2 Before executing the *my\_rs\_create\_test\_db\_sub.sql* script, change the “{pds}.{pdb}” and “{rds}.{rdb}” to the appropriate connection name, where:
  - *pds* is *rs\_source\_ds* (located in *\$SYBASE/RAX\_15\_1/init/oracle.rs*).
  - *pdb* is *rs\_source\_db* (located in *\$SYBASE/RAX\_15\_1/init/oracle.rs*).
  - *rds* is the DirectConnect server name
  - *rdb* is any valid identifier. (Sybase recommends that you use the Oracle SID name.)

Here is an example of create database replication subscription:

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

- 3 Execute the script in Replication Server:

```
isql -Usa -P -SSAMPLE_RS -imy_rs_create_test_db_sub.sql
```

A message is displayed that indicates the subscription is in the process of being created.

- 4 Use *isql* to log in to the Replication Server:

```
isql -Usa -P -SSAMPLE_RS
```

- 5 Replication Server comes with users who have the permissions necessary for Mirror Replication Agents. However, because Mirror Activator has the capability to create replication definitions, you must grant CREATE OBJECT permission to the Mirror Replication Agent user:

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

Permission is granted to user SAMPLE\_RS\_ra.

---

**Note** This user will be entered in the `rs_username` parameter when building the Mirror Replication Agent for Oracle instance. If you want a user other than the default replication agent user, make sure that the user is granted `CONNECT SOURCE` and `CREATE OBJECT` permission.

---

## Configuring Mirror Replication Agent

This section describes the steps, in order, for configuring Mirror Replication Agent for Oracle:

- 1 Configuring and verifying the primary Oracle database for replication
- 2 Creating the Mirror Replication Agent instance
- 3 Verify the Mirror Replication Agent instance
- 4 Initializing the Mirror Replication Agent instance

### Configuring and verifying the primary Oracle database for replication

Before you install Mirror Replication Agent, configure the primary Oracle database for replication.

To ensure that the Oracle *redo* log data is not overwritten before it is read by the Mirror Replication Agent, you must complete the following tasks:

- Verify the current archiving setting of the redo logs.
- Verify the supplemental logging of primary key data.
- For Oracle 10g, verify that the flashback feature is disabled.
- Create an Oracle user and granting Oracle permissions.

#### ❖ To verify the current archiving setting of the redo logs

*Redo* logs are used by Oracle to maintain a log infrastructure. The steps described below verify if the feature is turned on and, if not, indicate how to switch on the *redo* logging. This feature needs to be turned on only on the primary database.

- 1 Use `SQLPLUS` to connect to Oracle as a system administrator.
- 2 Run this command from `SQLPLUS`:

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

- If the archive log is on, the result should be:

```
LOG_MODE

ARCHIVELOG
```

- 3 To turn on log archiving, run these commands:

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

For more information about Oracle *redo* logs and archiving, see the Mirror Replication Agent *Primary Database Guide*.

#### ❖ To verify the supplemental logging of primary key data

By default, Oracle does not log primary keys and unique indexes to its logging infrastructure. You must include the logging of these values for a successful replication of all table values. To verify whether this information is currently being recorded in the logging system:

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

- For Oracle 9, run this command from SQLPLUS:

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

For Oracle 9, if logging of primary key and unique index values is enabled, the return values should be:

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

- For Oracle 10, run this command from SQLPLUS

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

For Oracle 10, if logging of primary key and unique index values is enabled, the return values should be:

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

- 2 If the result is different, turn on supplemental logging by executing these commands for Oracle 9 and Oracle 10:

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

❖ **To verify that the Oracle 10g flashback feature is disabled**

- 1 Use SQLPLUS to connect to Oracle as a system administrator with sysdba privileges.

- 2 View the current recycle bin configuration:

```
select inst_id, value from v$parameter;
```

- 3 View the contents of the recycle bin:

```
select * from dba_recyclebin;
```

- 4 Disable the recycle bin with the following commands:

```
PURGE dba_recyclebin;
ALTER SYSTEM SET recyclebin = OFF;
```

---

**Note** If you are using Oracle RAC, you must disable the recycle bin for each instance in the cluster.

---

❖ **To create an Oracle user and grant permissions**

- An Oracle database user must be created for use by Mirror 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 the following commands to create an Oracle user named “RA\_USER” 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 "CONNECT" TO "RA_USER";
GRANT "RESOURCE" TO "RA_USER";
GRANT "SELECT_CATALOG_ROLE" TO "RA_USER";
GRANT ALTER SESSION TO "RA_USER";
```

```
GRANT ALTER SYSTEM TO "RA_USER";
GRANT EXECUTE ON "SYS"."DBMS_FLASHBACK" TO
"RA_USER";
GRANT ALTER ANY PROCEDURE TO "RA_USER";
GRANT CREATE SESSION TO "RA_USER";
GRANT CREATE TABLE TO "RA_USER";
GRANT ALTER ANY TABLE TO "RA_USER";
GRANT DROP ANY TABLE TO "RA_USER";
GRANT CREATE PROCEDURE TO "RA_USER";
GRANT DROP ANY PROCEDURE TO "RA_USER";
GRANT CREATE PUBLIC SYNONYM TO "RA_USER";
GRANT DROP PUBLIC SYNONYM TO "RA_USER";
GRANT SELECT ON SYS.OBJ$ TO "RA_USER";
GRANT SELECT ON SYS.LOB$ TO "RA_USER";
GRANT SELECT ON SYS.COLLECTION$ TO "RA_USER";
GRANT SELECT ON SYS.CON$ TO "RA_USER";
GRANT SELECT ON SYS.COLTYPE$ TO "RA_USER";
GRANT SELECT ON SYS.COL$ TO "RA_USER";
GRANT SELECT ON SYS.CDEF$ TO "RA_USER";
GRANT SELECT ON SYS.USER$ TO "RA_USER";
GRANT SELECT ON SYS.SEQ$ TO "RA_USER";
GRANT SELECT ON SYS.IND$ TO "RA_USER";
```

❖ **To verify the Oracle user roles created by the Mirror Replication Agent**

- As described in the previous section, the following Oracle commands can be issued to validate the required Oracle settings.

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

```
select GRANTED_ROLE from USER_ROLE_PRIVS;
GRANTED_ROLE

CONNECT
RESOURCE
SELECT_CATALOG_ROLE
```

In addition, the user who starts the Mirror Replication Agent instance must 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 Mirror 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 Mirror Replication Agent instance

A single installation of the Mirror Replication Agent can support replication from multiple databases, however, one Mirror Replication Agent instance is needed for each Oracle database that is to be replicated. This section describes how to prepare for the creation of an instance.

### ❖ To prepare the Mirror Replication Agent instance

- 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, usually found at

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

Mirror Replication Agent may require a newer version of the JDBC driver, which you can obtain at

[http://technet.oracle.com/software/tech/java/sqlj\\_jdbc/content.html](http://technet.oracle.com/software/tech/java/sqlj_jdbc/content.html).

- 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 the following command to the .login script of a UNIX and Linux user:

```
setenv CLASSPATH /path_name/ojdbc14.jar:$CLASSPATH
```

For a Windows user, add the following to the CLASSPATH environment variable:

```
set CLASSPATH=path_name\jdbc\lib\ojdbc14.jar;
%CLASSPATH%
```

- 3 Record the location of the *tnsnames.ora* file and record the connection name from that file for the Mirror Replication Agent to use to connect to the Oracle primary database. You can find these values from any machine where an Oracle SQLPLUS session can be used to connect to the desired Oracle instance. The ORACLE\_SID is from the environment variable, \$ORACLE\_SID.

Obtain host and port information from the file called *tnsnames.ora*, available at *\$ORACLE\_HOME/network/admin*.

Record the following:

- Host name of the Oracle database on which the TNS listener is operating.



- Port number the TNS listener is monitoring.
  - ORACLE\_SID value for the instance you wish to connect to.
- 4 (Optional) If your operating system has a Java 5.0.x-compatible JRE or JDK installed, you can use Oracle's iSQL (the browser-based interface to SQL\*PLUS) demo items to connect to Oracle using the JDBC driver and to verify the connection information.

To perform this validation, 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 Mirror Replication Agent can correctly process the Oracle timestamp with timezone datatype.

---

**Note** This step is required only if the Mirror 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 under the Oracle installation directory:

- The *timezone.dat* binary file is the default file that contains the most commonly used time zones. This file is located in the `$ORACLE_HOME/oracore/zoneinfo` directory.
- The *timezlrg.dat* binary file contains a larger set of defined time zones. This file is located in the `$ORACLE_HOME/oracore/zoneinfo` directory.

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 Mirror 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 Mirror Replication Agent for Oracle resource file template.

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

The resource file template (*oracle.rs*) can be found in the Mirror Replication Agent directory *\$SYBASE/MA-15\_1/init*.

#### 7 Create an instance resource file.

Copy the resource file template *\$SYBASE/MA-15\_1/init/oracle.rs* to another file that contains the configuration information for a new instance, for example:

```
cp oracle.rs myma.rs
```

Using a text editor, 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 tnsname location, user, and password information for the Oracle, Replication Server, and RSSD connections.

The resource file is self-documenting, with each parameter described. For more information, refer to the Mirror Activator *Administration Guide*, in Chapter 2, “Setup and Configuration,” the section titled, “Creating a Mirror Replication Agent instance using resource files.”

**Table 3-2: Resource file parameters**

| Parameter         | Description                                                                                                                                                     | Example values                                   |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------|
| instance_name     | Any valid name.                                                                                                                                                 | myma                                             |
| admin_port        | Port number that Mirror Replication Agent uses.                                                                                                                 | 9030 (if in use, select a different port number) |
| pds_host_name     | Machine (host) where Oracle is installed.<br><br><b>Note</b> To be used with pds_port_number, only if you have not set pds_tns_connection and pds_tns_filename. | jdoe_host1                                       |
| pds_port_number   | Port number for Oracle.<br><br><b>Note</b> To be used with pds_host_name, only if you have not set pds_tns_connection and pds_tns_filename.                     | 1521                                             |
| pds_database_name | \$ORACLE_SID                                                                                                                                                    | NYora92                                          |

| Parameter           | Description                                                                                                                                                                                                                                                                                                                                                   | Example values                                                                                                           |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| pds_tns_connection  | <p>Connection name found in the <i>tnsnames.ora</i> file which identifies the connection information for the primary database.</p> <p><b>Note</b> To be used with <i>pds_tns_filename</i>, only if you have not set <i>pds_host_name</i> and <i>pds_port_number</i>.</p>                                                                                      | ORA102.JDOE_HOST.COM                                                                                                     |
| pds_tns_filename    | <p>File name identifying the Oracle <i>tnsnames.ora</i> file to be used to identify the connection information for the primary database.</p> <p><b>Note</b> To be used with <i>pds_tns_connection</i>, only if you have not set <i>pds_host_name</i> and <i>pds_port_number</i>.</p>                                                                          | <p>A valid Oracle <i>tnsnames.ora</i> file.</p> <p>For example:</p> <p><i>/opt/oracle/network/admin/tnsnames.ora</i></p> |
| pds_username        | User ID that Mirror Replication Agent uses to access primary data server.                                                                                                                                                                                                                                                                                     | RA_USER                                                                                                                  |
| pds_password        | Password for <i>pds_username</i> .                                                                                                                                                                                                                                                                                                                            | sybase                                                                                                                   |
| create_pds_username | Specifies whether or not the <i>ma_admin</i> utility, invoked without the <i>-r</i> option, should create the Mirror Replication Agent primary database user login if the login does not already exist in the primary data server. For information about the <i>ma_admin</i> utility <i>-r</i> option, see the Mirror Activator <i>Administration Guide</i> . | yes                                                                                                                      |
| 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 capabilities.                                                                                                                                                                                                                                                                                   | SAMPLE_RS_ra                                                                                                             |
| rs_password         | Password for <i>rs_username</i> .                                                                                                                                                                                                                                                                                                                             | SAMPLE_RS_ra_ps                                                                                                          |
| rs_charset          | <p>Character set that Replication Server is using.</p> <p><b>Note</b> The value defined for the <i>rs_charset</i> configuration parameter MUST match the <i>RS_charset</i> value in Replication Server's configuration file, <i>\$SYBASE/REP-15_1/install/&lt;server&gt;.cfg</i>.</p>                                                                         | <ul style="list-style-type: none"> <li>Windows: cp850</li> <li>UNIX: iso_1</li> </ul>                                    |
| rs_source_ds        | Valid name representing data server of Oracle primary database.                                                                                                                                                                                                                                                                                               | NY                                                                                                                       |
| rs_source_db        | Valid name representing Oracle primary database.                                                                                                                                                                                                                                                                                                              | NYora92                                                                                                                  |

| Parameter            | Description                                                                                                                                                                                                                                                                                                                              | Example values                                                                                                                                     |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| rssd_host_name       | Machine where RSSD resides.                                                                                                                                                                                                                                                                                                              | jdoe_host1                                                                                                                                         |
| rssd_port_number     | Port number where RSSD resides.                                                                                                                                                                                                                                                                                                          | 11751                                                                                                                                              |
| rssd_database_name   | Database name for RSSD.                                                                                                                                                                                                                                                                                                                  | SAMPLE_RS_ERSSD                                                                                                                                    |
| rssd_username        | Valid user for RSSD.                                                                                                                                                                                                                                                                                                                     | SAMPLE_RS_RSSD_maint                                                                                                                               |
| rssd_password        | Password for rssd_username.                                                                                                                                                                                                                                                                                                              | SAMPLE_RS_RSSD_maint_ps                                                                                                                            |
| ddl_username         | Name of the user who will apply DDL at the target database.                                                                                                                                                                                                                                                                              | ddl_user                                                                                                                                           |
| ddl_password         | Password for ddl_username.                                                                                                                                                                                                                                                                                                               | password (created previously)                                                                                                                      |
| pdb_timezone_file    | Path to the pdb_timezone_file directory that contains the archive redo log files.                                                                                                                                                                                                                                                        | /software/oracle/Ora10g<br>/oracore/zoneinfo<br>/timezone.dat                                                                                      |
| start_instance       | Start the instance that was created.                                                                                                                                                                                                                                                                                                     | yes                                                                                                                                                |
| initialize_instance  | Initialize the Mirror Replication Agent instance.<br><br><b>Note</b> Set the initialize_instance parameter to yes (the default is no) to initialize the Mirror Replication Agent instance. If you retain the default value of no, you must manually issue the ra_devicepath command after initialization to set the correct mirror path. | yes                                                                                                                                                |
| pdb_include_archives | Enables or disables the use of Oracle archive log files.                                                                                                                                                                                                                                                                                 | USE_DEFAULT                                                                                                                                        |
| pdb_archive_path     | Identifies directory path where Mirror Replication Agent expects to find archived Oracle redo log files.                                                                                                                                                                                                                                 | A valid directory path on the machine hosting Mirror Replication Agent that points to a location where Oracle puts the archived redo log files     |
| pdb_archive_remove   | Enables or disables the removal of archived Oracle redo log files from the path specified by pdb_archive_path.                                                                                                                                                                                                                           | USE_DEFAULT                                                                                                                                        |
| asm_tns_connection   | Identifies the Oracle ASM connection name found in the tnsnames.ora file.                                                                                                                                                                                                                                                                | A valid Oracle <i>tnsnames.ora</i> connection name<br><br><b>Note</b> To be used only if your Oracle environment uses ASM to manage the redo logs. |

| Parameter        | Description                                                                                               | Example values                                                                                                                           |
|------------------|-----------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| asm_tns_filename | Identifies the Oracle ASM <i>tnsnames.ora</i> filename where the connection information is located.       | A valid Oracle <i>tnsnames.ora</i> file.<br><br><b>Note</b> To be used only if your Oracle environment uses ASM to manage the redo logs. |
| asm_username     | Identifies the Oracle user name to be used when connecting to an ASM server.                              | User name used for ASM access<br><br><b>Note</b> To be used only if your Oracle environment uses ASM to manage the redo logs.            |
| asm_password     | Password for Oracle automatic Storage Management (ASM) access for the user specified in the asm_username. | Password used for the ASM user name<br><br><b>Note</b> To be used only if your Oracle environment uses ASM to manage the redo logs.      |

---

**Warning!** The `rs_source_ds` and the `rs_source_db` values must match the “pds.pdb” values of your Replication Server primary connection name that you configured in step 3, in the procedure, “To create a Replication Server connection to the primary database” on page 93.

---

## 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 Using the Mirror Replication Agent `ma_admin` utility, validate the settings in the resource file using the `-vr` parameter, for example:

```
$SYBASE/MA-15_1/bin/ma_admin.sh -vr myma.rs
```

Here, *myma.rs* is the path name of the resource file. Validation results are returned as one of the following:

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

If any validation fails, the `ma_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 because of this process.

- b After the resource file has been validated, allow the `ma_admin` utility to create and configure the Mirror Replication Agent instance, using the `-r` parameter, for example:

```
$SYBASE/MA-15_1/bin/ma_admin.sh -r myma.rs
```

---

**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`

For more information about the resource file or its use, refer to the *Mirror Activator Administration Guide*, Chapter 2, “Setup and Configuration,” in the section titled, “Creating a Mirror Replication Agent instance using resource files.”

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

Change to the instance directory and run Mirror Replication Agent in the background:

```
cd $SYBASE/MA-15_1/myma
```

Execute the *RUN* file in the background, for example:

```
./RUN_myma &
```

Your Mirror Replication Agent for Oracle is now running.

## Verify the Mirror Replication Agent instance

This section describes how to verify a Mirror Replication Agent instance, which includes the following sequence of tasks:

- 1 Update the Replication Server *interfaces* file with the Mirror Replication Agent location.
- 2 Verify the connection to the Mirror Replication Agent.
- 3 Verify the Mirror Replication Agent connection to Oracle.

❖ **To verify the Mirror Replication Agent**

- 1 Using `dsedit`, update the Replication Server *interfaces* file to include an entry for the Mirror Replication Agent location.

---

**Note** You can use any TDS client utility (`isql`, `isqlApp`, or `SQLAdvantage`) that you prefer.

---

- 2 Verify the connection to the Mirror Replication Agent:
  - a Open a command window in the `$SYBASE` directory of your Replication Server installation.
  - b Set the environment variables by sourcing the `SYBASE.csh` file.
  - c Use the `isql` utility to log in to Mirror Replication Agent:

```
isql -Usa -P -Smyma
```

---

**Note** The following verification steps are optional because they were done when you verified the resource file.

---

- 3 Verify the Mirror Replication Agent connection to Replication Server:

- a Enter this command:

```
test_connection RS
go
```

- b The following is displayed:

```
Type Connection

RS succeeded
(1 row affected)
```

- c If the result indicates a failure, either the server is not responding or the connection properties (host, port, user, or password) are incorrect.

Verify the host and port configuration values, and manually log in to the Replication Server as the configured user to determine which property is incorrectly defined.

For more information, refer to the Mirror Activator *Administration Guide*, Chapter 2, “Setup and Configuration,” in the section titled, “Testing network connectivity.”

- 4 Verify the Mirror Replication Agent connection to the primary Oracle database:

- a Enter this command:

```
test_connection PDS
go
```

This message appears:

```
Type Connection

PDS succeeded
(1 row affected)
```

- b If the result indicates a failure:
  - The server is not responding, or
  - The connection properties; *tnsnames.ora* connection name 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.

For more information about connection failures, see the Mirror Activator *Administration Guide* section titled, “Testing network connectivity.”

## Initializing the Mirror Replication Agent instance

This section describes the commands you must issue to initialize a Mirror Replication Agent instance.

### ❖ To initialize the Mirror Replication Agent instance

- 1 The `pdb_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 Mirror Replication Agent has the necessary permissions. It also creates objects in the database to support stored procedure replication.



The `pdb_init` command includes the `move_truncpt` keyword that executes an archiving operation in Oracle to archive all *redo* log files. This establishes the current position in the *redo* log as the end of the log, from a replication standpoint.

To verify correct configuration for archived and supplemental logging, issue the `pdb_init` command:

```
pdb_init move_truncpt
go
```

A message appears indicating that the procedure was successful.

For more information, see Chapter 2 in the *Mirror Activator Administration Guide*.

- 2 The `ra_init` command initializes the Mirror Replication Agent system database by reading schema information and *redo log* location information from the primary Oracle database. If this is a production setup, this step should coincide with creating the dump, copy, or data that is used to materialize the standby database.

To initialize the Mirror Replication Agent to read schema and *redo log* location information from the primary Oracle database, issue the `ra_init` command:

```
ra_init
go
```

A message appears indicating that the procedure was successful.

The `ra_init` command also causes `pdb_automark_tables` and `pdb_auto_create_repdefs` settings to take effect.

---

**Warning!** The default value for the `pdb_auto_create_repdefs` configuration property is `true`, which will cause a replication definition to be created for each table that is marked for replication during processing of the `ra_init` command. If you have thousands of tables, this may result in significant additional execution time. To avoid this additional execution time, set the `pdb_auto_create_repdefs` configuration property to `false` before invoking the `ra_init` command. When `ra_init` execution completes and before replicating, execute the `rs_create_repdef all` command to create your replication definitions.

---

- 3 The resume command puts the Mirror Replication Agent in an *ACTIVE* state, reading the Oracle *redo* log and sending commands to Mirror Replication Agent.

To place the Mirror Replication Agent in an *ACTIVE* state:

```
resume
go
```

If the Mirror Replication Agent successfully transfers to a *REPLICATING* state, this result is returned:

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

If the state returned is *ADMIN*, an error prevented the *REPLICATING* state from being achieved. To determine the error, review the contents of the Mirror Replication Agent system log. Also, see the Mirror Activator *Administration Guide*, Chapter 2, “Troubleshooting Mirror Replication Agent,” in the section titled, “Examining the Mirror Replication Agent when replication failure occurs.”

- 4 The ra\_status command returns the state of the Mirror Replication Agent. It is good practice to verify that the Mirror Replication Agent remains in *REPLICATING* state, even after the resume command executes successfully.

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

```
ra_status
go
```

If the Mirror Replication Agent is in *REPLICATING* state, this result is returned:

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

If the state returned is *ADMIN*, an error prevented the *REPLICATING* state from being achieved. To determine the error, review the contents of the Mirror Replication Agent system log.

Also, see the Mirror Activator *Administration Guide*, Chapter 2, “Troubleshooting Mirror Replication Agent,” in the section titled, “Examining the Mirror Replication Agent when replication failure occurs.”

- 5 Validate that both primary and standby connections are active:

```
isql -Usa -P -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 Mirror Replication Agent connection, if established for the standby database connection, is usually down, because you are not replicating data from the standby database.
- 

Do not proceed until admin who returns similar status for threads to the following:

| 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) DCOServer.oratest2              |
| 53   | DSI       | Awaiting Message | 102 DCOServer.oratest2                 |
| 17   | SQM       | Awaiting Message | 102:0 DCOServer.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

After you finish setting up the replication system, you can test replication.

### ❖ To test replication

- 1 Connect to the primary Oracle instance as a regular user, not the maintenance user. Make sure the regular user also exists in the standby database.

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

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

This statement is replicated by the user defined by the `ddl_username` Mirror Replication Agent configuration parameter. After the elapsed time referenced by the `scan_sleep_max` Mirror Replication Agent setting, the T1 table should be replicated to the standby database.

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

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

This statement is also replicated to the standby database by the user defined by the `ddl_username` Mirror Replication Agent configuration parameter.

- 4 After the table is marked and the Mirror 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. See “Configuring Replication Server for replication from the primary database” on page 93.

Mirror Replication Agent applies only committed transactions to a standby 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;
```

- 5 Use your preferred Oracle query tool to examine the standby site for results and compare the contents of your test table from both the primary and standby sites.

## Resetting the primary Oracle database for replication

In a test environment, you may occasionally need to reset the replication environment. Instead of deleting and re-creating a new Mirror Replication Agent instance, use this procedure to facilitate resetting the environment:

**❖ To reset the primary database for replication**

- 1 Protect your new environment from old log information by using the following command in the Mirror Replication Agent to archive all current *redo* log data:

```
pdb_init move_truncpt
```

Alternately, you can use the following Oracle command:

```
alter system archive log current ;
```

- 2 To retain marking and log device information, re-initialize the Mirror Replication Agent using `ra_init` with the `force` option, which forces the Mirror Replication Agent repository to be refreshed instead of overwritten:

```
ra_init force
```

---

**Note** If you prefer to delete and replace all the information in the Mirror Replication Agent repository, issue the `ra_deinit` command followed by a normal `ra_init` command (without the `force` option).

---

For more information about the `ra_init` command, see the Mirror Replication Agent *Reference Manual*.

- 3 Zero-out the locator stored in Replication Server using the `rs_zeroltm` command. For example:

```
isql -USAMPLE_RS_RSSD_prim -PSAMPLE_RS_RSSD_prim_ps
-SSAMPLE_RS_ERSSD -DSAMPLE_RS_ERSSD

rs_zeroltm NY, NYora92
go
```

The following is displayed:

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



# Glossary

This glossary describes Replication Agent terms used in this book.

## **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.
- *Adaptive Server IQ* manages multiple, large relational databases with special indexing algorithms to support high-speed, high-volume business intelligence, decision support, and reporting client applications.
- *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 standby 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. It provides an alternative to using SQL insert and select commands to transfer data.

## **bulk materialization**

A materialization method whereby subscription data in a standby 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**.

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>client</b>              | In client/server systems, the part of the system that sends requests to servers and processes the results of those requests. See also <b>client application</b> .                                                                                                                                                                                                                                                                    |
| <b>client application</b>  | Software that is responsible for the user interface, including menus, data entry screens, and report formats. See also <b>client</b> .                                                                                                                                                                                                                                                                                               |
| <b>commit</b>              | An instruction to the DBMS to make permanent the changes requested in a transaction. See also <b>transaction</b> . Contrast with <b>rollback</b> .                                                                                                                                                                                                                                                                                   |
| <b>data client</b>         | A client application that provides access to data by connecting to a data server. See also <b>client</b> , <b>client application</b> , and <b>data server</b> .                                                                                                                                                                                                                                                                      |
| <b>data distribution</b>   | 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 <b>data replication</b> .                                                                                                                            |
| <b>data replication</b>    | The process of copying data to remote locations, and then keeping the replicated data synchronized with the primary data. Data replication is different from data distribution. Replicated data is stored copies of data at one or more remote sites throughout a system, and it is not necessarily distributed data. Contrast with <b>data distribution</b> . See also <b>disk replication</b> and <b>transaction replication</b> . |
| <b>data server</b>         | 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 <b>client</b> , <b>client application</b> , and <b>data client</b> .                                                                                  |
| <b>database</b>            | A collection of data with a specific structure (or schema) for accepting, storing, and providing data for users. See also <b>data server</b> , <b>DBMS</b> , and <b>RDBMS</b> .                                                                                                                                                                                                                                                      |
| <b>database connection</b> | 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 <b>Replication Server</b> and <b>route</b> .                                                                                                                                                             |
| <b>datatype</b>            | 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.                                                                                                                                                                                                               |
| <b>DBMS</b>                | An abbreviation for <i>database management system</i> , 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 stand-alone data server system. Compare with <b>RDBMS</b> .                                                                                                                         |



|                          |                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>disaster recovery</b> | A method or process used to restore the critical business functions interrupted by a catastrophic event. A disaster recovery (or business continuity) plan defines the resources and procedures required for an organization to recover from a disaster, based on specified recovery objectives.                                                      |
| <b>ERSSD</b>             | An abbreviation for embedded <i>Replication Server System Database</i> , which manages replication system information for a Replication Server. See also <b>Replication Server</b> .                                                                                                                                                                  |
| <b>failback</b>          | A procedure that restores the normal user and client access to a primary database, after a failover procedure switched access from the primary database to a standby database. See also <b>failover</b> .                                                                                                                                             |
| <b>failover</b>          | A procedure that switches user and client access from a primary database to a standby 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 <b>failback</b> . |
| <b>function</b>          | A Replication Server object that represents a data server operation such as insert, delete, or begin transaction. Replication Server distributes operations to standby databases as functions. See also <b>function string</b> .                                                                                                                      |
| <b>function string</b>   | 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 standby databases are different types, with different SQL extensions and different command features. See also <b>function</b> .             |
| <b>gateway</b>           | Connectivity software that allows two or more computer systems with different network architectures to communicate.                                                                                                                                                                                                                                   |
| <b>inbound queue</b>     | A stable queue managed by Replication Server to spool messages received from a Mirror Replication Agent. See also <b>outbound queue</b> and <b>stable queue</b> .                                                                                                                                                                                     |
| <b>interfaces file</b>   | 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 <b>Open Client</b> and <b>Open Server</b> .                                                                                                                             |
| <b>isql</b>              | An interactive SQL client application that can connect and communicate with any Sybase Open Server application, including Adaptive Server, Mirror Replication Agent, and Replication Server. See also <b>Open Client</b> and <b>Open Server</b> .                                                                                                     |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Java</b>                   | An object-oriented programming language developed by Sun Microsystems. A platform-independent, “write once, run anywhere” programming language.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Java VM</b>                | 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 <b>Java</b> and <b>JRE</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>JDBC</b>                   | An abbreviation for <i>Java Database Connectivity</i> . JDBC is the standard communication protocol for connectivity between Java clients and data servers. See also <b>data server</b> and <b>Java</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>JRE</b>                    | An abbreviation for <i>Java Runtime Environment</i> . 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 the Mirror Replication Agent. See also <b>Java VM</b> .                                                                                                                                                                                                                                                                                                                                         |
| <b>LAN</b>                    | An abbreviation for “local area network,” a computer network located on the user’s premises that covers 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 <b>WAN</b> .                                                                                                                                                                                                                                                                        |
| <b>latency</b>                | <p>In transaction replication, the time it takes to replicate a transaction from a primary database to a standby database. Specifically, latency is the time elapsed between committing an original transaction in the primary database and committing the replicated transaction in the standby database.</p> <p>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 mirror (or standby) device.</p> <p>See also <b>disk replication</b> and <b>transaction replication</b>.</p> |
| <b>LOB</b>                    | An abbreviation for <i>large object</i> , a type of data element that is associated with a column that contains extremely large quantities of data.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Log Reader</b>             | An internal component of the Mirror Replication Agent that interacts with the primary database and mirror log devices to capture transactions for replication. See also <b>Log Transfer Interface</b> and <b>Log Transfer Manager</b> .                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Log Transfer Interface</b> | An internal component of the Mirror Replication Agent that interacts with Replication Server to forward transactions for distribution to a standby database. See also <b>Log Reader</b> and <b>Log Transfer Manager</b> .                                                                                                                                                                                                                                                                                                                                                                                                                  |

|                                  |                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Log Transfer Manager</b>      | An internal component of the Mirror Replication Agent that interacts with the other Mirror Replication Agent internal components to control and coordinate Mirror Replication Agent operations. See also <b>Log Reader</b> and <b>Log Transfer Interface</b> .                                                                     |
| <b>Maintenance User</b>          | A special user login name in the standby database that Replication Server uses to apply replicated transactions to the database. See also <b>Replication Server</b> .                                                                                                                                                              |
| <b>materialization</b>           | The process of copying the data from a primary database to a standby database, initializing the standby database so that the Replication Agent system can begin replicating transactions. See also <b>atomic materialization</b> , <b>bulk materialization</b> , and <b>non-atomic materialization</b> .                           |
| <b>nonatomic materialization</b> | 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 standby databases. Contrast with <b>atomic materialization</b> . See also <b>bulk materialization</b> . |
| <b>ODBC</b>                      | An abbreviation for <i>Open Database Connectivity</i> , an industry standard communication protocol for clients connecting to data servers. See also <b>JDBC</b> .                                                                                                                                                                 |
| <b>Open Client</b>               | 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 <b>Open Server</b> .                                                                                                                       |
| <b>Open Client application</b>   | An application that uses Sybase Open Client libraries to implement Open Client communication protocols. See also <b>Open Client</b> and <b>Open Server</b> .                                                                                                                                                                       |
| <b>Open Server</b>               | A Sybase product that provides the tools and interfaces required to create a custom server. See also <b>Open Client</b> .                                                                                                                                                                                                          |
| <b>Open Server application</b>   | A server application that uses Sybase Open Server libraries to implement Open Server communication protocols. See also <b>Open Client</b> and <b>Open Server</b> .                                                                                                                                                                 |
| <b>outbound queue</b>            | A stable queue managed by Replication Server to spool messages to a standby database. See also <b>inbound queue</b> and <b>stable queue</b> .                                                                                                                                                                                      |
| <b>primary data</b>              | The version of a set of data that is the source used for replication. Primary data is stored and managed by the primary database. See also <b>Mirror Replication Agent</b> , <b>primary database</b> , and <b>Replication Server</b> .                                                                                             |

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>primary database</b>    | The database that contains the data to be replicated to another database (the standby database) through a replication system. The primary database is the database that is the source of replicated data in a replication system. Sometimes called the <i>active database</i> . Contrast with <b>standby database</b> . See also <b>primary data</b> .                                                                                                                                            |
| <b>primary key</b>         | The column or columns whose data uniquely identify each row in a table.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>primary site</b>        | The location or facility at which primary data servers and primary databases are deployed to support normal business operations. Sometimes called the <i>active site</i> or <i>main site</i> . See also <b>primary database</b> and <b>standby site</b> .                                                                                                                                                                                                                                         |
| <b>primary table</b>       | A table used as a source for replication. Primary tables are defined in the primary database schema. See also <b>primary data</b> and <b>primary database</b> .                                                                                                                                                                                                                                                                                                                                   |
| <b>primary transaction</b> | A transaction that is committed in the primary database and recorded in the primary database transaction log. See also <b>primary database</b> , <b>replicated transaction</b> , and <b>transaction log</b> .                                                                                                                                                                                                                                                                                     |
| <b>quiesce</b>             | To cause a system to go into a state in which further data changes are not allowed. See also <b>quiescent</b> .                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>quiescent</b>           | <p>In a replication system, a state in which all updates have been propagated to their destinations. Some Mirror Replication Agent and Replication Server commands require that you first quiesce the replication system.</p> <p>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.</p> <p>This term is interchangeable with <i>quiesced</i> and <i>in quiesce</i>. See also <b>quiesce</b>.</p> |
| <b>MRASD</b>               | An abbreviation for <i>Mirror Replication Agent System Database</i> . Information in the MRASD is used by the primary database to recognize database structure or schema objects in the transaction log.                                                                                                                                                                                                                                                                                          |
| <b>RCL</b>                 | An abbreviation for <i>Replication Command Language</i> , the command language used to manage Replication Server.                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>RDBMS</b>               | An abbreviation for <i>relational database management system</i> , an application that manages and controls relational databases. Compare with <b>DBMS</b> . See also <b>relational database</b> .                                                                                                                                                                                                                                                                                                |
| <b>relational database</b> | 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. See also <b>SQL</b> .                                                                                                                                                                                                                                                                        |

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>replicated data</b>        | A set of data that is replicated from a primary database to a standby database by a replication system. See also <b>primary database</b> , <b>replication system</b> , and <b>standby database</b> .                                                                                                                                                                                                                                   |
| <b>replicated transaction</b> | A primary transaction that is replicated from a primary database to a standby database by a transaction replication system. See also <b>primary database</b> , <b>primary transaction</b> , <b>standby database</b> , and <b>transaction replication</b> .                                                                                                                                                                             |
| <b>Replication Agent</b>      | 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 standby database. See also <b>primary database</b> and <b>Replication Server</b> .                                                                                                        |
| <b>replication definition</b> | 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 <b>Replication Server</b> and <b>subscription</b> .                                                                              |
| <b>Replication Server</b>     | The Sybase software product that provides the infrastructure for a robust transaction replication system. See also <b>Replication Agent</b> .                                                                                                                                                                                                                                                                                          |
| <b>RSSD</b>                   | An abbreviation for <i>Replication Server System Database</i> , which manages replication system information for a Replication Server. See also <b>Replication Server</b> .                                                                                                                                                                                                                                                            |
| <b>replication system</b>     | 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 <b>disk replication</b> and <b>transaction replication</b> .                                                                                                                                                   |
| <b>rollback</b>               | An instruction to a database to back out of the changes requested in a unit of work (called a transaction). Contrast with <b>commit</b> . See also <b>transaction</b> .                                                                                                                                                                                                                                                                |
| <b>SQL</b>                    | An abbreviation for <i>Structured Query Language</i> , a non-procedural programming language used to process data in a relational database. ANSI SQL is an industry standard. See also <b>transaction</b> .                                                                                                                                                                                                                            |
| <b>stable queue</b>           | 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 standby database. Replication Server provides a stable queue for both incoming messages (the inbound queue) and outgoing messages (the outbound queue). See also <b>database connection</b> , <b>Replication Server</b> , and <b>route</b> . |

|                                  |                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>standby data</b>              | The data managed by a standby database, which is the destination (or target) of a replication system. See also <b>data replication</b> and <b>standby database</b> .                                                                                                                                                                               |
| <b>standby database</b>          | A database that contains data replicated from another database (the primary database) through a replication system. The standby database is the database that receives replicated data in a replication system. Sometimes called the <i>replicate database</i> . Contrast with <b>primary database</b> . See also <b>standby data</b> .            |
| <b>standby site</b>              | The location or facility at which standby data servers and standby databases are deployed to support disaster recovery, and normal business operations during scheduled downtime at the primary site. Sometimes called the <i>alternate site</i> or <i>replicate site</i> . Contrast with <b>primary site</b> . See also <b>standby database</b> . |
| <b>subscription</b>              | A request for Replication Server to maintain a replicated copy of a table, or a set of rows from a table, in a standby database at a specified location. See also <b>replication definition</b> and <b>Replication Server</b> .                                                                                                                    |
| <b>table</b>                     | 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 <b>database</b> .                                                                                                                          |
| <b>transaction</b>               | 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 <b>SQL</b> .                          |
| <b>transaction log</b>           | Generally, the log of transactions that affect the data managed by a data server. Mirror Replication Agent reads the transaction log to identify and acquire the transactions to be replicated from the primary database. See also <b>Mirror Replication Agent</b> , <b>primary database</b> , and <b>Replication Server</b> .                     |
| <b>transaction replication</b>   | A data replication method that copies data-changing operations from a primary database transaction log to a standby database. See also <b>data replication</b> and <b>disk replication</b> .                                                                                                                                                       |
| <b>transactional consistency</b> | A condition in which all transactions in the primary database are applied in the standby database, in the same order that they were applied in the primary database.                                                                                                                                                                               |
| <b>WAN</b>                       | An abbreviation for “wide area network,” a system of local-area networks (LANs) connected together with data communication lines. Contrast with <b>LAN</b> .                                                                                                                                                                                       |

# Index

## A

- applying HDS support scripts to RSSD 48, 90
- ASE
  - configuring Mirror Activator 17
  - configuring Mirror Activator components 11
  - configuring Replication Server for remote access 11
  - configuring Replication Server for replication from the primary database 13
  - configuring Replication Server for replication to standby 12
  - initializing the Mirror Replication Agent instance 23
  - installing Mirror Activator 3
  - installing Mirror Replication Agent 5
  - installing Replication Server 8
  - preparing and creating Mirror Replication Agent instance 17
  - prerequisites 2
  - resetting the primary database for replication 25
  - system configuration 1
  - testing replication 24
  - verifying the Mirror Replication Agent instance installation 21

## B

- before you begin 28, 74

## C

- configuring
  - components for Microsoft SQL Server 40
  - DirectConnect server for Microsoft SQL Server 41
  - ECDA 86
  - ECDA Option for ODBC 40

- Mirror Activator components in ASE 11
- Mirror Activator for ASE 17
- Replication Server for Microsoft SQL Server 48
- Replication Server for remote access with ASE 11
- Replication Server for replication from the primary database 13
- Replication Server for replication to standby 49
- Replication Server for replication to standby for ASE 12
- configuring Mirror Replication Agent 54, 96
- configuring Replication Server 90
  - for Microsoft SQL Server replication 48
  - for Oracle replication 90
  - for replication from the primary database 49, 91
  - for replication to standby 91
- creating
  - a DirectConnect server instance 42, 87
  - database replication definitions 94
  - database replication subscriptions 52, 95
  - Mirror Replication Agent instance 17, 100
  - objects 47, 89
  - Oracle user 98
  - public synonyms 47, 89
  - Replication Server connection to standby database 49, 91
  - Replication Server connection to standby database for Microsoft SQL Server 49
  - Replication Server from primary database 51, 93
  - resource file 105

## D

- DirectConnect server
  - configuring for Microsoft SQL Server 41
- dsedit utility 107

## E

- executing
  - resource file 105

## I

- identifying the Sybase installation directory 30
- initializing
  - Mirror Replication Agent instance 23
- installation directory
  - identifying 30
- installation guidelines 38, 78, 81, 84
  - for ECDA 84
  - for ECDA Option for ODBC 38
  - for Mirror Replication Agent 32, 78
  - for Replication Server 35, 81
  - Mirror Replication Agent for Microsoft SQL Server 32
  - Replication Server for Microsoft SQL Server 35
- installing
  - Mirror Activator for ASE 3
  - Mirror Replication Agent for ASE 5
  - Replication Server for ASE 8
- instance
  - creating a Mirror Replication Agent instance for ASE 17

## M

- ma\_admin utility 106
- Microsoft SQL Server
  - configuring Replication Server 48
  - configuring the DirectConnect server 41
  - creating a DirectConnect server instance for 42
  - installing ECDA Option for ODBC 38
  - installing Mirror Replication Agent 32
  - installing Replication Server 35
  - system requirements 29
  - verifying the Mirror Replication Agent instance 64
- Mirror Activator
  - components 30, 76
- Mirror Activator components
  - configuration 40, 86
- Mirror Replication Agent

- sample resource file 101
- verifying connection to Oracle 108
- verifying the installation in Oracle 106

## O

- Oracle
  - installing ECDA 84
  - installing Mirror Activator components 76
  - installing Mirror Replication Agent 77
  - installing Replication Server 80
  - verifying the Mirror Replication Agent instance 106

## P

- parameters for resource file 102
- prerequisites
  - ASE 2

## R

- replication
  - overview 30, 76
- replication from the primary Microsoft SQL Server database
  - configuring Replication Server 49
- replication system
  - diagram of 74
  - diagram of Microsoft SQL Server 28
  - diagram of Mirror Activator for ASE 1
- resetting primary Microsoft SQL Server database for replication 71
- resetting primary Oracle database for replication 112
- resource file
  - creating and executing 105
  - parameters 102

## S

- sample resource file
  - access for Mirror Replication Agent 101



- setup prerequisites 28, 74
- system configuration
  - Mirror Activator for ASE 1
- system requirements 74, 75
  - Microsoft SQL Server 29

## T

- testing
  - ASE replication 24
- testing replication 70, 111

## U

- utilities
  - dsedit 107
  - ma\_admin 106

## V

- verifying
  - compatibility level of Microsoft SQL Server 55
  - current archiving setting of the redo logs 96
  - Mirror Replication Agent connection to Oracle 108
  - Mirror Replication Agent installation in ASE 7
  - Mirror Replication Agent instance in ASE 21
  - Mirror Replication Agent instance in Microsoft SQL Server 64
  - Mirror Replication Agent instance in Oracle 106
  - primary Oracle configuration for replication 99
  - supplemental logging of primary key data 97

