



**Primary Database Guide**

---

**Replication Agent™ for Oracle**

**15.7.1 ESD #2**

Linux, Microsoft Windows, and UNIX

DOCUMENT ID: DC01848-01-1571-01

LAST REVISED: October 2012

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

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

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

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

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

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

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

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

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

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

# Contents

<b>Conventions</b> .....	<b>1</b>
<b>Replication Agent for Oracle</b> .....	<b>3</b>
Oracle-Specific Considerations .....	3
Unsupported Oracle Software Features .....	3
Unsupported Oracle Datatypes, Data, and Structures .....	4
Supported Oracle Datatypes, Data, and Structures .....	5
Replication Agent Connectivity .....	6
Replication Agent Permissions .....	6
Multiple Replication Agents .....	10
High Availability .....	12
Redo and Archive Log Setup .....	12
Supplemental Logging .....	16
DDL Replication .....	17
Character Case of Database Object Names .....	19
Format of Origin Queue ID .....	20
LTL Origin Commit Time Granularity .....	20
Replication Server and RSSD Scripts .....	20
Oracle Datatype Compatibility .....	21
Oracle Datatype Restrictions .....	28
Oracle Large Object (LOB) Support .....	31
Oracle User-Defined Types .....	33
Sequence Marking and Unmarking .....	35
Sequence Replication Enabling and Disabling .....	38
Setting Up Replication Agent and Oracle on Different Machines .....	39
Real Application Clusters (RAC) .....	40
Automatic Storage Management .....	42
Replication Server set autocorrection Command .....	43

Partitioned Tables .....	44
Materialized Views .....	45
Index-Organized Tables .....	50
Replicate Database Trigger Execution Control ....	51
Alteration of Replication Definitions from the Primary Data Server .....	51
Oracle Data Guard .....	52
Database Resynchronization .....	53
Oracle Transaction and Operation Troubleshooting .....	53
Stored Procedure Replication with BOOLEAN Arguments .....	53
Oracle Warm Standby .....	55
Oracle Flashback .....	56
XMLTYPE Data Replication .....	59
Oracle 11g Release 2 .....	62
Oracle 9i .....	64
Replication Agent Objects in the Oracle Primary Database .....	65
Replication Agent Object Names .....	65
Table Objects .....	66
Marker Objects .....	68
Sequences .....	68
Marked Procedures .....	69
Transaction Log Truncation .....	69
<b>Upgrading and Downgrading Replication Agent .....</b>	<b>71</b>
Upgrade and Migration Procedures for Replication Agent for Oracle .....	71
Upgrading Replication Agent for Oracle to 15.7.1 ESD #2 .....	72
Migrating Replication Agent for Oracle 15.7.1 ESD #2 When Upgrading Oracle 10g to 11g .....	72
Downgrading Replication Agent for Oracle .....	73
<b>Obtaining Help and Additional Information .....</b>	<b>75</b>

Technical Support .....	75
Downloading Sybase EBFs and Maintenance Reports .....	75
Sybase Product and Component Certifications .....	76
Creating a MySybase Profile .....	76
Accessibility Features .....	76
<b>Glossary .....</b>	<b>79</b>
<b>Index .....</b>	<b>87</b>

# Contents

# Conventions

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

## Style conventions

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

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

Run:

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

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

### *Syntax conventions*

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

### *Case-sensitivity*

- All command syntax and command examples are shown in lowercase. However, replication command names are not case-sensitive. For example, **RA\_CONFIG**, **Ra\_Config**, and **ra\_config** are equivalent.
- Names of configuration parameters are case-sensitive. For example, **Scan\_Sleep\_Max** is not the same as **scan\_sleep\_max**, and the former would be interpreted as an invalid parameter name.
- Database object names are not case-sensitive in replication commands. However, to use a mixed-case object name in a replication command (to match a mixed-case object name in the primary database), delimit the object name with double quote characters. For example: **pdb\_get\_tables "TableName"**
- Identifiers and character data may be case-sensitive, depending on the sort order that is in effect.
  - If you are using a case-sensitive sort order, such as “binary,” you must enter identifiers and character data with the correct combination of uppercase and lowercase letters.
  - If you are using a sort order that is not case-sensitive, such as “nocase,” you can enter identifiers and character data with any combination of uppercase or lowercase letters.

### *Terminology*

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



# Replication Agent for Oracle

Review the characteristics of Replication Agent that are unique to Replication Agent for Oracle.

The term "Replication Agent for Oracle" refers to an instance of Replication Agent software that is installed and configured for a primary database that resides in an Oracle data server.

---

**Note:** For information on the basic features and operation of Replication Agent, see the *Replication Agent Administration Guide* and *Replication Agent Reference Manual*.

---

## Oracle-Specific Considerations

---

These general issues and considerations are specific to using Replication Agent with the Oracle data server.

### Unsupported Oracle Software Features

Review the Oracle features that are not supported by Replication Agent.

These features are not supported:

- Oracle virtual columns
- Oracle label security
- Oracle-packaged stored procedures and functions (standalone procedures and functions are supported)
- Oracle schema objects in encrypted tablespaces
- Replication Server<sup>®</sup> parallel DSI
- Replication Server **rs\_init** utility
- Replication Server **rs\_subcomp** utility
- Replication Server automatic materialization
- Replication Server when replicating in an environment where other vendors are replicating

### Replication of Deferred Updates on Primary Keys

Updates on tables that have a unique column index are not supported by traditional replication, and the Replication Server reports errors.

The replication of updates to tables that have a unique column index is not supported, and Replication Server reports errors. For example, table *t* has a unique index on column *c*, with values 1, 2, 3, 4 and 5. A single **update** statement is applied to the table:

```
update t set c = c+1
```

Using traditional replication, this statement results in:

```
update t set c = 2 where c = 1
update t set c = 3 where c = 2
update t set c = 4 where c = 3
update t set c = 5 where c = 4
update t set c = 6 where c = 5
```

The first update attempts to insert a value of `c=2` into the table. However, this value already exists in the table. Replication Server displays error 2601—an attempt to insert a duplicate key.

Adaptive Server<sup>®</sup> Enterprise solves this problem with the SQL statement feature, which avoids a duplicate-key error by deferring the updates and executing them in reverse order. However, non-Sybase data servers do not provide this feature, and the Replication Server DSI will crash if you attempt to replicate updates to a non-Sybase table that has a unique column index. To work around this problem, broaden the unique index definition.

### **Unsupported Oracle Datatypes, Data, and Structures**

Review the Oracle datatypes that are not supported by Replication Agent.

These datatypes are not supported:

- Oracle-supplied datatypes:
  - “Any” Types (`SYS.ANYTYPE`, `SYS.ANYDATASET`), except for `SYS.ANYDATA`.
  - `MLSLABEL`
  - Spatial Types (`MDSYS.SDO_GEOMETRY`, `SDO_TOPO_GEOMETRY`, `SDO_GEORASTER`)
  - Media Types (`ORDSYS.ORDAudio`, `ORDSYS.ORDImage`, `ORDSYS.ORDImageSignature`, `ORDSYS.ORDVideo`, `ORDSYS.ORDDoc`, `SI_StillImage`, `SI_Color`, `SI_AverageColor`, `SI_ColorHistogram`, `SI_PositionalColor`, `SI_Texture`, `SI_FeatureList`)
  - Datatypes used with Oracle Expression Filter
  - `ANYDATA`, if it is being replicated to a non-`ANYDATA` column or if the data exceeds 16KB, which is the size constraint of the Replication Server `OPAQUE` datatype.
  - `BFILE`, `UROWID`, or `REF` data stored in an `ANYDATA` column
  - `BLOB` or object user-defined datatype data being replicated to a replicate database other than Oracle or Adaptive Server<sup>®</sup> Enterprise (ASE). When replicating to Oracle, use ExpressConnect for Oracle since data of these types is not supported by Enterprise Connect<sup>™</sup> Data Access (ECDA) for Oracle.
  - `REF`
  - `UROWID`
- User-defined datatypes containing LOB data
- User-defined datatypes defined as `NOT FINAL`
- Partial updates to Oracle LOBs are not supported.
- Virtual columns – Replication Agent supports the replication of tables containing computed (or virtual) columns in Oracle 11g. However, the replication of individual

computed columns is not supported. You can mark tables with virtual columns for replication using the **force** option, but the virtual columns are not replicated.

Data stored in the Oracle XML DB Repository is not supported if it is accessed with these protocols or methods:

- Internet protocols like FTP, HTTP, HTTPS, and WebDAV
- The Oracle XML DB Repository API

These user-defined object types and structures are not supported:

- Associative arrays
- Nested tables
- VARRAY
- Nested tables or VARRAY stored in an ANYDATA column

### **Predefined PL/SQL Numeric Datatypes**

Replication Agent does not support marking procedures containing the `PLS_INTEGER` or `BINARY_INTEGER` predefined PL/SQL numeric datatypes.

Replication Agent does support marking procedures with the `SIMPLE_INTEGER` datatype, which is a subtype of `PLS_INTEGER`. However, Replication Agent does not support marking procedures containing any other subtypes of the aforementioned types, including `NATURAL`, `NATURALN`, `POSITIVE`, `POSITIVEN`, and `SIGNTYPE`.

## **Supported Oracle Datatypes, Data, and Structures**

Replication Agent supports the replication of encrypted data, compressed data, and SecureFiles.

### *Encrypted Data*

Replication Agent transparently supports the replication of data in encrypted columns, partitions, and tablespaces for Oracle versions 11g Release 1 and later.

If your primary Oracle database version is earlier than 11g Release 1, you can still mark tables containing encrypted columns, partitions, and tablespaces using the **force** option, but the columns are not replicated.

### *Compressed Data*

Replication Agent supports the replication of compressed tables, compressed tablespaces, and data compressed for Oracle direct-load operations (bypassing I/O buffers) for Oracle versions 11g Release 2 and later.

---

**Note:** If you want to replicate tables compressed for online transaction processing (OLTP), apply first the Oracle patch# 129050503.

---

### *SecureFile Data*

Replication Agent supports the replication of Oracle data stored with the SecureFile option.

Replication Agent supports the replication of Oracle LOB data stored with the SecureFile option for Oracle versions 11g Release 2 and later. If your primary Oracle database version is earlier than 11g Release 2, you can still mark tables containing SecureFile LOB columns using the **force** option, but the columns are not replicated.

### Replication Agent Connectivity

Replication Agent for Oracle uses JDBC™ for communications with all replication system components.

The Oracle JDBC driver must be installed on the Replication Agent host machine, and the directory this driver is installed in must be in the CLASSPATH environment variable.

The TNS Listener Service must be installed and running on the primary database so the Replication Agent instance can connect to it. See the *Oracle Database Net Services Administrator's Guide*. See the *Replication Agent Installation Guide* for the specific JDBC driver and version to install.

### Replication Agent Permissions

Replication Agent for Oracle uses the **pds\_username** to connect to Oracle and must have certain Oracle permissions.

These permissions are required:

- **GRANT ALTER ANY PROCEDURE** – required to manage procedures for replication.
- **GRANT ALTER DATABASE** – required for Replication Agent to read from a Data Guard standby database transaction log.
- **GRANT ALTER ON *TABLE\_NAME*** – required to replicate user-defined datatypes if table-level supplemental logging has not been enabled for the specified *TABLE\_NAME*.
- **GRANT ALTER SYSTEM** – required to perform redo log archive operations.
- **GRANT CREATE ANY PROCEDURE** – required to mark procedures for replication.
- **GRANT CREATE PROCEDURE** – required to create **rs\_marker** and **rs\_dump proc** procedures.
- **GRANT CREATE PUBLIC SYNONYM** – required to create synonyms for created tables in the primary database.
- **GRANT CREATE SEQUENCE** – required to support replication.
- **GRANT CREATE SESSION** – required to connect to Oracle.
- **GRANT CREATE TABLE** – required to create tables in the primary database.
- **GRANT DROP PUBLIC SYNONYM** – required to drop created synonyms.
- **GRANT EXECUTE\_CATALOG\_ROLE** – required to use Oracle LogMiner.
- **GRANT EXECUTE ON DBMS\_FLASHBACK** – required to execute **DBMS\_FLASHBACK.get\_system\_change\_number**.
- **GRANT EXECUTE ON SYS.DBMS\_LOCK** – required to generate **commit** log records at the primary database.

- **GRANT SELECT ANY TRANSACTION** – required to use Oracle LogMiner.
- **GRANT SELECT\_CATALOG\_ROLE** – required to select from the **DBA\_\* views**.
- **GRANT SELECT ON SYS.OPQTYPE\$** – required for DDL replication and XMLTYPE data replication.
- **GRANT SELECT ON SYS.RECYCLEBIN\$** – required to use Oracle Flashback with Replication Agent.
- **GRANT SELECT ON SYS.ARGUMENT\$** – required to process procedure DDL commands.
- **GRANT SELECT ON SYS.CCOL\$** – required to support table replication (column constraint information).
- **GRANT SELECT ON SYS.CDEF\$** – required to support table replication (constraint information).
- **GRANT SELECT ON SYS.COL\$** – required to support table replication (column information).
- **GRANT SELECT ON SYS.COLLECTION\$** – required to support table replication.
- **GRANT SELECT ON SYS.COLTYPE\$** – required to support table replication.
- **GRANT SELECT ON SYS.CON\$** – required to support table replication (constraint information).
- **GRANT SELECT ON SYS.DEFERRED\_STG\$** – required to suppress replication of compressed tables on Oracle 11g Release 2, on which LogMiner does not support compressed tables.
- **GRANT SELECT ON SYS.IND\$** – required to identify indexes.
- **GRANT SELECT ON SYS.INDCOMPART\$** – required to identify indexes.
- **GRANT SELECT ON SYS.INDPART\$** – required to identify indexes.
- **GRANT SELECT ON SYS.INDSUBPART\$** – required to identify indexes.
- **GRANT SELECT ON SYS.LOB\$** – required for LOB replication support.
- **GRANT SELECT ON SYS.LOBCOMPPART\$** – required to support partitioned LOB replication.
- **GRANT SELECT ON SYS.LOBFrag\$** – required to support partitioned LOB replication.
- **GRANT SELECT ON SYS.MLOG\$** – required to filter materialized view log tables.
- **GRANT SELECT ON SYS.NTAB\$** – required to support table replication.
- **GRANT SELECT ON SYS.OBJ\$** – required for processing procedure DDL commands in the repository.
- **GRANT SELECT ON SYS.PROCEDUREINFO\$** – required for procedure replication support.
- **GRANT SELECT ON SYS.SEG\$** – required to suppress replication of compressed tables on versions of Oracle where LogMiner does not support compressed tables.
- **GRANT SELECT ON SYS.SEQ\$** – required to support sequence replication.
- **GRANT SELECT ON SYS.SNAP\$** – required to filter out materialized view tables.
- **GRANT SELECT ON SYS.TAB\$** – required to support table replication.
- **GRANT SELECT ON SYS.TABCOMPART\$** – required to support partitioned table replication.

- **GRANT SELECT ON SYS.TABPART\$** – required to support partitioned table replication.
- **GRANT SELECT ON SYS.TABSUBPART\$** – required to support partitioned table replication.
- **GRANT SELECT ON SYS.TS\$** – required to identify tablespace encryption in Oracle 11g.
- **GRANT SELECT ON SYS.TYPE\$** – required to process Oracle predefined and user-defined types.
- **GRANT SELECT ON SYS.USER\$** – required for Oracle user identification.
- **GRANT SELECT ON SYS.ATTRIBUTE\$** – required to process Oracle types.
- **GRANT SELECT ON V\_\$LOGMNR\_CONTENTS** – required to use Oracle LogMiner.
- **GRANT SELECT ON V\_\$LOGMNR\_LOGS** – required to use Oracle LogMiner.

If the Replication Agent is configured to remove old archive files, the user must have **UPDATE** authority to the directory and the archive log files.

Replication Agent for Oracle requires the **ALTER SYSTEM** privilege to issue the **ALTER SYSTEM ARCHIVE LOG** command. If Replication Agent is configured to access only online Oracle redo logs, Replication Agent issues the **ALTER SYSTEM ARCHIVE LOG SEQUENCE** command when the online redo log is no longer needed for replication (as when all data from the log has been replicated). Regardless of online or archive log processing, Replication Agent uses the **ALTER SYSTEM** privilege to issue the **ALTER SYSTEM ARCHIVE LOG CURRENT** command when Replication Agent is instructed to move processing to the end of the Oracle log. By issuing the **ALTER SYSTEM ARCHIVE LOG CURRENT** command, Replication Agent ensures that the current redo log file does not contain old data. Replication Agent moves processing to the end of the Oracle redo log when requested by the **move\_truncpt** option of the **ra\_locator** command. Replication Agent may also move processing to the end of the Oracle redo log during migration from one version of Replication Agent to another.

### *Oracle 10g and 11g Privileges for DDL Replication*

---

**Note:** Issuing **GRANT ALL PRIVILEGES TO DDLUSER** turns the DDL user into a superuser, like the SYS or SYSTEM user.

---

Different versions of Oracle have different permission requirements. For Oracle 10g and 11g, grant the DDL user permission to execute these commands:

- **GRANT ALTER ANY INDEX**
- **GRANT ALTER ANY INDEXTYPE**
- **GRANT ALTER ANY PROCEDURE**
- **GRANT ALTER ANY TABLE**
- **GRANT ALTER ANY TRIGGER**
- **GRANT ALTER ANY TYPE**
- **GRANT ALTER SESSION**
- **GRANT BECOME USER**
- **GRANT CREATE ANY INDEX**

- **GRANT CREATE ANY INDEXTYPE**
- **GRANT CREATE ANY PROCEDURE**
- **GRANT CREATE ANY SYNONYM**
- **GRANT CREATE ANY TABLE**
- **GRANT CREATE ANY TRIGGER**
- **GRANT CREATE ANY TYPE**
- **GRANT CREATE ANY VIEW**
- **GRANT CREATE INDEXTYPE**
- **GRANT CREATE MATERIALIZED VIEW**
- **GRANT CREATE PROCEDURE**
- **GRANT CREATE PUBLIC SYNONYM**
- **GRANT CREATE SYNONYM**
- **GRANT CREATE TABLE**
- **GRANT CREATE TRIGGER**
- **GRANT CREATE TYPE**
- **GRANT CREATE VIEW**
- **GRANT DELETE ANY TABLE**
- **GRANT DROP ANY INDEX**
- **GRANT DROP ANY INDEXTYPE**
- **GRANT DROP ANY MATERIALIZED VIEW**
- **GRANT DROP ANY PROCEDURE**
- **GRANT DROP ANY SYNONYM**
- **GRANT DROP ANY TABLE**
- **GRANT DROP ANY TRIGGER**
- **GRANT DROP ANY TYPE**
- **GRANT DROP ANY VIEW**
- **GRANT DROP PUBLIC SYNONYM**
- **GRANT INSERT ANY TABLE**
- **GRANT SELECT ANY TABLE**
- **GRANT UPDATE ANY TABLE**

Revoke user permission from the DDL user to execute:

- **GRANT ALTER DATABASE**
- **GRANT ALTER ROLLBACK SEGMENT**
- **GRANT ALTER SYSTEM**
- **GRANT ALTER TABLESPACE**
- **GRANT ANALYZE ANY**
- **GRANT AUDIT ANY**
- **GRANT AUDIT SYSTEM**

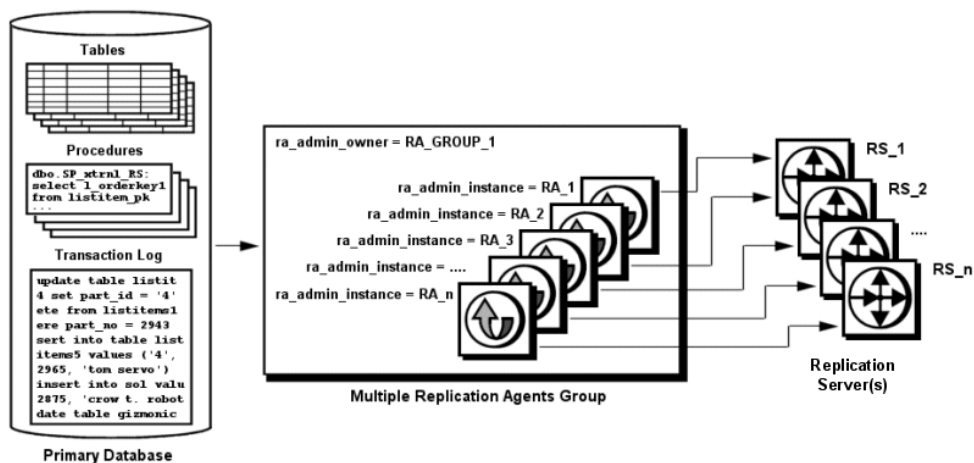
- GRANT CREATE DATABASE LINK
- GRANT CREATE ROLLBACK SEGMENT
- GRANT CREATE TABLESPACE
- GRANT DROP PUBLIC DATABASE LINK
- GRANT DROP ROLLBACK SEGMENT
- GRANT DROP TABLESPACE
- GRANT LOCK ANY TABLE

## Multiple Replication Agents

Multiple Replication Agent instances allow transactions to be replicated in parallel along multiple independent paths, increasing replication throughput and performance, and reducing resource contention.

Using multiple Replication Agent instances with Replication Server Multi-Path Replication™ enables the replication of data through different streams, maintaining data consistency within a path but not adhering to the commit order across different paths. Replication Agent for Oracle supports the use of multiple Replication Agent instances to replicate Oracle tables and procedures.

Figure 1: Multipath Replication



**Note:** All Replication Agent instances in a common Replication Agent group must be of the same version.

For a description of end-to-end Multi-Path Replication scenarios, see the *Replication Server Heterogeneous Replication Guide*.



### Configuration

In a Replication Server Multi-Path Replication scenario, multiple Replication Agent for Oracle instances are grouped together and share common system resources to coordinate in a Multi-Path Replication solution:

- Every Replication Agent for Oracle instance in a Replication Agent group must use the same value for the **ra\_admin\_owner** and **ra\_admin\_prefix** parameters.
- Each Replication Agent for Oracle instance in a Replication Agent group must use a unique value for the **ra\_admin\_instance\_prefix** parameter. This parameter distinguishes Replication Agent for Oracle instances within a group.

See the *Replication Agent Reference Manual*.

### DDL Replication with Multiple Replication Agents

You can configure each Replication Agent instance to replicate all DDL operations or no DDL operations in the primary database.

In a Replication Agent group, you can also configure each Replication Agent instance to replicate:

- DDL operations for all objects marked for replication by the instance
- DDL operations for all objects marked for replication by the instance and DDL operations for objects not marked for replication by any instances

To avoid DDL synchronization errors with multiple Replication Agents, use these keywords with the **pdb\_setrepddl** command:

- **all** specifies that all DDL operations are to be replicated by the instance.
- **marked** specifies that only DDL operations for objects that have been marked for replication by the instance are to be replicated.
- **unmarked** specifies that DDL operations for objects that have not been marked for replication by any instance are to be replicated. Nonschema DDL operations are also replicated by an instance that specifies this keyword.

### Example

A Replication Agent group containing three instances replicates DDL operations performed on these tables in the primary database:

- `table_A` is marked for replication by instance `ra_instance_1`.
- `table_B` is marked for replication by instance `ra_instance_2`.
- `table_C` is marked for replication by instance `ra_instance_3`.
- `table_D` is not marked for replication by any instance.

The Replication Agent instances use the **pdb\_setrepddl** command to specify which DDL to replicate:

```
ra_instance_1>pdb_setrepddl enable, marked
ra_instance_1>pdb_setrepddl enable, unmarked
```

## Replication Agent for Oracle

```
ra_instance_2>pdb_setrepddl enable, marked  
ra_instance_3>pdb_setrepddl enable, marked
```

The DDL is replicated as follows:

Replication Agent Instance	DDL Replicated
ra_instance_1	DDL for table_A, table_D, and nonschema DDL operations
ra_instance_2	DDL for table_B
ra_instance_3	DDL for table_C

### High Availability

A primary data server may employ a high-availability solution, such as failover clustering, to minimize downtime in the event of hardware or software failure.

Although Replication Agent does not provide any high-availability solutions, it works with a third-party high-availability solution for the primary database if:

- Replication Agent is installed on a shared file system, such as OCFS, a Network File System (NFS), or a Veritas Cluster Server (VCS). The Replication Agent binaries, configuration files, and RASD files must be installed on the system.
- A third-party cluster-management solution—such as Sun Cluster Manager, Veritas Cluster Manager, or Oracle Cluster Ready Services (CRS)—is used to automatically start Replication Agent in the event of failover.

### Redo and Archive Log Setup

By default, you can access both online and archive logs. You can configure Replication Agent to access only the online logs, but doing so requires you to turn Oracle automatic archiving off and requires Replication Agent to issue manual archive log commands to Oracle.

---

**Note:** If Replication Agent for Oracle is configured to truncate Oracle archive logs directly—`rman_enabled` is set to false and `pdb_archive_remove` is set to true—it must be installed on a host that has direct access to the Oracle archive log files.

---

#### *Archive Log Access*

When you are using the default, for archive log files to be accessed, configure Replication Agent to use the directory path where archive log files are located. By default, an Oracle instance creates multiple directories under the flash recovery area specified by the `DB_RECOVERY_FILE_DEST` parameter of the Oracle **ALTER SYSTEM** command, each directory corresponding to and named after a separate day. However, Replication Agent requires archived redo log files to reside in a single directory. Consequently, you must configure Oracle to archive to a single directory to be read by Replication Agent.

---

**Note:** To prevent conflicts with other archive file processes, you may want to configure Oracle to duplex the archive log files into an additional destination directory that is used only for replication.

---

For information on specifying archive log destinations for your Oracle environment, see the Oracle **ALTER SYSTEM** command and **LOG\_ARCHIVE\_DEST\_n** parameter.

---

**Note:** Accessing Oracle archived redo logs that are stored as file-system files is discussed in this section. If the archived redo logs are stored using the Oracle ASM, see the discussion of Automatic Storage Management.

---

Replication Agent for Oracle requires that redo log archiving is enabled at your Oracle database:

```
alter database ARCHIVELOG;
```

---

**Note:** If you are using Oracle Real Application Clusters (RAC), you must enable redo log archiving for each instance in the cluster.

---

Verify that log archiving is enabled:

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

---

If you are using Oracle RAC, use this SQL statement to verify that log archiving has been enabled:

```
select instance, name, log_mode from gv$database;
```

---

If ARCHIVELOG (ARCHIVELOG or MANUAL in Oracle 10g) is returned, then log archiving is enabled.

### *Archive Log File Access*

In the Replication Agent, set the **pdb\_archive\_path** configuration property to the expected location of archived redo log files. You can also set the Replication Agent **pdb\_archive\_remove** configuration parameter to true to allow the Replication Agent to remove these archive log files when they are no longer needed to support replication.

The **rman\_enabled** parameter enables Replication Agent to use the Oracle **RMAN** utility to truncate old archive log files. See the *Replication Agent Reference Manual*.

### *Replication Agent Archive Setup*

When **pdb\_include\_archives** is set to true (the default), Replication Agent does not archive, and Sybase recommends that you configure Oracle to perform automatic archiving of redo logs.

When the **pdb\_include\_archives** configuration parameter is set to false, Replication Agent for Oracle also requires you to disable automatic archiving of Oracle redo logs. Archiving is performed manually by Replication Agent as the data in the online redo log files is replicated.

Replication Agent for Oracle requires these settings in your Oracle database depending on the Oracle version.

### See also

- *Automatic Storage Management* on page 42

### **Disabling Automatic Archiving for Oracle 10g**

Disable automatic archiving for Oracle 10g.

1. Make sure you have **sysdba** administrator privileges, and close the database.
2. Enter:

```
alter database ARCHIVELOG MANUAL;
```

3. To verify that log archiving is disabled, enter:

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

If MANUAL is returned, then automatic log archiving is disabled.

### **Disabling Automatic Archiving for Oracle 11g**

Disable automatic archiving for Oracle 11g.

1. To change the LOG\_ARCHIVE\_START parameter, either manually edit the server start-up parameter file, or enter:

```
alter system set log_archive_start=false scope=spfile;
```

2. To check the setting of the LOG\_ARCHIVE\_START parameter, enter:

```
select value from v$system_parameter where name =  
'log_archive_start';
```

3. If false is returned, the value in the server parameter file has been correctly modified to prevent automatic archiving when you restart the Oracle server. For information about the LOG\_ARCHIVE\_START parameter or the ALTER SYSTEM commands, see the *Oracle Database Reference Guide*.
4. Automatic archiving must be disabled in the active server and when you restart the Oracle server. To stop automatic archiving in the active server, enter:

```
alter system archive log stop;
```

5. To disable automatic archiving when you restart the Oracle server, change the value of the server LOG\_ARCHIVE\_START parameter to false.

---

**Note:** If **pdb\_include\_archives** is set to false: For redo log file processing after Replication Agent for Oracle is initialized, automatic archiving must never be enabled, even temporarily. If automatic archiving is enabled or if manual archiving is performed, causing a redo log file not yet processed by the Replication Agent to be overwritten, the data in the lost redo log file is not replicated. You can recover from this situation by reconfiguring the Replication Agent to access archive log files. Set **pdb\_include\_archives** to true, set **pdb\_archive\_path** to the directory location that contains the archive of the file that has been overwritten, and resume. After catching up, suspend the Replication Agent, and reset **pdb\_include\_archives** to false.

---

### **Forcing the Logging of All Database Changes**

Sybase recommends that you enable forced logging of all database changes to the Oracle redo log file to ensure that all data that should be replicated is logged.

1. In the primary database, execute:

```
alter database FORCE LOGGING;
```

2. Verify the current setting by executing:

```
select force_logging from v$database;
```

### **UNC Paths for Windows Archive and Online Redo Log Paths**

If Replication Agent for Oracle is running as a Windows service and the primary Oracle data server is installed on a separate machine, configure the archive and online redo log paths according to the Microsoft Windows Universal Naming Convention (UNC):

```
\\oracle_server_machine\oracle_log_path
```

where *oracle\_server\_machine* is where the primary Oracle data server resides, and *oracle\_log\_path* is the archive or redo log file. For example, to set the location of archive redo log files to the `oracle` directory on the machine named `labratx64`, enter:

```
1> ra_config pdb_archive_path, \\labratx64\oracle
2> go
```

### **Configuring Oracle LogMiner**

Configure Oracle LogMiner on the primary Oracle database.

1. Go to `$(ORACLE_HOME)/rdbms/admin`.
2. Log in as a “sys as sysdba” user.
3. Execute the Oracle LogMiner installation script:

```
@dbmslm.sql
```

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

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

---

**Note:** This step is required if you are using Oracle 10g.

---

5. Grant these privileges to `pds_username`:
  - EXECUTE\_CATALOG\_ROLE
  - SELECT ON V\_\$LOGMNR\_CONTENTS
  - SELECT ON V\_\$LOGMNR\_LOGS
  - SELECT ANY TRANSACTION

---

**Note:** The **ra\_migrate** command verifies that these privileges have been granted to **pds\_username**. If these privileges have not been granted at the time **ra\_migrate** is invoked, a warning message is returned and logged in the Replication Agent log file.

---

### **Supplemental Logging**

Enable supplemental logging and supplemental logging of primary key data and index columns.

To enable supplemental logging, execute these Oracle commands:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
```

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

To verify that minimal supplemental logging and supplemental logging of primary key and unique index information is enabled, enter:

```
select SUPPLEMENTAL_LOG_DATA_MIN, SUPPLEMENTAL_LOG_DATA_PK,  
SUPPLEMENTAL_LOG_DATA_UI from v$database;
```

If YES is returned for each column, supplemental logging of primary key information is enabled.

### **Table-Level Supplemental Logging**

To replicate updates to user-defined object type attributes, Replication Agent must enable table-level supplemental logging.

Manually enable table-level supplemental logging by entering:

```
ALTER TABLE THE_TABLE ADD SUPPLEMENTAL LOG DATA (ALL)  
COLUMNS;
```

where *THE\_TABLE* is the name of the table on which supplemental logging is being enabled. Verify that table-level supplemental logging has been enabled by executing:

```
select count(*) from ALL_LOG_GROUPS where  
LOG_GROUP_TYPE='ALL COLUMN LOGGING' and OWNER=THE_OWNER  
and TABLE_NAME=THE_TABLE
```

where *THE\_OWNER* is the table owner. If this command returns a value of 1, table-level supplemental logging has been enabled for this table.

You can also enable supplemental logging from Replication Agent for Oracle using the **ra\_set\_autocorrection** command as described in the *Replication Agent Reference Manual*.

## DDL Replication

Replication of data definition language (DDL) commands is supported, but only to Oracle databases. You cannot replicate DDL commands from Oracle to non-Oracle replicate databases.

Replication of DDL commands is enabled or disabled in Replication Agent using the **pdb\_setrepddl** command. Replication Agent for Oracle can disable or enable the replication of specific DDL commands by object, owner, statement, or user. Replication Server uses the **ddl\_username** parameter to execute DDL commands in the replicate database as the same user who executed the DDL commands in the primary database.

---

**Note:** By default, Replication Agent for Oracle filters SYS user DDL and if you want to replicate the SYS user DDL, you can manually remove the filter.

---

See *Replication Agent Reference Manual > Command Reference > pdb\_setrepddl* and *Replication Agent Reference Manual > Configuration Parameters > ddl\_username* for details on using **pdb\_setrepddl** and **ddl\_username**.

### DDL parameters

Set the **ddl\_username** and **ddl\_password** parameters.

To replicate DDL in Oracle, use **pdb\_setrepddl** to set filtering rules accordingly. You must also set the Replication Agent **ddl\_username** and **ddl\_password** parameters. **ddl\_username** is the database user name used to execute the replicated DDL command at the target database. This user must have permission to execute all replicated DDL commands at the target database. **ddl\_password** is the corresponding password for **ddl\_username**. In addition, the **ddl\_username** database user must have permission to issue the **ALTER SESSION SET CURRENT\_SCHEMA** command for any primary database user that might issue a DDL command to be replicated. See the *Replication Agent Reference Manual*.

### *Special Usage Notes*

The value of **ddl\_username** cannot be the same as the maintenance user defined in Replication Server for the replicate connection. If these names are the same, a Replication Server error results.

The value of the **ddl\_username** parameter is sent in the LTL for all replicated DDL statements. When DDL is replicated, Replication Server connects to the replicate database using the user ID and password specified by the **ddl\_username** and **ddl\_password** parameters. Replication Server then issues:

```
ALTER SESSION SET CURRENT_SCHEMA=user
```

where *user* is the user ID that generated the DDL operation at the primary database. The actual DDL command is then executed against the replicate database. If the user ID specified in **ddl\_username** does not have permission to issue the **ALTER SESSION SET**

**CURRENT\_SCHEMA** or to execute the DDL command against the user schema, the command fails.

---

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

---

### **DDL Commands and Objects Filtered from Replication**

Some Oracle DDL commands and objects are not replicated.

These DDL commands are not replicated:

- **alter database**
- **alter rollback segment**
- **alter session**
- **alter snapshot**
- **alter snapshot log**
- **alter system**
- **alter tablespace**
- **analyze**
- **audit**
- **create control file**
- **create database link**
- **create pfile from spfile**
- **create rollback segment**
- **create schema authorization**
- **create snapshot**
- **create snapshot log**
- **create spfile from pfile**
- **create tablespace**
- **drop database link**
- **drop rollback segment**
- **drop snapshot**
- **drop snapshot/log**
- **drop tablespace**
- **explain**
- **lock table**
- **no audit**
- **rename**
- **set constraints**
- **set role**
- **set transaction**

Any objects that are owned by SYS are not replicated. Any object owned by users defined in the list of nonreplicated users is not replicated. You can modify this list using the



**pdb\_ownerfilter** command. In addition, Sybase has provided a default list of owners whose objects are not replicated. However, you cannot remove the SYS owner. Use the **pdb\_ownerfilter** command to return, add, or remove the list of owners whose objects are not replicated. See the *Replication Agent Reference Manual*.

---

**Note:** The **truncate table** command is replicated as **rs\_truncate**.

---

## **Character Case of Database Object Names**

Database object names must be delivered to the primary Replication Server in the same format as specified in replication definitions; otherwise, replication fails. For example, if a replication definition specifies a table name in all lowercase, then that table name must appear in all lowercase when it is sent to the primary Replication Server by the Replication Agent.

To control the way Replication Agent treats the character case of database object names sent to the primary Replication Server, set the **ltl\_character\_case** configuration parameter to one of these values:

- **asis** – (the default) database object names are passed to Replication Server in the same format as stored in the primary data server.
- **lower** – database object names are passed to Replication Server in all lowercase, regardless of how they are stored in the primary data server.
- **upper** – database object names are passed to Replication Server in all uppercase, regardless of how they are stored in the primary data server.

In the Oracle data server, by default, database object names are stored in all uppercase. However, if you create a case-sensitive name, the case-sensitivity is retained in Oracle.

These examples use the **asis** option:

- `create table tabA` is stored as TABA
- `create table TabB` is stored as TABB
- `create table 'TaBc'` is stored as TaBc

These examples use the **upper** option:

- `create table tabA` is rendered in LTL as TABA
- `create table TabB` is rendered in LTL as TABB
- `create table 'TaBc'` is rendered in LTL as TABC

## **Format of Origin Queue ID**

Each record in the transaction log is identified by an origin queue ID that consists of 64 hexadecimal characters (32 bytes). The format of the origin queue ID is determined by the Replication Agent instance and varies according to the primary database type.

**Table 1. Replication Agent for Oracle Origin Queue ID Format**

Character	Bytes	Description
0–3	2	Database generation ID
4–15	6	System change number
16–19	2	System change number generation ID
20–23	2	Redo log thread
24–43	10	Redo log record block address
44–55	6	System change number of the oldest active transaction begin
56–63	4	Locator ID

## **LTL Origin Commit Time Granularity**

For Oracle, the precision of the origin commit time does not include milliseconds.

Replication Agent retrieves the origin commit time from the Oracle redo log. Timestamps in the redo log have a granularity only of seconds, not milliseconds.

## **Replication Server and RSSD Scripts**

Replication Agent provides supplemental scripts to support additional Replication Server user-defined datatypes for Oracle datatypes and the replication of DDL commands.

These Replication Server scripts are shipped with Replication Agent and must be applied when the installed Replication Server is version 15.0.1 or earlier:

- `$$SYBASE/RAX-15_5/scripts/oracle/hds_oracle_new_setup_for_replicate.sql`
- `$$SYBASE/RAX-15_5/scripts/oracle/oracle_create_error_class_1_rs.sql`
- `$$SYBASE/RAX-15_5/scripts/oracle/oracle_create_error_class_2_rssd.sql`
- `$$SYBASE/RAX-15_5/scripts/oracle/oracle_create_error_class_3_rs.sql`

Manually run these Replication Server scripts against the RSSD when the installed Replication Server is version 15.0.1 or earlier:

- \$SYBASE/RAX-15\_5/scripts/oracle/hds\_oracle\_funcstrings.sql
- \$SYBASE/RAX-15\_5/hds\_oracle\_udds.sql
- \$SYBASE/RAX-15\_5/hds\_clt\_ase\_to\_oracle.sql

### **Applying Script Changes for User-Defined Datatypes**

Apply these script changes to use Oracle user-defined datatypes.

To use Oracle user-defined datatypes:

1. If your Replication Server is version 15.0.1 or earlier, apply this script to support replication of DDL to an Oracle replicate database:

```
$SYBASE/RAX-15_5/scripts/oracle/
hds_oracle_new_setup_replicate.sql
```

This script defines Replication Server objects that must be created in the replicate database. Use this script instead of the `hds_oracle_setup_replicate.sql` script provided in the Replication Server install directory. This revised script contains additional changes to support Oracle-to-Oracle DDL replication.

2. To correctly define the Oracle error class for Replication Server 15.0.1 or an earlier version:

- Apply this script at Replication Server:  
\$SYBASE/RAX-15\_5/scripts/oracle/  
oracle\_create\_error\_class\_1\_rs.sql
- Apply this script against your RSSD:  
\$SYBASE/RAX-15\_5/scripts/oracle/  
oracle\_create\_error\_class\_2\_rssd.sql
- Apply this script at Replication Server:  
\$SYBASE/RAX-15\_5/scripts/oracle/  
oracle\_create\_error\_class\_3\_rs.sql

See *Replication Server Heterogeneous Replication Guide > Oracle Primary Data Server Issues*.

## **Oracle Datatype Compatibility**

Replication Agent for Oracle processes Oracle transactions and passes data to the primary Replication Server. In turn, the primary Replication Server uses the datatype formats specified in the replication definition to receive the data from Replication Agent for Oracle.

**Table 2. Recommended Oracle Datatype Mapping**

Oracle Datatype	Oracle Length/Range	Sybase Datatype	Sybase Length/Range	Notes
BINARY_DOUBLE	9 bytes, 64-bit single precision floating point number datatype	double	8 bytes	<ul style="list-style-type: none"> <li>Maximum positive finite value is 1.79769313486231E+308.</li> <li>Minimum positive finite value is 2.22507485850720-308.</li> </ul>
BINARY_FLOAT	5 bytes, 32-bit single precision floating point number datatype	rs_oracle_float	4 or 8 bytes, depending on precision	<ul style="list-style-type: none"> <li>Maximum positive finite value is 3.40282E+38F.</li> <li>Minimum positive finite value is 1.17549E-38 F.</li> </ul>
BFILE	4GB, locator points to large binary file	image	2GB	
BLOB	4GB, variable-length binary large object	image	2GB	
BOOLEAN	1 byte	rs_oracle_decimal	17 bytes.	The BOOLEAN datatype is only for use with PL/SQL.
CHAR	255 bytes	char	32K	

Oracle Datatype	Oracle Length/Range	Sybase Datatype	Sybase Length/Range	Notes
CLOB	4GB, variable-length character large object	image or unitext	2GB	For Replication Server 15.0 and later versions, the CLOB datatype maps to unitext. For earlier versions of Replication Server, the NCLOB datatype maps to image.
DATE	8 bytes, fixed-length	datetime or rs_oracle_datetime	8 bytes	<p>Replication Server supports dates from January 1, 1753 to December 31, 9999.</p> <p>Oracle supports dates from January 1, 4712 BC to December 31, 9999 AD.</p> <hr/> <p><b>Note:</b> Use the Replication Server heterogeneous datatype support (HDS) feature for datatype conversion and translation.</p>
INTERVAL DAY(n) TO SECOND(n)	Variable-length	rs_oracle_interval		

Oracle Datatype	Oracle Length/Range	Sybase Datatype	Sybase Length/Range	Notes
INTERVAL YEAR (n) TO MONTH	Variable-length	rs_oracle_interval		
LONG	2GB, variable-length character data	text		
LONG RAW	2GB, variable-length binary data	image		
NCHAR	255 bytes, multibyte characters	unichar or char	32K	
NCLOB	4GB, variable-length multi-byte character large object	unitext or text	2GB	For Replication Server 15.0 and later versions, the NCLOB datatype maps to unitext. For earlier versions of Replication Server, the NCLOB datatype maps to image.

Oracle Datatype	Oracle Length/Range	Sybase Datatype	Sybase Length/Range	Notes
NUMBER (p,s)	21 bytes, variable-length numeric data	float, int, real, number, decimal, or rs_oracle_decimal	float is 4 or 8 bytes. int is 4 bytes. real is 4 bytes. number and decimal are 2 to 17 bytes.	<p>The float datatype can convert to scientific notation if the range is exceeded.</p> <p>Integers (int) are truncated if they exceed the Replication Server range of 2,147,483,647 to -2,147,483,648 or <math>1 \times 10^{-130}</math> to <math>9.99 \times 10^{25}</math>.</p> <p>The number and decimal datatypes are truncated if they exceed the range of <math>-10^{38}</math> to <math>10^{38}-1</math>.</p> <p>Oracle precision ranges from 1 to 38 digits. Default precision is 18 digits.</p> <p>Oracle scale ranges from -84 to 127. Default scale is 0.</p>
NVARCHAR2	2000 bytes, variable-length, multi-byte character data	nvarchar or varchar	32K	

Oracle Datatype	Oracle Length/Range	Sybase Datatype	Sybase Length/Range	Notes
RAW	2000 bytes, variable-length binary data	rs_oracle_binary	32K	
ROWID	6 bytes, binary data representing row addresses	rs_oracle_rowid	32K	
SIMPLE_INTEGER	4 bytes representing signed integers	integer		<p>SIMPLE_INTEGER is new as of Oracle 11g and is only for use with PL/SQL.</p> <hr/> <p><b>Note:</b> Marking procedures with PLS_INTEGER and predefined PL/SQL numeric datatypes other than SIMPLE_INTEGER is not supported.</p>



Oracle Datatype	Oracle Length/Range	Sybase Datatype	Sybase Length/Range	Notes
TIMESTAMP (n)	21-31 bytes, variable-length	datetime or rs_oracle_timestamp9	8 bytes	<p>Replication Server supports dates from January 1, 1753 to December 31, 9999.</p> <p>Oracle supports dates from January 1, 4712 BC to December 31, 4712 AD.</p> <hr/> <p><b>Note:</b> Use the Replication Server heterogeneous datatype support (HDS) feature for datatype conversion and translation.</p>
TIMESTAMP (n) WITH [LOCAL] TIME ZONE	Variable-length	rs_oracle_timestampz		
UDD object type	Variable-length character data	rs_rs_char_raw	32K	
VARCHAR2	2000 bytes, variable-length character data	varchar	32K	
XMLTYPE	4GB, variable-length character large object	text	2GB	XMLTYPE data is implicitly handled as Oracle CLOB data.

### See also

- *Oracle User-Defined Types* on page 33

**Replication Server 15.0 Unsigned Datatype Mapping**

For Replication Server 15.0 and later, unsigned datatypes are supported and can be specified in the replication definitions.

**Table 3. Unsigned Integer Replication Definition Datatype Mapping**

RepServer 15.0 Unsigned Datatypes	Replication Definition Datatypes
unsigned bigint	numeric (20)
unsigned int	numeric (10)
unsigned smallint	int
unsigned tinyint	tinyint

**Oracle ANYDATA Datatype Compatibility**

Replication Agent supports the replication of data stored in ANYDATA columns.

- Both the primary and replicate databases must be Oracle databases.
- Both the primary database table and replicate database tables must have the same ANYDATA columns.

The `pdb_ignore_unsupported_anydata` configuration parameter determines how Replication Agent handles data of unsupported datatypes stored in columns of type ANYDATA. See the *Replication Agent Reference Manual*.

**Oracle XMLTYPE Datatype Compatibility**

Replication Agent supports the replication of XMLTYPE column data stored as CLOB or XML from an Oracle 10g or 11g primary database to an Oracle 10g or Adaptive Server Enterprise replicate database.

**See also**

- *Oracle 10g and 11g XMLTYPE Restrictions* on page 30
- *XMLTYPE Data Replication* on page 59

**Oracle Datatype Restrictions**

Replication Server and Replication Agent impose some constraints on the Oracle NUMBER datatype.

See the *Replication Server Option for Oracle 15.7.1 ESD #2 Release Bulletin* for the latest information on datatype restrictions.

These constraints are:

- In the integer representation:

- The corresponding Sybase `int` datatype has a smaller absolute maximum value. The Oracle `NUMBER` absolute maximum value is 38 digits of precision, between  $9.9 \times 10^{125}$  and  $1 \times 10^{-130}$ . The Sybase `int` value is between  $2^{31} - 1$  and  $-2^{31}$  (2,147,483,647 and -2,147,483,648), inclusive.
- Oracle `NUMBER` values greater than the Sybase `int` maximum are rejected by Replication Server.
- In the floating point representation:
  - The precision of the floating point representation has the same range limitation as the integer representation.
  - If the floating point value is outside the Sybase range of  $2^{31} - 1$  and  $-2^{31}$  (2,147,483,647 and -2,147,483,648), Replication Agent for Oracle converts the number into exponential format to make it compatible with Replication Server. No loss of precision or scale occurs.

Replication Agent does not support replication of these special values for `BINARY_FLOAT` and `BINARY_DOUBLE` datatypes:

- NaN (not a number)
- Inf (positive infinity)
- -Inf (negative infinity)

Replication Server and Replication Agent impose these constraints on the Oracle `TIMESTAMP WITH [LOCAL] TIME ZONE` datatype.

When a `TIMESTAMP WITH TIME ZONE` datatype is replicated, the time zone information is used to resolve the timestamp value to the “local” time zone and then the resolved value is replicated. The time zone information itself is not replicated.

For example, if a `TIMESTAMP WITH TIME ZONE` datatype is recorded in Oracle as “01-JAN-05 09:00:00.000000 AM -8:00” and the “local” time zone is -6:00, the replicated value is “01-JAN-05 11:00:00.000000”. The timestamp value is adjusted for the difference between the recorded time zone of -8:00 and the local time zone of -6:00, and the adjusted value is replicated.

### **Oracle ANYDATA Datatype Restrictions**

Replication Agent does not support the replication of data stored in `ANYDATA` columns under some circumstances.

Replication of `ANYDATA` is not supported when:

- The replicate database table column is not of type `ANYDATA`. An attempt to replicate data stored in an `ANYDATA` column to a column that is not of the `ANYDATA` type causes the Replication Server Data Server Interface (DSI) thread to fail.
- The size of the data stored in an `ANYDATA` column exceeds the maximum size of the Replication Server `opaque` datatype, which is 16K.

## Replication Agent for Oracle

- Replication Agent does not replicate data of these Oracle datatypes or structures stored in a column of type ANYDATA:
  - BFILE
  - Nested tables
  - REF
  - UROWID
  - VARRAY

The **pdb\_ignore\_unsupported\_anydata** configuration parameter determines how Replication Agent handles data of unsupported datatypes stored in columns of type ANYDATA.

See the *Replication Server Reference Manual* for information on replication definitions and **create replication definition**.

See the Oracle SQL Reference guide for a complete list of Oracle-supplied types.

See the *Replication Agent Reference Manual* for information on **pdb\_ignore\_unsupported\_anydata**.

### **Oracle 10g and 11g XMLTYPE Restrictions**

Replication Agent supports the replication of XMLTYPE columns and tables if they are stored as CLOB or XML data. Replication Agent does not support the replication of XMLTYPE data in object-relational XML storage or binary XML storage.

Replication Agent supports the replication of XMLTYPE columns from an Oracle primary database to an Oracle or Adaptive Server Enterprise replicate database, but no other platforms are supported. Replication Agent replicates XMLTYPE tables only from an Oracle primary database to an Oracle replicate database.

---

**Note:** The XMLTYPE datatype is supported only for replication to Oracle 10g replicate databases and from Oracle 10g and 11g primary databases.

---

### **See also**

- *Oracle XMLTYPE Datatype Compatibility* on page 28
- *XMLTYPE Data Replication* on page 59

### **Oracle ROWID Datatype Restrictions**

When Replication Agent replicates ROWID data, the value replicated always represents the value stored in the table in the primary database and has no relationship to the ROWID value in the replicate database. There is no attempt to convert or adjust ROWID data to match the data in the replicate database.

## **Oracle Large Object (LOB) Support**

Oracle LOB data can exist in several formats in Oracle.

The LOB datatypes in Oracle are:

- **Character:**
  - LONG
  - CLOB
  - NCLOB
- **Binary:**
  - LONG RAW
  - BLOB
  - BFILE—points to file contents stored outside of the Oracle database

For those types stored in the database (all types except BFILE), Oracle records the content of the LOB in the redo log. The Replication Agent reads the LOB data from the redo log and submits the data for replication.

Because BFILE type data is stored outside of the database, the BFILE contents are not recorded in the redo log. To replicate the content of a BFILE, the Replication Agent connects to the primary Oracle database and issues a query to select the data from the BFILE. Selecting the BFILE data separately from other data in the redo log can provide a temporary out-of-sync condition if the BFILE contents are changed multiple times.

### **See also**

- *Replication of LOB Columns* on page 31

### **Replication of LOB Columns**

Oracle logs all LOB data (except for BFILE datatypes) in the Oracle redo log. This allows the Replication Agent to apply each individual LOB change. However, for BFILE data, the same technique is used and the same limitation exists—BFILE data is not logged but read from the database at the time the rest of the transaction is processed.

For instructions on enabling and disabling replication for LOB columns, see the *Replication Agent Administration Guide*.

### ***Transaction Integrity and LOB Data***

Because of the way the Replication Agent processes the LOB column data when replicating transactions, transaction integrity may be compromised. For example, if two transactions change the data in a LOB column and the Log Reader does not process the first transaction until after the second transaction has been committed, when the LOB data is read from the primary database, the value of that data is the result of the second transaction. In this event, the value of the LOB data in the first transaction is never sent to the replicate database. After the second transaction is processed by the Log Reader, the primary and replicate databases are

synchronized again, but for a period of time between processing the first and second transactions, the replicate database contains data that does not match the originating transaction.

This problem occurs only when a LOB column is changed more than once by a sequence of transactions. The period of time over which the problem exists may be significant if the replication system throughput is slow or if a replication system component fails. As soon as the last transaction that changes the LOB column is processed at the replicate site, the problem is corrected.

### **Replicating CLOB and NCLOB Datatypes**

Oracle NCLOB (National Character Large Object) is a datatype that stores large character data using a multibyte national character set. Similarly, the CLOB datatype may also store character data using a multibyte national character set, when the Oracle database is defined with a double-byte or variable-width character set.

By default, the byte order of the multibyte characters stored in the NCLOB datatype (and CLOB when the database is defined with a double-byte or variable-width character set) is converted during replication to big-endian byte order. This allows the data to be transmitted over networks using big-endian order, which is the common network byte order.

The datatype in a replication definition for an NCLOB or CLOB should be `untext`. This prevents Replication Server from attempting character set conversion on the data. If the Replication Server version does not support `untext`, use the `image` datatype.

If the target database that is to receive this NCLOB or CLOB data is installed on a little-endian platform, the database may not automatically convert the replicated data from the sent big-endian order to the little-endian order. To support replicating NCLOB or CLOB data to a database server that does not provide the necessary conversion from big-endian (network order) to little-endian, force the byte order to be sent by the Replication Agent using the `lr_ntext_byte_order` parameter to set a value of **big** (for big-endian) or **little** (for little-endian).

The `lr_ntext_byte_order` parameter is important for replication between two databases that reside on different platforms.

---

**Note:** The default behavior of Replication Agent for Oracle is to force any Unicode data to big-endian order as defined by the `ltl_big_endian_untext` configuration parameter. To allow the `lr_ntext_byte_order` configuration parameter to successfully override the Oracle byte order, you must also set the `ltl_big_endian_untext` configuration parameter to false whenever the `lr_ntext_byte_order` parameter is used.

---

The `ltl_big_endian_untext` parameter specifies whether `untext` data should be converted from little-endian to big-endian before sending LTL to the Replication Server. Valid values are true and false. When setting this parameter, you must know how `lr_ntext_byte_order` is set. If `lr_ntext_byte_order` is set to send the correct byte order for the replicate database, the `ltl_big_endian_untext` parameter must be set to false so that the byte order is not changed.

`ltl_big_endian_uni` is true, by default. The `ltl_big_endian_uni` and `lr_nxtxt_byte_order` configuration parameters have differences:

- When `ltl_big_endian_uni` is true, Replication Agent for Oracle sends all Unicode data in big-endian order.
- When `ltl_big_endian_uni` is false, Replication Agent for Oracle allows Unicode data to be sent in a byte order that is used when the data is stored in the transaction log file.

`lr_nxtxt_byte_order` forces the result of Unicode data that is read from the transaction log to be in the correct byte order, regardless of how it normally exists in the transaction log file.

## Oracle User-Defined Types

User-defined datatypes (UDDs) use Oracle built-in datatypes and other user-defined datatypes as building blocks that model the structure and behavior of data in applications.

Replication Agent for Oracle supports replication of user-defined object types. Object types are abstractions of real-world entities, such as purchase orders, that application programs deal with. An object type is a schema object with three kinds of components:

- A name, which identifies the object type uniquely within that schema.
- Attributes, which are built-in types or other user-defined types. Attributes model the structure of the object.
- Methods, which are functions or procedures written in PL/SQL and stored in the database, or written in a language such as C or Java and stored externally. Methods implement operations the application can perform on the object.

### Creating a Datatype Definition in Replication Server

Create a definition for your user-defined datatype.

#### Prerequisites

You must have Replication Server administrator privileges or permission. Also, if you are using Replication Server 15.1 or earlier, see "Replication Server and RSSD Scripts" first.

#### Task

To replicate user-defined datatypes in Oracle, the datatype specified in the replication definition must be `rs_char_raw`.

1. Log in to the RSSD.
2. Add a row to the `rs_datatype` table using this example as a guide:

```
/* rs_oracle_udd raw - char with no delimiters */
insert into rs_datatype values(
0, /* prsid */
0x0000000001000008, /* classid */
'rs_oracle_udd', /* name */
0x0000000000010210, /* dtid */
0, /* base_coltype */
```

## Replication Agent for Oracle

```
255,          /* length */
0,           /* status */
1,           /* length_err_act */
'CHAR',      /* mask */
0,           /* scale */
0,           /* default_len */
'',         /* default_val */
0,           /*-delim_pre_len-*/
'',         /* delim_pre */
0,           /*-delim_post_len-*/
'',         /* delim_post */
0,           /* min_boundary_len */
'',         /* min_boundary */
3,           /* min_boundary_err_act */
0,           /* max_boundary_len */
'',         /* max_boundary_err_act */
0           /* rowtype */
)
go
```

### 3. Restart Replication Server.

### 4. In Replication Server, test the new type:

```
admin translate, 'The quick brown fox jumped over the lazy
dog.', 'char(255)', 'rs_oracle_udd'
```

```
go
```

Delimiter Prefix	Translated	Value	Delimiter Postfix
------------------	------------	-------	-------------------

```
-----
-----
```

NULL	The quick brown fox jumped over the lazy dog.		
NULL			

The new type is defined correctly if the sentence translates correctly.

### See also

- *Replication Server and RSSD Scripts* on page 20

### **Example: Create a Replication Definition**

This example demonstrates how to create a replication definition using the `rs_char_raw` type defined in Replication Server.

These Oracle table and type definitions are used in the example:

- Oracle UDD object type name: `NAME_T`
- Oracle table name: `USE_NAME_T`
- Oracle table columns: `PKEY INT, PNAME NAME_T`

```
create replication definition use_name_t_repdef
with primary at ra_source_db.ra_source_ds
with all tables named 'USE_NAME_T'
```



```
(
  PKEY int,
  PNAME rs_rs_char_raw
)
primary key (PKEY)
searchable columns (PKEY)
go
```

---

**Note:** For this example, `lfl_character_case` must be `upper`.

---

### **Object Type Attribute Replication**

To replicate updates to user-defined object type attributes, Replication Agent must enable table-level supplemental logging. Table-level supplemental logging can be enabled manually.

Replication Agent also attempts to enable this logging when marking a table that contains a user-defined object type. However, for Replication Agent to mark such a table, there must already be an Oracle user specified by the `pds_username` parameter that has ALTER permission granted for the table.

If table-level supplemental logging has not been enabled for a table containing a user-defined object type and Replication Agent encounters an update log record in the Oracle log, Replication Agent changes its status from Replicating to Admin with this error:

```
There is insufficient column data in the log to support Oracle UDD
update command processing. Please make sure table-level supplemental
logging is enabled.
```

In this case, use the `pdb_skip_op` to skip this log record. See the *Replication Agent Reference Manual*.

## **Sequence Marking and Unmarking**

Support for Oracle sequence replication is supported only for replication to Oracle. No support is provided for replicating a sequence value to a non-Oracle replicate database.

Replication Agent supports replication of sequences in the primary database. To replicate a sequence invoked in a primary database, the sequence must be marked for replication, and replication must be enabled for that sequence. This is analogous to marking and enabling replication for tables.

---

**Note:** Marking a sequence for replication is separate from enabling replication for the sequence. If the value of the `pdb_dflt_object_repl` parameter is true, replication is enabled automatically at the time a sequence is marked.

---

Oracle does not log information every time a sequence is incremented. Sequence replication occurs when the Replication Agent captures the system table updates that occur when the sequence's cache is refreshed. Therefore, the sequence value replicated when a sequence is marked for replication is the “next” sequence value to be used when the current cache expires. The result is that not every individual increment of a sequence is replicated, but the replicate site always has a value greater than the primary site's currently available cached values.

To temporarily suspend replication of a marked sequence, you can disable replication for the sequence.

### See also

- *Sequence Replication Enabling and Disabling* on page 38

### **Replication Server Changes to Support Sequence Replication**

By default, Replication Server does not support replication of Oracle sequence objects. You must make changes to Replication Server and the replicate Oracle database before you can replicate Oracle sequences.

For Replication Server, you must create a replication definition that defines a stored procedure to assist with sequence replication. Execute the `$$SYBASE/RAX-15_5/scripts/oracle/oracle_create_rs_sequence_repdef.sql` script against your primary Replication Server after editing the script to replace values `{pds}` and `{pdb}` with the name of your primary Replication Server connection. You can find these values in the `rs_source_ds` and `rs_source_db` Replication Agent configuration properties.

---

**Note:** The replication definition assumes that a database replication definition exists. You may need to alter the definition if a database replication definition does not exist. For details, see comments in the `oracle_create_rs_sequence_repdef.sql` script.

---

In the replicate Oracle database, you must create a stored procedure to support sequence replication. Log in to the replicate Oracle database as the maintenance user defined in your Replication Server connection to the replicate database. Execute the `$$SYBASE/RAX-15_5/scripts/oracle/oracle_create_replicate_sequence_proc.sql` script to create the necessary stored procedure.

---

**Note:** The maintenance user defined in the Replication Server connection to the replicate database must have sufficient privileges to execute functions in the Oracle DBMS\_SQL package. Also, this maintenance user must have authority at the replicate Oracle database to update any replicated sequence.

---

### **Marking a Sequence for Replication**

Mark a sequence for replication.

1. Log in to the Replication Agent instance with the administrator login.
2. Determine whether or not the sequence is marked in the primary database:

```
pdb_setrepseq pdb_seq
```

where `pdb_seq` is the name of the sequence you want to mark for replication.

- If `pdb_setrepseq` returns information that the specified sequence is marked, you do not need to continue this procedure.

- If **pdb\_setrepseq** returns information that the specified sequence is not marked, continue this procedure to mark the sequence for replication.

### 3. Mark the sequence for replication.

**pdb\_setrepseq** allows you to mark the primary sequence to be replicated and specify a different sequence name to use in the replicate database.

- If the sequence name you want to increment at the replicate site has the same name as at the primary site, use this command to mark the sequence for replication:

```
pdb_setrepseq pdb_seq, mark
```

---

**Note:** Replicating a sequence with a different name than is provided is consistent with other marking commands but is not a typical configuration.

---

- To mark the sequence for replication using a different sequence name, use:

```
pdb_setrepseq pdb_seq, rep_seq, mark
```

where *rep\_seq* is the name of the sequence that you want to increment in the replicate database.

---

**Note:** Replicating sequence values to a sequence with a different name at the replicate site assumes that the replicate site sequence has the same attributes and starting value as the primary site's sequence.

---

- If the value of **pdb\_dflt\_object\_repl** is true, the sequence marked for replication with **pdb\_setrepseq** is ready for replication after you successfully invoke **pdb\_setrepseq**.
- If the value of **pdb\_dflt\_object\_repl** is true (the default value), skip step 4.
- If the value of **pdb\_dflt\_object\_repl** is false, you must enable replication for the sequence before replication can take place.

### 4. Enable replication for the sequence:

```
pdb_setrepseq pdb_seq, enable
```

After replication is enabled for the sequence, you can begin replicating invocations of that sequence in the primary database.

---

**Note:** To replicate a sequence, you must also run the `oracle_create_replicate_sequence_proc.sql` script in the `$$SYBASE/RAX-15_5/scripts/oracle` directory at the replicate site to create a procedure named **rs\_update\_sequence**.

---

## **Unmarking a Sequence**

Unmark a sequence.

1. Log in to the Replication Agent instance with the administrator login.
2. Determine whether or not the sequence is marked in the primary database:

```
pdb_setrepseq pdb_seq
```

where *pdb\_seq* is the name of the sequence you want to unmark.

- If **pdb\_setrepseq** returns information that the specified sequence is marked, continue this procedure to unmark the sequence.
- If **pdb\_setrepseq** does not return information that the specified sequence is marked, you do not need to continue this procedure.

### 3. Disable replication for the sequence:

```
pdb_setrepseq pdb_seq, disable
```

### 4. Remove the replication marking from the sequence:

```
pdb_setrepseq pdb_seq, unmark
```

To force the unmarking, use:

```
pdb_setrepseq pdb_seq, unmark, force
```

### 5. Confirm that the sequence is no longer marked for replication:

```
pdb_setrepseq pdb_seq
```

## **Sequence Replication Enabling and Disabling**

To temporarily suspend replication of a sequence, use **pdb\_setrepseq** to disable replication for the marked sequence. When you are ready to resume replication of the marked sequence, use **pdb\_setrepseq** again to enable replication.

---

**Note:** By default, no sequences are marked for replication.

---

To replicate updates of a sequence in the primary database, the sequence must be marked for replication and, replication must be enabled for that sequence.

Marking a sequence for replication is separate from enabling replication for the sequence.

### **See also**

- *Marking a Sequence for Replication* on page 36

### **Enabling Replication for a Marked Sequence**

Enable replication for a marked sequence.

1. Log in to the Replication Agent instance with the administrator login.
2. Determine whether or not replication is enabled for the sequence:

```
pdb_setrepseq pdb_seq
```

where *pdb\_seq* is the name of the sequence for which you want to enable replication.

If **pdb\_setrepseq** returns information that the specified sequence is marked and has replication disabled, continue this procedure to enable replication for the sequence.

---

**Note:** A sequence must be marked for replication before replication can be enabled or disabled for the sequence.

---

3. Enable replication for the sequence:

```
pdb_setrepseq pdb_seq, enable
```

After replication is enabled for the sequence, any invocation of that sequence is replicated.

4. Use **pdb\_setrepseq** again to verify that replication is now enabled for the sequence:

```
pdb_setrepseq pdb_seq
```

### **Disabling Replication for a Marked Sequence**

Disable replication for a marked sequence.

1. Log in to the Replication Agent instance with the administrator login.
2. Determine whether or not replication is enabled for the sequence:

```
pdb_setrepseq pdb_seq
```

where *pdb\_seq* is the name of the sequence for which you want to disable replication.

If **pdb\_setrepseq** returns information that the specified sequence is marked and has replication enabled, continue this procedure to disable replication for the sequence.

---

**Note:** A sequence must be marked for replication before replication can be enabled or disabled for the sequence.

---

3. Disable replication for the sequence:

```
pdb_setrepseq pdb_seq, disable
```

After replication is disabled for the sequence, any invocation of that sequence is not captured for replication until replication is reenabled.

4. Use **pdb\_setrepseq** again to verify that replication is now disabled for the sequence:

```
pdb_setrepseq pdb_seq
```

## **Setting Up Replication Agent and Oracle on Different Machines**

Run Replication Agent and the primary data server on different machines.

1. Install Replication Agent on a machine of the same hardware and operating system as the machine on which the primary data server is running.
2. Install the JDBC driver on the same machine as Replication Agent and add its path to the CLASSPATH environment variable.
3. If the `timezone.dat` file is not accessible to both machines, copy the `$ORACLE_HOME/oracle/timezone.dat` file to the Replication Agent machine.

---

**Note:** Be sure to copy the `timezone.dat` file of the Oracle server that Replication Agent is replicating from.

---

4. Set the Replication Agent **pdb\_timezone\_file** configuration parameter to the full path name of the `timezone.dat` file.
5. If Replication Agent for Oracle is configured to truncate Oracle archive logs directly, make sure the Oracle archive logs are accessible to both machines. Use the **ra\_devicepath** command to point Replication Agent to the log files.

### Real Application Clusters (RAC)

Replication Agent for Oracle provides support for Oracle 10g and 11g RAC environments. When a Replication Agent for Oracle instance is initialized, the Oracle database is queried to determine how many nodes are supported by the cluster. Based on this information, Replication Agent automatically configures itself to process the redo log information from all nodes.

To process the redo log data from all nodes in an Oracle RAC cluster, the Replication Agent must execute from a location that has access to the same shared storage used by the Oracle nodes to store their redo data.

Configure Replication Agent to connect to a single Oracle instance by supplying the required host, port, and Oracle SID values to the **pds\_host\_name**, **pds\_port\_number** and **pds\_database\_name** configuration parameters. However, in an Oracle RAC environment, Replication Agent must be able to connect to any node in the cluster in the event that a node fails or otherwise becomes unavailable. To support the configuration of multiple node locations, Replication Agent supports connectivity to all possible RAC nodes by obtaining needed information from an Oracle `tnsnames.ora` file for one specified entry. As a result, instead of configuring individual host, port and instance names for all nodes, Replication Agent only requires the location of a `tnsnames.ora` file and the name of the TNS connection to use.

Sybase recommends that you point Replication Agent to a `tnsnames.ora` entry that contains the address for all nodes in the cluster.

For example, if this entry exists in a `tnsnames.ora` file for a three-node cluster, instruct Replication Agent to use that entry by providing the `tnsnames.ora` file location to the **pds\_tns\_filename** configuration property and specifying **RAC10G** as the value for the **pds\_tns\_connection** configuration property:

```
RAC10G =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (LOAD_BALANCE = yes)
      (FAILOVER = ON)
      (ADDRESS = (PROTOCOL = TCP) (HOST = www.xxx.yyy.zz1)
        (PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP) (HOST = www.xxx.yyy.zz2)
        (PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP) (HOST = www.xxx.yyy.zz3)
        (PORT = 1521))
    )
  )
```

```

(CONNECT_DATA =
  (SERVER = DEDICATED)
  (SERVICE_NAME = rac10g)
)
)

```

The `tnsnames.ora` file must also contain a connect descriptor for each node in the cluster:

```

NODE1-VIP =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = www.xxx.yyy.zz1)
    (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = rac10g)
      (INSTANCE_NAME = node1-vip)
    )
  )
)
NODE2-VIP =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = www.xxx.yyy.zz2)
    (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = rac10g)
      (INSTANCE_NAME = node2-vip)
    )
  )
)
NODE3-VIP =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = www.xxx.yyy.zz3)
    (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = rac10g)
      (INSTANCE_NAME = node3-vip)
    )
  )
)
)

```

If the `tnsnames.ora` file does not contain a connect descriptor for each node, Replication Agent cannot generate commit records correctly.

See the *Replication Agent Reference Manual* for details about **pds\_tns\_filename** and **pds\_tns\_connection**.

---

**Note:** Replication Agent must have read access to the `tnsnames.ora` file.

---

### **pdb\_archive\_path**

The **pdb\_archive\_path** configuration parameter identifies where Replication Agent expects to find archived Oracle redo log files.

In an Oracle RAC environment, each Oracle instance can be configured to point to one or more archive log destinations. To support replication, all instances in the Oracle RAC cluster must provide a copy of their archive log files to a shared location that Replication Agent for Oracle

can use to access all archived redo logs. Configure **pdb\_archive\_path** to point to a location to which all Oracle instances write archived log data. Replication Agent must have read access to this directory and all archived redo logs within that directory.

---

**Note:** Archived redo logs can also be stored within ASM.

---

You can configure Replication Agent to remove archived logs from the location specified by **pdb\_archive\_path**, using the **pdb\_archive\_remove** configuration parameter. This allows Replication Agent to remove archived log files that are no longer needed to support replication. If **pdb\_archive\_remove** is set to true, Replication Agent must have update authority to the archive log directory and delete authority on the individual archive log files. If **rman\_enabled** or the **pdb\_archive\_path** configuration parameter is pointed to an ASM path, then update authority to the archive log directory and delete authority on the individual archive logs are not required.

---

**Note:** The **rman\_enabled** parameter enables Replication Agent to use the Oracle **RMAN** utility to truncate old archive log files. See the *Replication Agent Reference Manual*.

---

### See also

- *Automatic Storage Management* on page 42
- *Marking a Sequence for Replication* on page 36

### Oracle Instance Failover

If the Oracle instance to which Replication Agent is connected fails for any reason, Replication Agent attempts to reconnect to any surviving instance, choosing from the list of instances defined in the `tnsnames.ora` file entry.

No manual intervention or configuration is required. If none of the instances are available, Replication Agent reports an error and continues processing as long as redo log file information is still available.

## Automatic Storage Management

Replication Agent for Oracle supports the use of the Oracle Automatic Storage Management (ASM) feature. ASM provides file system and volume management support for an Oracle database environment. ASM can be used in both Real Application Cluster (RAC) and non-RAC environments.

### Archive Log Removal and Configuration

Archive logs that are managed by ASM can be removed from ASM when they are no longer needed by Replication Agent for Oracle.

When **pdb\_archive\_remove** is set to true and the archive logs are managed by ASM, **pdb\_archive\_path** must be set to the name of the ASM disk group in which the archive logs are stored. The disk group name must be preceded with a plus sign (+) indicating that the path is an ASM path. For example:



```

pdb_archive_remove=true
pdb_archive_path=+DISK_GROUP1

```

Archive logs stored in and managed by ASM are owned by the corresponding unique Oracle database name. If the Oracle database name differs from the global unique database name, **pdb\_archive\_path** must be set to both the name of the ASM disk group and the globally unique name of the database in which the archive logs are stored:

```

pdb_archive_path=+DISK_GROUP1/database_name

```

In addition to automatic removal of archive logs from ASM, manual removal is supported with **pdb\_truncate\_xlog**. **pdb\_archive\_path** must be set to the ASM disk group name and preceded with a plus (+) sign for archive logs to be manually removed.

---

**Note:** **rman\_enabled** enables Replication Agent to use the Oracle **RMAN** utility to truncate old archive log files. See the *Replication Agent Reference Manual*.

---

### **Configuration Parameters**

Some configuration parameters must be set if your log files are being managed by ASM.

These configuration parameters are:

- **asm\_password**
- **asm\_tns\_connection**
- **asm\_tns\_filename**
- **asm\_username**

The ASM user ID for **asm\_username** must have **sysdba** permission. For Oracle 10g or 11g, set **asm\_username** as follows:

```

asm_username="sys as sysdba"

```

Alternately, for Oracle 11g, you can set **asm\_username** as follows:

```

asm_username="sys as sysasm"

```

See the *Replication Agent Reference Manual*.

### **Replication Server set autocorrection Command**

The Replication Server **set autocorrection** command prevents failures that would otherwise be caused by missing or duplicate rows in a replicated table.

The **set autocorrection** command corrects discrepancies that may occur during materialization by converting each **update** or **insert** operation into a **delete** followed by an **insert**.

To set autocorrection from:

- Replication Agent for one or all tables in the primary database, use the Replication Agent **ra\_set autocorrection** command as described in the *Replication Agent Reference Manual*.

## Replication Agent for Oracle

- Replication Server, use the **set autocorrection** command in a replication definition. You must do this from Replication Server, however, because Replication Agent cannot alter the autocorrection setting on a replication definition.

Replication Agent for Oracle does not support use of the autocorrection feature for large-object (LOB), LONG, LONG RAW, or user-defined datatypes. Also, **pds\_username** must have the **ALTER ANY TABLE** privilege to execute these commands:

- ALTER TABLE *tablename* ADD SUPPLEMENTAL LOG DATA (ALL) COLUMNS;
- ALTER TABLE *tablename* DROP SUPPLEMENTAL LOG DATA (ALL) COLUMNS;

## Partitioned Tables

Replication Agent supports Oracle partitioning functionality.

Partitioning allows a table, index, or index-organized table to be subdivided into smaller pieces, where each piece of such a database object is called a partition. Each partition has its own name and may optionally have its own storage characteristics. Any table can be partitioned into many separate partitions except those tables containing columns with LONG or LONG RAW datatypes.

Unstructured data (such as images and documents) stored in a LOB column in the database can also be partitioned. When a table is partitioned, all the columns reside in the tablespace for that partition, with the exception of LOB columns, which can be stored in their own tablespace. For additional information about Oracle partitioning, see the Oracle Database VLDB and Partitioning Guide at [http://download.oracle.com/docs/cd/B28359\\_01/server.111/b32024/toc.htm](http://download.oracle.com/docs/cd/B28359_01/server.111/b32024/toc.htm).

### Replication of the truncate partition Command

Replication Agent supports replication of the **truncate partition** command.

Replicate the **truncate partition** command either by:

- Using **lr\_send\_trunc\_partition\_ddl**
- Wrapping **truncate partition** in a stored procedure and replicating the procedure

#### *Use lr\_send\_trunc\_partition\_ddl*

Use the Replication Agent configuration parameter **lr\_send\_trunc\_partition\_ddl** to determine whether **truncate partition** commands are sent as DDL or DML to the replicate database. The configuration can be:

- true (default) – the truncate partition command is sent as a DDL command (**alter table**). Use this setting to replicate to Oracle.
- false – the truncate partition is sent as a DML operation. Use this setting when replicating to databases that treat **truncate partition** commands as DML.

For information about Replication Agent configuration properties, see the *Replication Agent Reference Manual*.

### *Wrap the **truncate partition** command*

You can wrap the **truncate partition** command in a stored procedure definition and replicate the procedure.

For example, to replicate **truncate partition** commands from an Oracle primary to an Adaptive Server Enterprise replicate, create this stored procedure at the primary database:

```
create procedure sp_truncate_partition
as
begin
execute immediate 'ALTER TABLE myTable TRUNCATE PARTITION part1';
end;
```

Create a corresponding stored procedure at the replicate database:

```
create proc sp_truncate_partition as
truncate table myTable part1
```

Mark the **sp\_truncate\_partition** procedure for replication. When **sp\_truncate\_partition** is executed at the primary database, the **truncate partition** command is replicated to the replicate database.

## **Materialized Views**

A materialized view is a stored view query result.

The data on which the view is defined is referred to as the master table (or tables). The materialized view is stored in its own table, which is refreshed based on changes to the master table. A materialized view may be local, in which it is defined on the same database as the master table, or remote, in which the materialized view is defined on a different database than the master table.

Oracle supports these types of materialized views:

- Read-only – materialized view content is derived from the corresponding master table or tables, and the view content cannot be changed.
- Writeable – materialized view content can be changed temporarily, but any changes are overwritten when the table containing the materialized view is refreshed based on changes to the corresponding master table or tables.
- Updatable – updates made to a materialized view are written back to the corresponding master table or tables when the materialized view is refreshed.

For a complete description of materialized views, see the Oracle documentation.

### **Replication and Materialized Views**

An Oracle materialized view allocates space to hold the result set of its base query. Replication Agent can replicate transactions involving the data on which a materialized view is defined as well as on the materialized view itself.

#### *Materialized View DDL*

By default, Replication Agent does not replicate Oracle DDL commands used for materialized views, for example, **CREATE MATERIALIZED VIEW**, **ALTER MATERIALIZED VIEW**, or **DROP MATERIALIZED VIEW**. Materialized view DDL commands are disabled from replication unless otherwise specified using the **pbdd\_setrepddl** command. See the *Replication Agent Reference Manual > Command Reference > pbdd\_setrepddl*.

#### *Materialized Views at the Primary and Replicate Databases*

A materialized view may exist on both the primary database and the replicate database. Such a situation might arise, for example, if materialized view DDL has been enabled for replication with the **pbdd\_setrepddl** command or if the replicate database has been materialized from a primary database dump.

If the master table on which the materialized view is defined exists in the primary database, Replication Agent replicates this master table. The materialized view at the replicate database refreshes according to the contents of the replicated master table. Under no circumstances does Replication Agent replicate the table in which a materialized view is stored in the primary database, and you should not attempt to replicate such a table.

If the materialized view is remote—meaning that the master table on which the materialized view is defined does not exist in the primary database—the materialized view at the replicate database must be redirected so that it points to the database on which the master table is located. If the replicate database is not redirected, a refresh of the materialized view fails at the replicate database. In redirecting the replicate database, re-create the Oracle database link that the replicate database uses to connect to the database containing the master table.

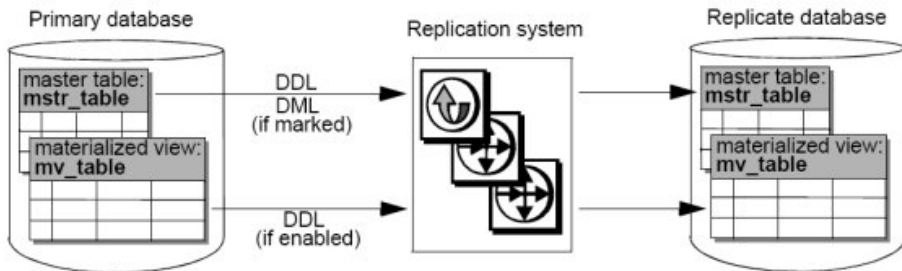
#### *Writeable and Updatable Materialized Views*

Instead of replicating changes to the table containing a materialized view, Replication Agent replicates changes to the master table if the master table has been marked for replication. Replication Agent therefore does not replicate changes made to a writeable materialized view. However, because changes made to an updatable materialized view are written back to the corresponding master table or tables when the materialized view is refreshed, Replication Agent replicates changes made to an updatable materialized view on the primary database to the corresponding master table on the replicate database. Changes made to an updatable materialized view on the replicate database affect only the local master table unless bidirectional replication has been enabled.

### Materialized View Replication Scenarios

In this figure, a materialized view and a corresponding master table reside on both the primary database and the replicate database.

**Figure 2: Master Table and Materialized View on Primary Database**



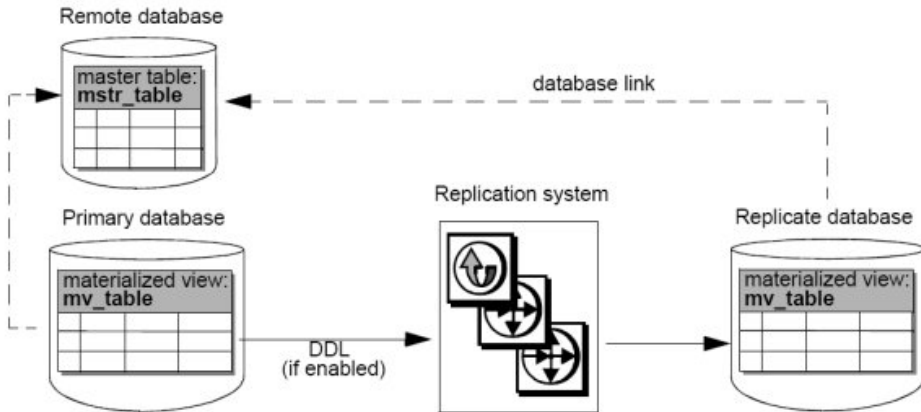
In this situation, DDL commands affecting the master table can be replicated as well as objects affected by the DML that are marked for replication.

DDL commands affecting the materialized view are not replicated unless such DDL is enabled with the **pdb\_setrepddl** command. Since a materialized view also exists on the replicate database, all master tables on which the materialized view is defined must also be replicated. Otherwise, the contents of the materialized view on the replicate database may become invalid.

If the materialized view on the primary database is updatable, changes made to this view are written back to the corresponding master table and, if the master table has been marked for replication, replicated to the replicate database. If the materialized view on the replicate database is updatable, changes made to this view are written back to the corresponding master table on the replicate database, but the master table on the primary database is not changed accordingly unless bidirectional replication has been enabled.

In this figure, the master table on which the primary database materialized view is defined resides on a different, or remote, database.

**Figure 3: Master Table on Remote Database, Materialized View on Primary Database**

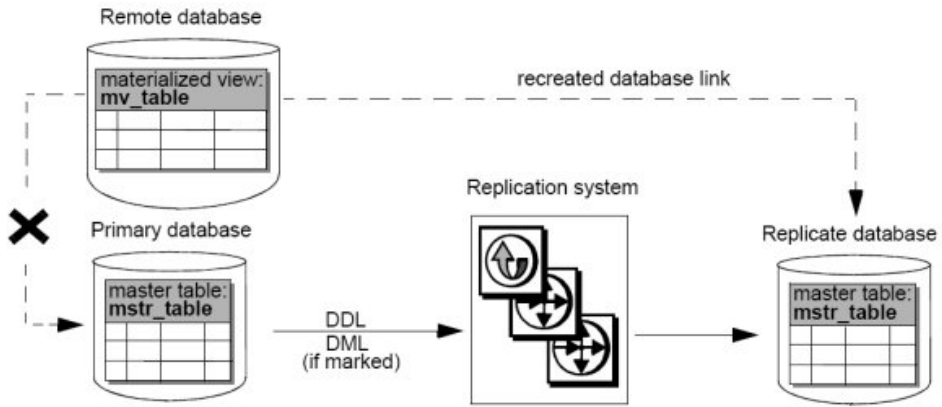


In this situation, neither DML nor DDL affecting the master table is replicated. DDL commands affecting the materialized view are not replicated unless such DDL is enabled with the **pdb\_setrepddl** command. Since the materialized view also exists on the replicate database, a database link must be created so that it points to the database containing the master table on which the materialized view is defined.

If the materialized views on the primary and replicate databases are both updatable and are properly linked to the master table at the remote database, changes made to one of these views are written back to the master table, and the changes are reflected in both materialized views upon refresh.

In the figure below, a materialized view resides on a remote database, and the master table on which the materialized view is defined resides on the primary database. A copy of this master table also resides on the replicate database, and the database link between the remote and primary databases is subsequently broken.

**Figure 4: Master Table on Primary Database, Materialized View on Remote Database**

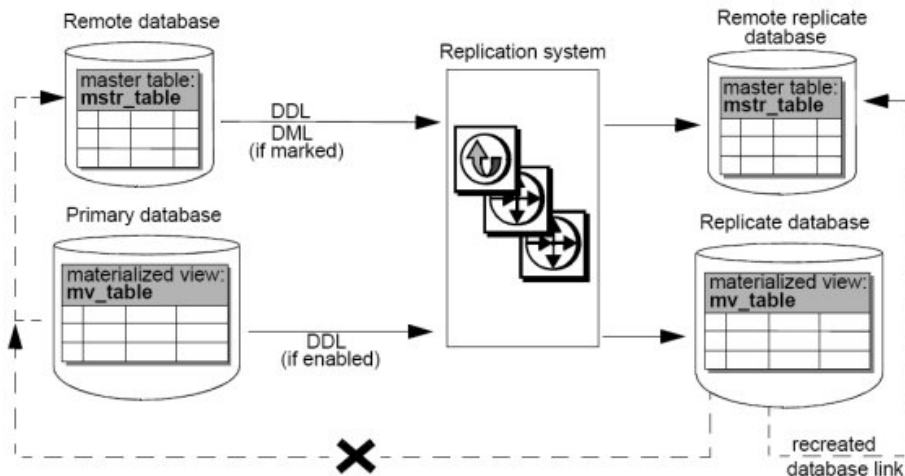


In this situation, DDL commands affecting the master table at the primary database can be replicated as well as DML commands that are marked for replication. DDL commands affecting the materialized view cannot be replicated because there is no corresponding materialized view on the replicate database. If the database link between the remote and primary databases is broken, the remote database must create a link to the replicate database before a refresh occurs.

If the materialized view on the remote database is updatable, changes made to this view are written back to the master table on the database to which the remote database is currently linked.

In the figure below, a master table resides on two different remote databases, one of which is a replicate database. A materialized view resides on the primary database and the replicate database. The materialized view at the replicate database is initially defined by the master table on the remote database, but its database link becomes broken, and the replicate database re-creates a link to the remote replicate database instead.

**Figure 5: Master Tables on Remote Databases, Materialized Views on Primary and Replicate Databases**



If the materialized view on the replicate database is updatable, changes made to this view are written back to the master table on the database to which the replicate database is currently linked. Before the database link between the replicate database and the remote database becomes broken, updates to the materialized view on the replicate database are written back only to the master table on the remote database. After a link is created between the replicate database and the remote replicate database, updates to the materialized view on the replicate database are written back only to the master table on the remote replicate database.

## Index-Organized Tables

Replication Agent can replicate DML for Oracle index-organized tables (IOTs).

DML for these types of index-organized tables can be replicated for Oracle 10g and 11g databases:

- Simple IOTs
- IOTs with **including** and **overflow** clauses
- IOTs with composite partitions
- IOTs with mapping tables
- Index-compressed IOTs
- IOTs with row dependencies
- IOTs with large objects
- IOTs with secondary indexes

Replication Agent cannot replicate IOTs with nested tables and columns of type `VARRAY`.



## Replicate Database Trigger Execution Control

Trigger execution in a replication system can cause problems when both the transaction that caused a trigger to fire and the transactions that resulted from the trigger are replicated.

This may result in data duplication in the case where a trigger causes data to be recorded twice for a single operation performed on an audited table. It may also result in data inconsistency in the case where a trigger results in DML commands being performed twice: once as a result of the trigger firing at the primary database and a second time as a result of the trigger firing at the replicate database to which trigger-altered data has already been replicated. To avoid data duplication and inconsistency, it is important to control trigger execution in the replication system. However, Oracle provides no session-level command to disable trigger execution.

Replication Server allows you to disable trigger execution at the session or connection level. You can control trigger firing each time a PL/SQL command is executed against the replicate database. Controlling trigger execution at the replicate database eliminates data duplication and inconsistency caused by the absence of any trigger control at the replicate database.

For a complete description of the Replication Server **rs\_triggers\_reset** function, see the *Replication Server Reference Manual*. For complete instructions on controlling trigger execution at the replicate database, see "Oracle Replicate Data Server Issues" in the *Replication Server Heterogeneous Replication Guide*.

## Alteration of Replication Definitions from the Primary Data Server

You can alter replication definitions from the primary data server.

To avoid having to quiesce the replication system before altering a replication definition, you can issue the Replication Server **alter replication definition** command from the primary data server and make schema changes to primary database objects at the same time. The propagation of changes to a replication definition can be automatically coordinated with data replication without having to stop the replication process.

To issue the Replication Server **alter replication definition** command from the primary data server, create a stored procedure named **rs\_send\_repserver\_cmd** in the primary Oracle database. The SQL for creating this procedure is contained in the appropriate connection profile on Replication Server. For a list of connection profiles, use the Replication Server **admin show\_connection\_profiles** command.

For a full description of **rs\_send\_repserver\_cmd** and the **alter replication definition** Replication Server commands, see the *Replication Server Reference Manual*.

### Security Considerations

When the **rs\_send\_repserver\_cmd** procedure is invoked at the primary data server, Replication Agent passes corresponding Replication Command Language (RCL) directly to Replication Server. You should therefore consider carefully to whom execution privileges are

## Replication Agent for Oracle

assigned for the `rs_send_repserver_cmd` procedure, and assign privileges as appropriate for your environment and security policy.

### Limitations

You cannot use the `rs_send_repserver_cmd` procedure to alter replication definitions for tables that contain columns of certain datatypes.

These types are:

- BINARY ROWID
- BINARY UROWID
- DATE
- INTERVAL YEAR TO MONTH
- INTERVAL DAY TO SECOND
- MLSLABEL
- RAW
- REF
- TIMESTAMP
- TIMESTAMP WITH LOCAL TZ
- TIMESTAMP WITH TZ

---

**Note:** If you manually change a table-level replication definition in Replication Server, you must then suspend and resume replication in the Replication Agent to ensure that the Replication Agent clears and refreshes its cache.

---

## Oracle Data Guard

Data Guard is a disaster-protection architecture consisting of a primary Oracle database and one or more standby Oracle databases to which the primary database copies its data. These standby copies can be used in the event that the primary Oracle database fails. Replication Agent supports replication of data from an Oracle database system that uses Data Guard.

### *Recommended Configuration*

Although you can configure Replication Agent to replicate data from either the Data Guard primary database or a Data Guard standby database, Sybase recommends that you configure Replication Agent to read from a Data Guard standby database transaction log. This way, if the primary Data Guard database fails over to the standby Data Guard database, Replication Agent is already connected to a working Oracle database, and replication is not disrupted.

If you configure Replication Agent to read from a Data Guard standby database transaction log, `pds_username` must have **alter database** permission.

## **Database Resynchronization**

You can avoid having to quiesce the primary database when you initialize Replication Agent for Oracle if your replication system is configured for database resynchronization.

For information on configuring database resynchronization, see the *Replication Server Heterogeneous Replication Guide > Oracle Replicate Database Resynchronization*.

## **Oracle Transaction and Operation Troubleshooting**

The **ra\_dumptran** and **ra\_helpop** commands return information for use in troubleshooting a specified Oracle database transaction or database operation, respectively.

The **ra\_dumptran** and **ra\_helpop** commands use information gathered by Oracle LogMiner to help you troubleshoot Replication Agent for Oracle. Oracle LogMiner consists of Oracle procedures and views that allow you to obtain detailed information about database activities from the Oracle redo logs. To use **ra\_dumptran** and **ra\_helpop**, you must install Oracle LogMiner, or these commands return errors. For details on using these commands, see the *Replication Agent Reference Manual*.

### **See also**

- *Configuring Oracle LogMiner* on page 15

## **Stored Procedure Replication with BOOLEAN Arguments**

Learn how to replicate Oracle stored procedures with arguments of boolean type.

To replicate an Oracle stored procedure with an argument of type `BOOLEAN`, Replication Agent sends the `BOOLEAN` argument to Replication Server as an integer. Replication Server then uses a function string to convert the argument back to a `BOOLEAN` value so that the stored procedure can be executed on the replicate database. On Replication Server, you must manually create this function string for each Oracle stored procedure that has an argument of type `BOOLEAN`.

Some replicate databases do not support `BOOLEAN` stored procedure arguments. In these cases, the Oracle stored procedure `BOOLEAN` argument should be mapped to an integer argument in the corresponding stored procedure at the replicate database. A function string is then unnecessary.

These examples illustrate how to replicate an Oracle stored procedure with an argument of type `BOOLEAN` to a Oracle replicate database and to a non-Oracle replicate database.

### **Example: Replicating to an Oracle Replicate Database**

Replicate a stored procedure with `BOOLEAN` arguments to an Oracle replicate database.

To replicate a stored procedure defined by these PL/SQL statements:

```
CREATE PROCEDURE boolproc (a IN BOOLEAN, b INT) AS
BEGIN
```

```
        IF (a = true) THEN
            DBMS_OUTPUT.PUT_LINE('True');
        ELSE
            DBMS_OUTPUT.PUT_LINE('False or NULL');
        ENDIF;
    END;
```

1. Manually create a replication definition on Replication Server using this RCL command:

```
create function replication definition ra$xxx_boolproc
    with primary at myprimary.pdb
    with all functions named boolproc (
        @"a" rs_oracle_decimal
        @"b" rs_oracle_decimal )
    searchable parameters(@"a", @"b")
    send standby all parameters
```

**Note:** If the Replication Agent `pdb_auto_create_repdefs` configuration parameter is set to true, a replication definition will be created automatically.

2. Mark the stored procedure for replication:

```
pdb_setrepproc boolproc, mark
```

3. Create a function string on Replication Server:

```
create function string ra$xxx_boolproc.boolproc
    for rs_oracle_function_class
    output language
        'begin execute immediate "begin
ra_user.boolproc
(?a!param?=1,?b!param?);;end;;";;end;;'
    go
```

4. Create a subscription on Replication Server for the replication definition:

```
create subscription sub_intproc
    for ra$xxx_boolproc
    with replicate at myreplicate.rdb
    go
```

The stored procedure is replicated when it is executed at the primary database:

```
EXECUTE boolproc(true,1);
```

### **Example: Replicating to a Non-Oracle Database**

Replicate a stored procedure with BOOLEAN arguments to a non-Oracle replicate database.

To replicate a stored procedure defined by these PL/SQL statements:

```
CREATE PROCEDURE boolproc (a IN BOOLEAN, b INT) AS
BEGIN
    IF (a = true) THEN
        DBMS_OUTPUT.PUT_LINE('True');
    ELSE
        DBMS_OUTPUT.PUT_LINE('False or NULL');
    ENDIF;
END;
```

1. Manually create a replication definition on Replication Server using this RCL command:

```
create function replication definition ra$xxx_boolproc
  with primary at myprimary.pdb
  with all functions named boolproc (
    @"a" rs_oracle_decimal
    @"b" rs_oracle_decimal )
  searchable parameters(@"a", @"b")
  send standby all parameters
```

**Note:** If the Replication Agent `pdb_auto_create_repdefs` configuration parameter is set to true, a replication definition will be created automatically.

2. Mark the stored procedure for replication:

```
pdb_setrepproc boolproc, mark
```

3. Adaptive Server Enterprise does not support BOOLEAN stored procedure arguments, so you must map the Oracle stored procedure BOOLEAN argument to an integer argument for the corresponding stored procedure at the replicate database.
4. Create a stored procedure, defined by this Transact-SQL® statement, at the replicate database:

```
create proc boolproc (@a int, @b int) as
  if @a = 1
    print 'True'
  else
    print 'False or NULL'
go
```

5. Create a subscription on Replication Server for the replication definition:

```
create subscription sub_intproc
  for ra$xxx_boolproc
  with replicate at myreplicate.rdb
go
```

The stored procedure is replicated when it is executed at the primary database:

```
EXECUTE boolproc(true,1);
```

However, the `boolproc` procedure at the replicate Adaptive Server Enterprise will be invoked with an integer value instead of a BOOLEAN argument:

```
boolproc 1, 1
go
```

## Oracle Warm Standby

Using Replication Server, you can create and maintain a warm standby environment for an Oracle database.

In standby mode, Replication Agent:

- Scans the transaction log and keeps the Replication Agent System Database (RASD) current

## Replication Agent for Oracle

- Does not send any LTL to Replication Server
- Continues to perform log truncation

When the active database fails or when you want to perform maintenance on the active database, you can switch to the standby database. For instructions on switching the active and standby database, see “Managing Heterogeneous Warm Standby for Oracle” in the *Replication Server Heterogeneous Replication Guide*.

To function in warm standby mode:

- Replication Agent must be installed on both the primary and standby side and must also be successfully initialized. The Replication Agent on the standby side should be running in standby mode with the **ra\_standby** parameter set to true.
- Replication Agent should have the **rs\_source\_ds** and **rs\_source\_db** parameters configured as physical connections to Replication Server.
- Replication Agent should enable or disable the replication of DDL statements as desired using the **pdb\_setrep\_ddl** command.
- Replication Agent should set the **pdb\_auto\_create\_repdefs**, **pdb\_dflt\_column\_repl**, **pdb\_dflt\_object\_repl**, and **pdb\_automark\_tables** parameters to true.

For details on using these commands and configuration parameters, see the *Replication Agent Reference Manual*. For detailed steps involved in creating and managing warm standby for Oracle, see the *Replication Agent Heterogeneous Replication Guide*.

## Oracle Flashback

Replication Agent can replicate Oracle Flashback operations performed at the table level and the transaction level.

- Replication Agent can replicate Flashback Table commands to restore a database back to a System Change Number (SCN), timestamp, or restore point within the threshold specified by the Oracle **UNDO\_RETENTION** parameter. If a table has been marked for replication, Replication Agent can replicate any DML changes that result from executing the contents of the **UNDO\_SQL** column of the Oracle **FLASHBACK\_TRANSACTION\_QUERY** view.
- Replication Agent replicates these Oracle DDL commands:
  - **DROP TABLE** *table* (when the Recycle Bin is enabled)
  - **FLASHBACK TABLE** *table* **TO BEFORE DROP**
  - **PURGE TABLE** *table*
  - **PURGE INDEX** *index*
  - **PURGE TABLESPACE** *tablespace*
  - **PURGE RECYCLEBIN**
  - **PURGE DBA\_RECYCLEBIN**

---

**Note:** Since the replicate database may not be an exact copy of the primary database, these DDL commands may not execute successfully at the replicate database or may have a different result at the replicate database. For example, the **PURGE DBA\_RECYCLEBIN**

command purges more objects at the replicate database if the Recycle Bin for the replicate database contains more objects than the primary database Recycle Bin.

- Replication Agent can read from the Oracle Flash Recovery Area if the **pdb\_archive\_path** configuration parameter is set to that location.

### **Requirements for Oracle Flashback**

Observe these requirements when using Oracle Flashback with Replication Agent.

To use Oracle Flashback with Replication Agent:

- The user ID for **pds\_username** must have **select** permission on **SYS.RECYCLEBIN\$**.
- For Replication Agent to replicate the Oracle **PURGE DBA\_RECYCLEBIN** command, the user ID for **ddl\_username** must have **sysdba** permission and should be suffixed with "as sysdba". For example:

```
ra_config ddl_username, "myuser as sysdba"
go
```

- Replication Agent replicates Flashback operations from Oracle 10g and 11g databases but not from earlier versions of Oracle.

### **Limitations to Oracle Flashback**

Replication of Oracle Flashback commands is subject to some limitations.

- If the replicate database does not have the Recycle Bin enabled, Flashback commands, the **PURGE** command, and any command that accesses Recycle Bin objects fails, even if a **DROP TABLE** command has been successfully replicated.
- Replication Agent does not support the translation of DDL commands between a primary and replicate database of different types. Consequently, DDL replication must be disabled when replicating from an Oracle primary database to a non-Oracle replicate database, and Flashback DDL commands cannot be replicated in this case.
- If an Oracle **FLASHBACK TABLE** command is issued with the **RENAME TO** clause, Replication Agent does not automatically update the replication definition with the new table name. You must do this manually.
- Replication Agent reconstructs Flashback commands based on the original object name, not on the object Recycle Bin name. When there are multiple versions of an object in the Recycle Bin, Replication Agent reconstructs a Flashback command to use the most recent version of the object that exists in the Recycle Bin at the replicate database. Consequently, subsequent Oracle commands that affect the Recycle Bin may result in inconsistency between the primary and replicate databases.

For example, the primary Oracle database contains these dropped versions of table TAB1:

```
SQL> SELECT object_name as recycle_name, original_name,
           FROM recyclebin;
RECYCLE_NAME                                ORIGINAL_NAME  TYPE
-----
BIN$zyxwvutsrqponmlkjihgfedcba$1  TAB1          TABLE
```

## Replication Agent for Oracle

```
BIN$zyxwvutsrqponmlkjihgfedcba$2  TAB1  TABLE
BIN$zyxwvutsrqponmlkjihgfedcba$3  TAB1  TABLE
```

The replicate Oracle database contains these dropped versions of table TAB1:

```
SQL> SELECT object_name as recycle_name, original_name, type
        FROM recyclebin;
RECYCLE_NAME                                ORIGINAL_NAME  TYPE
-----
BIN$abcdefghijklmnopqrstuvwxy$1            TAB1          TABLE
BIN$abcdefghijklmnopqrstuvwxy$2            TAB1          TABLE
BIN$abcdefghijklmnopqrstuvwxy$3            TAB1          TABLE
```

This Flashback command is executed at the primary Oracle database:

```
FLASHBACK TABLE "BIN$zyxwvutsrqponmlkjihgfedcba$2" TO BEFORE
DROP;
```

Because Replication Agent reconstructs Flashback commands based on the original object name and uses the most recent version of a dropped object in the Flashback command, this command is executed at the replicate Oracle database:

```
FLASHBACK TABLE "BIN$abcdefghijklmnopqrstuvwxy$3" TO BEFORE
DROP;
```

If BIN\$zyxwvutsrqponmlkjihgfedcba\$2 differs in content from BIN\$abcdefghijklmnopqrstuvwxy\$3, the primary and replicate databases have become inconsistent.

### **Disabling Oracle Recycle Bin**

If you do not intend to use the Recycle Bin at the replicate Oracle database, you can manually disable it.

### **Prerequisites**

Disabling the Recycle Bin requires the **sysdba** privilege.

### **Task**

Enter this command, and then restart the replicate Oracle database:

```
ALTER SYSTEM SET RECYCLEBIN=OFF SCOPE=SPFILE;
```

Or, if you are using a version of Oracle that does not have the RECYCLEBIN parameter, enter:

```
ALTER SYSTEM SET "_recyclebin"=FALSE SCOPE=BOTH;
```

---

**Note:** If you are using Oracle RAC, disable the recycle bin for each instance in the cluster.

---

### **Dropped Objects and Article Status**

If a marked table is dropped while the Recycle Bin is enabled at the primary database, the **ra\_helparticle** command still reports the status of the corresponding article as Current. **ra\_helparticle** reports the status of the corresponding article as Dropped only after the dropped table is purged from the primary database Recycle Bin.



### **Disabling Flashback Replication with Recycle Bin Disabled**

Disable the replication of Oracle Flashback DDL commands when the Recycle Bin is disabled.

If the replicate Oracle database does not have the Recycle Bin enabled, any command that accesses Recycle Bin objects at the replicate database fails. You should therefore disable the replication of Flashback DDL commands.

Disable the replication of Flashback DDL commands in one of these ways:

- Use the **pdb\_setrepddl** command to prevent the replication of Flashback DDL commands. See the *Replication Agent Reference Manual*.
- Add a “warning” error action for Flashback DDL execution failure to Replication Server so that replication can continue with the replicate Recycle Bin disabled. Run this Replication Server script against the RSSD:

```
$SYBASE/RAX-15_5/scripts/oracle/  
hds_oracle_setup_flashback_errors.sql
```

After you run this script, you must restart Replication Server.

---

**Note:** This script uses the **rs\_oracle\_error\_class** default error class as a template. If you are using a custom error class and want Replication Server to continue replicating without interruption, you must instruct Replication Server to display warning messages 38305 and 38307 in its error log:

```
assign action warn for your_error_class to 38305, 38307
```

where *your\_error\_class* is the name of your custom error class.

---

To remove the “warning” error action for Flashback DDL execution failure to Replication Server, run this Replication Server script against the RSSD:

```
$SYBASE/RAX-15_5/scripts/oracle/  
hds_oracle_remove_flashback_errors.sql
```

### **XMLTYPE Data Replication**

When an Oracle table is created with an XMLTYPE column but without any XML schema specification, a hidden CLOB column is automatically created to store the XML data. The XMLTYPE column becomes a virtual column for the hidden CLOB column. In the corresponding Oracle base table, the hidden column immediately follows the XMLTYPE column that it represents and is named **SYS\_NCnnnnn\$**, where *nnnnn* represents the position of the hidden column in the base table.

For example, a table is created in the Oracle database with this DDL command:

```
CREATE TABLE sampletable  
( col1 INT,  
  , col2 INT,  
  , xml1 XMLTYPE  
  , xml2 XMLTYPE);
```

The Oracle database creates hidden columns named `SYS_NC00004$` and `SYS_NC00006$`, which respectively correspond to the `xml1` and `xml2` columns. These hidden CLOB columns cannot be accessed directly. However, they can be viewed by querying the `col$` and `obj$` base tables of the Oracle data dictionary. See the Oracle documentation.

### See also

- *Oracle XMLTYPE Datatype Compatibility* on page 28
- *Oracle 10g and 11g XMLTYPE Restrictions* on page 30

### **Example: Replicating XMLTYPE Column Data from Oracle to Oracle**

Replicate XMLTYPE column data from an Oracle primary database to an Oracle replicate database.

To replicate a table defined by this DDL statement:

```
CREATE TABLE sampletable
( col1 INT,
  col2 INT,
  xml1 XMLTYPE,
  xml2 XMLTYPE);
```

1. Manually create a replication definition on Replication Server this RCL command:

```
create replication definition ra$xxx_sampletable
with primary at myprimary.pdb
with all tables named sampletable (
  col1 int,
  col2 int,
  xml1 as SYS_NC00004$ text,
  xml2 as SYS_NC00006$ text )
primary key (col1, col2)
go
```

---

**Note:** If the Replication Agent `pdb_auto_create_repdefs` configuration parameter is set to true, a replication definition is created automatically.

---

2. Mark the table for replication:

```
pdb_setreptable sampletable, mark
```

3. Because of the hidden CLOB columns, you must enable replication for the table using **pdb\_setrepcol**:

```
pdb_setrepcol sampletable, enable
```

4. Create a corresponding table at the replicate Oracle database:

```
CREATE TABLE sampletable
( col1 INT,
  col2 INT,
  xml1 XMLTYPE,
  xml2 XMLTYPE);
```

---

**Note:** If DDL replication has been enabled, manually create the replicate table.

---

### **Example: Replicating XMLTYPE Column Data from Oracle to Adaptive Server Enterprise**

Replicate XMLTYPE column data from an Oracle primary database to an Adaptive Server Enterprise replicate database.

To replicate a table defined by this DDL statement:

```
CREATE TABLE sampletable
( col1 INT,
  col2 INT,
  xml1 XMLTYPE,
  xml2 XMLTYPE);
```

1. Manually create a replication definition on Replication Server using this RCL command:

```
create replication definition ra$xxx_sampletable
with primary at myprimary.pdb
with all tables named sampletable (
col1 int,
col2 int,
xml1 as SYS_NC00004$ text,
xml2 as SYS_NC00006$ text )
primary key (col1, col2)
go
```

---

**Note:** If the Replication Agent **pdb\_auto\_create\_repdefs** configuration parameter is set to true, a replication definition is created automatically.

---

2. Mark the table for replication:

```
pdb_setreptable sampletable, mark
```

3. Because of the hidden CLOB columns, you must enable replication for the table using **pdb\_setrepcol**:

```
pdb_setrepcol sampletable, enable
```

4. Create a corresponding table at the replicate Adaptive Server Enterprise database using the hidden column names from the primary database table:

```
create table sampletable
( col1 int,
  col2 int,
  SYS_NC00004$ text,
  SYS_NC00006$ text)
go
```

---

**Note:** The XMLTYPE columns from the Oracle primary database table map to text columns in the Adaptive Server Enterprise replicate database table.

---

### **Example: Replicating an XMLTYPE Table from Oracle to Oracle**

Replicate an XMLTYPE table from an Oracle primary database to an Oracle replicate database.

This statement creates a simple XMLTYPE table with one implicit CLOB column that can be accessed through a default pseudocolumn named XMLDATA:

```
CREATE TABLE sampletable OF XMLTYPE;
```

To replicate a table defined by this DDL statement:

1. Manually create a replication definition on Replication Server using this RCL command and the hidden column SYS\_NC\_OID\$, which contains the object ID for sampletable:

```
create replication definition ra$xxx_sampletable
with primary at myprimary.pdb
with all tables named sampletable (
SYS_NC_OID$ rs_oracle_binary,
XMLDATA text )
primary key (SYS_NC_OID$)
go
```

The Replication Server name for the Oracle RAW datatype is rs\_oracle\_binary.

---

**Note:** If the Replication Agent **pdb\_auto\_create\_repdefs** configuration parameter is set to true, a replication definition is created automatically.

---

2. Mark the table for replication:

```
pdb_setreptable sampletable, mark
```

3. Create a corresponding table at the replicate Oracle database:

```
CREATE TABLE sampletable OF XMLTYPE;
```

## **Oracle 11g Release 2**

All of the functionality that Replication Agent supports for Oracle Database 11g Release 1 is also supported by Replication Agent for Oracle Database 11g Release 2. Replication Agent also supports some functionality introduced by Oracle Database 11g Release 2.

---

**Note:** If you want to replicate tables compressed for online transaction processing (OLTP), apply first the Oracle patch# 129050503.

---

Replication Agent supports functionality common to both Oracle Database 11g Release 1 and Oracle Database 11g Release 2, including:

- Oracle DDL and DML replication in systems with and without Automatic Storage Management (ASM) and Real Application Clusters (RAC).
- Use of the Oracle Recovery Manager (RMAN) utility to truncate old archive log files.
- Use of the Oracle Recycle Bin and replication of Oracle Flashback operations.

- Oracle Data Guard.

Replication Agent also supports some features that are new as of Oracle Database 11g Release 2:

- Use of the **FORCE** option with a **CREATE OR REPLACE TYPE** statement on types with type dependencies.
- DDL statements on tables enabled for Flashback Data Archive.
- Version 11.2 time zone files and new time zone behavior.

### **Network Configuration File Location and Structure**

The `tnsnames.ora` file is located in `ORACLE_HOME\network\admin`. In an Oracle Database 11g Release 2 instance running ASM or RAC, the `tnsnames.ora` file is read by default from the grid infrastructure home directory at `Grid_home\network\admin`.

If you are using an Oracle Database 11g Release 2 instance running ASM or RAC, set the Replication Agent **asm\_tns\_filename** parameter to `Grid_home\network\admin\tnsnames.ora`.

By default, the `tnsnames.ora` file at `Grid_home\network\admin` contains an incomplete ASM entry that lacks information in the `DESCRIPTION` and `SERVICE_NAME` fields. If you are using an Oracle Database 11g Release 2 instance running ASM, set the Replication Agent **asm\_tns\_connection** parameter to the ASM connection name specified in this incomplete ASM entry. Replication Agent completes the `DESCRIPTION` and `SERVICE_NAME` fields, and you can then use the `tnsnames.ora` file in `Grid_home\network\admin` to connect to the ASM instance server.

### **Time Zone File**

By default, Oracle Database 11g Release 2 uses the large time zone file, `timezone_11.dat`. This file contains all the time zones defined in the database.

If you are using Oracle Database 11g Release 2, set the Replication Agent **pdb\_timezone\_file** parameter to the location of the `timezone_11.dat` file:

```
ra_config pdb_timezone_file, $ORACLE_HOME/oracore/zoneinfo/
timezone_11.dat
```

### **User-Defined Type Dependencies**

You can use the **CREATE OR REPLACE TYPE** command to change the definition for an existing user-defined type. However, this command throws an error if the referenced type has table or type dependencies.

Oracle Database 11g Release 2 allows you to use **FORCE** with the **CREATE OR REPLACE TYPE** command to replace a type that has a type dependency:

```
CREATE TYPE mytype1 AS OBJECT (a number) NOT FINAL;
CREATE TYPE mytype2 UNDER mytype1 (b varchar(10));
CREATE OR REPLACE TYPE mytype1 FORCE AS OBJECT (c varchar(20));
```

## Replication Agent for Oracle

Oracle Database 11g Release 2 does not allow you to use **FORCE** with the **CREATE OR REPLACE TYPE** command to replace a type that has a table dependency:

```
CREATE TABLE mytable1 (cola mytype1);  
CREATE OR REPLACE TYPE mytype1 FORCE AS OBJECT (d number);
```

The last command results in an error because `mytype1` has a table dependency on `mytable1`:

```
ERROR at line 1:  
ORA-22866: cannot replace a type with table dependents
```

Replication Agent supports use of the **FORCE** option with the **CREATE OR REPLACE TYPE** command in Oracle Database 11g Release 2 to replace types with type dependencies but not for types with table dependencies.

### **Flashback Data Archive Support for DDL Commands**

Oracle Flashback allows database administrators and users to view past states of database objects and to restore database objects to a previous state without using point-in-time recovery. Users can query past data, query metadata to build a detailed history of changes, recover data to a previous point in time, and roll back transactions while the database is online.

Replication Agent supports the replication of DDL commands on tables that are being tracked with the Flashback Data Archive in Oracle Database 11g Release 2. These DDL commands include:

- **Add, Drop, Rename, Modify Column**
- **Drop, Truncate Partition**
- **Rename, Truncate Table**
- **Add, Drop, Rename, Modify Constraint**

## **Oracle 9i**

Replication Agent supports Oracle 9i when Oracle 10g or 11g is running in 9i compatibility mode.

### **Limitations to Oracle 9i**

Some Replication Agent features cannot be used with Oracle 9i.

These Replication Agent features cannot be used with Oracle 9i:

- Oracle index-organized tables and the `ANYDATA` datatype
- Autocorrection
- Replication of `XMLTYPE` data
- Flashback
- RAC
- ASM

## **Replication Agent Objects in the Oracle Primary Database**

Replication Agent creates objects in the primary database to assist with replication tasks.

The Replication Agent objects are created by invoking the **ra\_admin** command with the **init** keyword. When you invoke this command, Replication Agent generates a SQL script that contains the SQL statements for the objects created or modified in the primary database. This script is stored in the `partinit.sql` file in the `RAX-15_5\inst_name\scripts\alog` directory. You must create these objects before marking any primary database objects for replication.

---

**Note:** The generated scripts are for informational purposes only. You cannot run them manually to initialize the primary database or Replication Agent. This is also true for the procedure marking and unmarking scripts that are generated when you use **pdb\_setrepproc**. Scripts are no longer generated when marking and unmarking tables with **pdb\_setreptable**.

See the *Replication Agent Administration Guide*.

### **Replication Agent Object Names**

Replication Agent creates objects in the primary database to assist with replication tasks.

There are two variables in Replication Agent database object names:

- **prefix** – represents the one-to-three-character string value of the **ra\_admin\_prefix** parameter. The default is **ra\_**.
- **xxx** – represents an alphanumeric counter, a string of characters that is (or may be) added to a database object name to make that name unique in the database.

The value of the **ra\_admin\_instance\_prefix** parameter is the prefix string used in all Replication Agent object names.

The value of the **ra\_admin\_prefix\_chars** parameter is a list of the nonalphanumeric characters allowed in the prefix string specified by **ra\_admin\_instance\_prefix**. This list of allowed characters is database-specific. For example, in Oracle, the only nonalphanumeric characters allowed in a database object name are the \$, #, and \_ characters.

Use the **ra\_admin** command to view the names of Replication Agent transaction log components in the primary database.

See the *Replication Agent Administration Guide* for details on setting up object names.

### **Finding the Names of the Objects Created**

Find the names of Replication Agent objects created in the primary Oracle database.

Use the **ra\_admin** command to return a list of all the Replication Agent objects created in the primary database.

## **Table Objects**

Replication Agent creates table objects in the Oracle primary database.

These tables are considered Replication Agent objects.

**Table 4. Replication Agent Tables**

<b>Table</b>	<b>Database Name</b>
Procedure-active table	<prefix>PROACTIVE
Multiple Replication Agents instances table	<prefix>AGENT
Multiple Replication Agents marked-tables table	<prefix>TABLE
Multiple Replication Agents marked-procedures table	<prefix>PROCEDURE

### **Procedure-Active Table**

Maintains information about the actively running procedures during replication.

**Table 5. Procedure-Active Table**

<b>Column Name</b>	<b>Type</b>	<b>Contents</b>
callseq	NUMBER	The procedure call sequence.
sid	NUMBER	The ID of the session the procedure is executed on.
owner	VARCHAR2(256)	The procedure owner.
sproc	VARCHAR2(256)	The procedure name.
spid	NUMBER	The procedure object ID.
shadow	VARCHAR2(256)	The procedure shadow table name.
shid	NUMBER	The procedure shadow table object ID.
objtype	NUMBER	The procedure object type.
sqlerrm	VARCHAR2(256)	The text of the SQL error message for any error that occurred during procedure execution.

### **Multiple Replication Agents Instances Table**

Maintains information about each Replication Agent instance in a Replication Agent group.

The multiple Replication Agents instances table is named *ra\_admin\_prefix*AGENT, where *ra\_admin\_prefix* is the prefix string used for Replication Agent system object names.



**Table 6. Multiple Replication Agents Instances Table**

Column Name	Type	Contents
prefix	VARCHAR2(3)	The instance prefix for this Replication Agent instance. This prefix is determined by the setting of the <b>ra_admin_instance_prefix</b> parameter.
locator	VARCHAR2(128)	The truncation point for this Replication Agent instance.
inittime	TIMESTAMP	The time at which this Replication Agent instance was initialized.
version	VARCHAR2(8)	The version of this Replication Agent instance.
rasd_export_file	VARCHAR(255)	The file to which RASD data is exported.
upg_from_ver	VARCHAR(16)	The version from which Replication Agent has been upgraded.
upg_fin_ts	TIMESTAMP	The time at which the upgrade finished.

**Multiple Replication Agents Marked-Tables Table**

Maintains information about tables marked for replication in a Replication Agent group.

The multiple Replication Agents marked-tables table is named *ra\_admin\_prefix*TABLE, where *ra\_admin\_prefix* is the prefix string used for Replication Agent system object names.

**Table 7. Multiple Replication Agents Marked-Tables Table**

Column Name	Type	Contents
prefix	VARCHAR2(3)	The instance prefix for this Replication Agent instance. This prefix is determined by the setting of the <b>ra_admin_instance_prefix</b> parameter.
tableid	NUMBER	The object ID for this marked table.
autocorrection	NUMBER(38)	Indicates whether or not autocorrection is enabled for the marked table. If autocorrection is enabled, this column has a value of 1. Otherwise, the column has a value of 0 (default).
ruletype	NUMBER(38)	Not used.
rulevalue	VARCHAR2(30)	Not used.

**Multiple Replication Agents Marked-Procedures Table**

Maintains information about procedures marked for replication in a Replication Agent group.

The multiple Replication Agents marked-procedures table is named *ra\_admin\_prefix*PROCEDURE, where *ra\_admin\_prefix* is the prefix string used for Replication Agent system object names.

**Table 8. Multiple Replication Agents Marked-Procedures Table**

Column Name	Type	Contents
prefix	VARCHAR2(3)	The instance prefix for this Replication Agent instance. This prefix is determined by the setting of the <i>ra_admin_instance_prefix</i> parameter.
procid	NUMBER	The object ID for this marked procedure.
shadowtable	VARCHAR2(30)	The object ID of the shadow table corresponding to this procedure.
ruletype	NUMBER(38)	Not used.
rulevalue	VARCHAR2(30)	Not used.

**Marker Objects**

Replication Agent creates marker objects in the primary database.

These Replication Agent objects are related to Replication Server markers. No permissions are granted when these objects are created.

**Table 9. Replication Agent Marker Objects**

Object	Name
Transaction log marker procedure	RS_MARKER
Dump marker procedure	RS_DUMP
Transaction log marker shadow table	<prefix>MARKERSH[xxx]
Dump marker shadow table	<prefix>DUMPSH[xxx]

**Sequences**

Replication Agent creates sequences in the Oracle primary database.

These Oracle sequences are considered Replication Agent objects.

**Table 10. Replication Agent Sequences**

Sequence	Database Name
Assign procedure call	<prefix>PCALL_[xxx]

## **Marked Procedures**

Replication Agent creates an object for each primary procedure marked for replication in the Oracle primary database.

A shadow table is generated for every procedure that is marked for replication. A shadow table is a work table that is used during replication in conjunction with the procedure-active table to replicate procedure execution.

**Table 11. Replication Agent Objects for Each Marked Procedure**

Object	Name
Shadow table	<prefix><procedure_name>SH

### **Shadow Table**

Maintains information about the procedure call sequence and also maintains a column for every procedure argument where the column definition matches the definition of the argument.

**Table 12. Shadow Table**

Column Name	Type	Contents
callseq	NUMBER	The procedure call sequence.

## **Transaction Log Truncation**

Replication Agent supports both automatic and manual log truncation.

Replication Agent provides two options for automatic transaction log truncation:

- Periodic truncation, based on a time interval you specify
- Automatic truncation whenever Replication Agent receives a new LTM locator value from the primary Replication Server. You also have the option to switch off automatic log truncation. By default, automatic log truncation is enabled and is set to truncate the log whenever Replication Agent receives a new LTM locator value from the primary Replication Server.

To configure Replication Agent log truncation, observe these guidelines:

- When **pdb\_include\_archives** is set to true, the default, and **pdb\_archive\_remove** is set false, the Replication Agent does not perform any online or archived transaction log truncation. When **pdb\_include\_archives** is set to true, the default, and

**pdb\_archive\_remove** is set to true, Replication Agent deletes from the **pdb\_archive\_path** location the archive redo logs that have already been processed. The Replication Agent is not responsible for archiving online transaction logs.

---

**Note:** Sybase recommends that you configure the Replication Agent to remove archive log files only if an additional archive log directory is used.

---

- When the configuration parameter **pdb\_include\_archives** is set to false, Replication Agent performs online redo log truncation (either scheduled or manual) by issuing the **alter system** command with the archive log sequence keywords. The command uses the log sequence number of the redo log file whose contents have been processed by the Replication Agent and are ready to be archived.

---

**Note:** The **alter system** command syntax in Oracle allows redo log files to be archived in addition to the single log sequence specified in the command. To avoid unintentional archiving, Replication Agent issues this command only when it is processing the redo log file whose status is current.

---

- You can specify the automatic truncation option you want (including none) by using **ra\_config** to set the value of the **truncation\_type** configuration parameter. To truncate the transaction log automatically based on a time interval, use **ra\_config** to set the value of the **truncation\_interval** configuration parameter.
- You can truncate the Replication Agent transaction log manually, at any time, by invoking **pdb\_truncate\_xlog** at the Replication Agent administration port.

For more information on these properties, see the *Replication Agent Reference Manual*. For a more detailed description of truncating, see "Administering Replication Agent" in the *Replication Agent Administration Guide*.

# Upgrading and Downgrading Replication Agent

Review the procedures for upgrading and downgrading Replication Agent.

---

**Warning!** Replication Agent cannot be downgraded after you use any new features. Before you use any new features, test the new version of Replication Agent at the same functional level as your previous version of Replication Agent. Only when you are satisfied with the new version of Replication Agent at this existing functional level should you finalize the upgrade by using new features. You should also make a backup copy of an instance of the Replication Agent version to which you have upgraded before attempting to use any new feature.

---

## Upgrade and Migration Procedures for Replication Agent for Oracle

---

Replication Agent for Oracle 15.7.1 ESD #2 must be installed on the same host on which the primary Oracle server is running.

Using any of the upgrade procedures described here, the new Replication Agent for Oracle 15.7.1 ESD #2 instances will have the same configuration as previously existing instances, including instance names, administrative user IDs and passwords, and administrative port numbers.

Replication Agent for Oracle 15.7.1 ESD #2 does not have to be installed on the same machine as the primary Oracle data server. However:

- Replication Agent for Oracle must be installed on a host that has access to Oracle LogMiner.
- If Replication Agent for Oracle is configured to automatically truncate Oracle logs, it must be installed on a machine that has direct access to the Oracle logs.

---

**Note:** For upgrades within a common release level, as in the case of an ESD applied to a particular version of Replication Agent, use the **rao\_admin -u** option applied to a particular instance of Replication Agent or to all instances of Replication Agent. See the *Replication Agent Administration Guide*.

---

You can use the **ra\_finalize\_upgrade** to manually force upgrade finalization of an instance from a previous version and prevent downgrade to the previous version. You must finalize the upgrade to enable any new functionality. See the *Replication Agent Reference Manual* for details on using the **ra\_finalize\_upgrade** command.

### **Upgrading Replication Agent for Oracle to 15.7.1 ESD #2**

Upgrade Replication Agent for Oracle to version 15.7.1 ESD #2.

1. Create the 15.7.1 ESD #2 version of all valid existing Replication Agent instances.

---

**Note:** This step creates new Replication Agent 15.7.1 ESD #2 instances for all valid existing Oracle instances of the earlier version of Replication Agent. If you do not want to run a newly created instance on this host, simply delete the new instance directory.

---

- a) On UNIX, set the SYBASE environment variables by changing to the SYBASE directory in which Replication Agent 15.7.1 ESD #2 is installed and sourcing the SYBASE script:

- For C Shell: `source SYBASE.csh`
- For Bourne or Korn shell: `. SYBASE.sh`

- b) Change to the Replication Agent `bin` directory:

- On UNIX:

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

- On Windows:

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

---

**Note:** The Replication Agent directory name for release 15.5 and later is `RAX-15_5`.

---

- c) Create new versions of all valid existing instances:

```
rao_admin init
```

2. Run the `$SYBASE/RAX-15_5/<instance>/RUN_<instance>` script and log in to the Replication Agent 15.7.1 ESD #2 instance. Migrate the Replication Agent metadata by issuing the `ra_migrate` command.

### **Migrating Replication Agent for Oracle 15.7.1 ESD #2 When Upgrading Oracle 10g to 11g**

Migrate Replication Agent 15.7.1 ESD #2 when you are also upgrading Oracle 10g to 11g.

Replication Agent for Oracle migration to support upgrading Oracle 10g to Oracle 11g is similar to upgrading Replication Agent for Oracle 15.1 or 15.2 to Replication Agent for Oracle 15.7.1 ESD #2.

---

**Note:** Quiesce the Replication Agent before upgrading Oracle 10g to Oracle 11g. The replication environment must have completed processing of all transactions before upgrading Oracle because the Replication Agent moves the truncation point to the end of the log during Replication Agent migration.

---

1. Follow the steps that Oracle provides in their documentation for upgrading from Oracle 10g to Oracle 11g.
2. After upgrading Oracle, restart the Replication Agent, and issue the **ra\_migrate** command.
3. As with the log-based Replication Agent upgrade process, you may need to reconfigure the Replication Agent for Oracle instance to read archive logs depending on the configuration in Oracle. This may change following the Oracle upgrade.

If you are upgrading from log-based Replication Agent and upgrading Oracle 10g to Oracle 11g at the same time, migrate Replication Agent 15.7.1 ESD #2 only once.

## Downgrading Replication Agent for Oracle

You can downgrade Replication Agent for Oracle from version 15.7.1 ESD #2 to version 15.7.1 or later.

You may need to downgrade Replication Agent if the upgrade process fails or if replication fails after an upgrade. Replication may fail when new features fail to function as expected or if there are changes to:

- DDL and how it is handled by Replication Agent
- The content or structure of the Replication Agent System Database (RASD)

If you are using any new features from Replication Agent 15.7.1 ESD #2, you cannot downgrade. For a list of the new features for Replication Agent 15.7.1 ESD #2, see the *Replication Agent New Features*.

Replication Agent 15.7.1 ESD #2 must be installed on the same platform on which the primary Oracle server is running.

1. Change to the Replication Agent 15.7.1 ESD #2 `bin` directory:

- On UNIX:

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

- On Windows:

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

The Replication Agent directory name for release 15.5 and later is `RAX-15_5`.

2. Run the **ra\_downgrade** command at the Replication Agent instance from which you are downgrading (the current version):

```
ra_downgrade
```

The **ra\_downgrade** command extracts the contents of the Replication Agent System Database (RASD) to a file named *timestamp.export*, where *timestamp* is a timestamp taken at the moment **ra\_downgrade** was invoked. This file is located in the `import` subdirectory under the directory specified by the **rasd\_backup\_dir** configuration

## Upgrading and Downgrading Replication Agent

parameter of the Replication Agent instance to which you are downgrading (the earlier version). The absolute path to this file is returned if **ra\_downgrade** executes successfully.

---

**Note:** The **ra\_downgrade\_prepare** and **ra\_downgrade\_accept** commands are deprecated as of Replication Agent 15.7.1.

---

3. Complete the downgrade by running the **ra\_migrate** command at the Replication Agent instance to which you are downgrading (the earlier version):

```
ra_migrate timestamp.export
```

where *timestamp.export* is the file to which the **ra\_downgrade** command extracted RASD contents.

If the **ra\_migrate** command executes successfully, Replication Agent shuts down.

4. Start the Replication Agent instance to which you have downgraded (the earlier version), and resume replication:

```
resume purge
```

The **purge** keyword is needed here to purge data from the Replication Server inbound queue for the connection to which this Replication Agent is connected. Purging prevents any duplicate records from being created in Replication Server as a result of the change in OQID formats between the earlier and later versions of Replication Agent.



# Obtaining Help and Additional Information

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

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

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

---

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

---

## Technical Support

---

Get support for Sybase products.

If your organization has purchased a support contract for this product, then one or more of your colleagues is designated as an authorized support contact. If you have any questions, or if you need assistance during the installation process, ask a designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

## Downloading Sybase EBFs and Maintenance Reports

---

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

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

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

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

## Sybase Product and Component Certifications

Certification reports verify Sybase product performance on a particular platform.

To find the latest information about certifications:

- For partner product certifications, go to [http://www.sybase.com/detail\\_list?id=9784](http://www.sybase.com/detail_list?id=9784)
- For platform certifications, go to <http://certification.sybase.com/ucr/search.do>

## Creating a MySybase Profile

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

1. Go to <http://www.sybase.com/mysybase>.
2. Click **Register Now**.

## Accessibility Features

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

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

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

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

---

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

---

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

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



# Glossary

This glossary describes Replication Server Options terms.

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

See also *DBMS* and *RDBMS*.

- **atomic materialization** – a materialization method that copies subscription data from a primary database to a replicate database in a single, atomic operation. No changes to primary data are allowed until the subscription data is captured at the primary database. See also *bulk materialization* and *nonatomic materialization*.
- **BCP utility** – a bulk copy transfer utility that provides the ability to load multiple rows of data into a table in a target database. See also *bulk copy*.
- **bulk copy** – an Open Client™ interface for the high-speed transfer of data between a database table and program variables. Bulk copying provides an alternative to using SQL **insert** and **select** commands to transfer data.
- **bulk materialization** – a materialization method whereby subscription data in a replicate database is initialized outside of the replication system. You can use bulk materialization for subscriptions to table replication definitions or function replication definitions. See also *atomic materialization* and *nonatomic materialization*.
- **client** – in client/server systems, the part of the system that sends requests to servers and processes the results of those requests. See also *client application*.
- **client application** – software that is responsible for the user interface, including menus, data entry screens, and report formats. See also *client*.
- **commit** – an instruction to the DBMS to make permanent the changes requested in a transaction. See also *transaction*. Contrast with *rollback*.
- **database** – a collection of data with a specific structure (or schema) for accepting, storing, and providing data for users. See also *data server*, *DBMS*, and *RDBMS*.
- **database connection** – a connection that allows Replication Server to manage the database and distribute transactions to the database. Each database in a replication system

can have only one database connection in Replication Server. See also *Replication Server* and *route*.

- **data client** – a client application that provides access to data by connecting to a data server. See also *client*, *client application*, and *data server*.
- **data distribution** – a method of locating (or placing) discrete parts of a single set of data in multiple systems or at multiple sites. Data distribution is distinct from data replication, although a data replication system can be used to implement or support data distribution. Contrast with *data replication*.
- **data replication** – the process of copying primary data to remote locations and synchronizing the copied data with the primary data. Data replication is different from data distribution. Replicated data is a stored copy of data at one or more remote sites throughout a system, and it is not necessarily distributed data. Contrast with *data distribution*. See also *transaction replication*.
- **data server** – a server that provides the functionality necessary to maintain the physical representation of a table in a database. Data servers are usually database servers, but they can also be any data repository with the interface and functionality a data client requires. See also *client*, *client application*, and *data client*.
- **datatype** – a keyword that identifies the characteristics of stored information on a computer. Some common datatypes are: *char*, *int*, *smallint*, *date*, *time*, *numeric*, and *float*. Different data servers support different datatypes.
- **DBMS** – an abbreviation for database management system, a computer-based system for defining, creating, manipulating, controlling, managing, and using databases. The DBMS can include the user interface for using the database, or it can be a standalone data server system. Compare with *RDBMS*.
- **ERSSD** – an abbreviation for Embedded Replication Server System Database, which manages replication system information for a Replication Server. See also *Replication Server*.
- **failback** – a procedure that restores the normal user and client access to a primary database, after a failover procedure switches access from the primary database to a replicate database. See also *failover*.
- **failover** – a procedure that switches user and client access from a primary database to a replicate database, particularly in the event of a failure that interrupts operations at the primary database, or access to the primary database. Failover is an important fault-tolerance feature for systems that require high availability. See also *failback*.
- **function** – a data server object that represents an operation or set of operations. Replication Server distributes operations to replicate databases as functions. See also *stored procedure*.
- **function string** – a string that Replication Server uses to map a function and its parameters to a data server API. Function strings allow Replication Server to support heterogeneous replication, in which the primary and replicate databases are different types, with different SQL extensions and different command features. See also *function*.

- **gateway** – connectivity software that allows two or more computer systems with different network architectures to communicate.
- **inbound queue** – a stable queue managed by Replication Server to spool messages received from a Replication Agent. See also *outbound queue* and *stable queue*.
- **interfaces file** – a file containing information that Sybase Open Client and Open Server™ applications need to establish connections to other Open Client and Open Server applications. See also *Open Client* and *Open Server*.
- **isql** – an Interactive SQL client application that can connect and communicate with any Sybase Open Server application, including Adaptive Server, Replication Agent, and Replication Server. See also *Open Client* and *Open Server*.
- **Java** – an object-oriented programming language developed by Sun Microsystems. A platform-independent, “write once, run anywhere” programming language.
- **Java VM** – the Java Virtual Machine. The Java VM (or JVM) is the part of the Java Runtime Environment (JRE) that is responsible for interpreting Java byte codes. See also *Java* and *JRE*.
- **JDBC** – an abbreviation for Java Database Connectivity. JDBC is the standard communication protocol for connectivity between Java clients and data servers. See also *data server* and *Java*.
- **JRE** – an abbreviation for Java Runtime Environment. The JRE consists of the Java Virtual Machine (Java VM or JVM), the Java Core Classes, and supporting files. The JRE must be installed on a machine to run Java applications, such as Replication Agent. See also *Java VM*.
- **LAN** – an abbreviation for “local area network,” a computer network located on the user premises and covering a limited geographical area (usually a single site). Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary can be subject to some form of regulation. Contrast with *WAN*.
- **latency** – in transaction replication, the time it takes to replicate a transaction from a primary database to a replicate database. Specifically, latency is the time elapsed between committing an original transaction in the primary database and committing the replicated transaction in the replicate database.

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

See also *transaction replication*.

- **LOB** – an abbreviation for large object, a large collection of data stored as a single entity in a database.
- **Log Reader** – an internal component of Replication Agent that interacts with the primary database to capture transactions for replication. See also *Log Transfer Interface* and *Log Transfer Manager*.

- **Log Transfer Interface** – an internal component of Replication Agent that interacts with Replication Server to forward transactions for distribution to Replication Server. See also *Log Reader* and *Log Transfer Manager*.
- **Log Transfer Language** – the proprietary protocol used between Replication Agent and Replication Server to replicate data from the primary database to Replication Server. See also *Log Reader* and *Log Transfer Interface*.
- **Log Transfer Manager** – an internal component of Replication Agent that interacts with the other Replication Agent internal components to control and coordinate Replication Agent operations. See also *Log Reader* and *Log Transfer Interface*.
- **maintenance user** – a special user login name in the replicate database that Replication Server uses to apply replicated transactions to the database. See also *replicate database* and *Replication Server*.
- **materialization** – the process of copying the data from a primary database to a replicate database, initializing the replicate database so that the replication system can begin replicating transactions. See also *atomic materialization*, *bulk materialization*, and *nonatomic materialization*.
- **Multi-Path Replication™** – Replication Server feature that improves performance by enabling parallel paths of data from the source database to the target database. These multiple paths process data independently and are applicable when sets of data can be processed in parallel without transactional consistency requirements between them.
- **nonatomic materialization** – a materialization method that copies subscription data without a lock on the primary database. Changes to primary data are allowed during data transfer, which may cause temporary inconsistencies between the primary and replicate databases. Contrast with *atomic materialization*. See also *bulk materialization*.
- **ODBC** – an abbreviation for Open Database Connectivity, an industry-standard communication protocol for clients connecting to data servers. See also *client*, *data server*, and *JDBC*.
- **Open Client** – a Sybase product that provides customer applications, third-party products, and other Sybase products with the interfaces needed to communicate with Open Server applications. See also *Open Server*.
- **Open Client application** – An application that uses Sybase Open Client libraries to implement Open Client communication protocols. See also *Open Client* and *Open Server*.
- **Open Server** – a Sybase product that provides the tools and interfaces required to create a custom server. See also *Open Client*.
- **Open Server application** – a server application that uses Sybase Open Server libraries to implement Open Server communication protocols. See also *Open Client* and *Open Server*.
- **outbound queue** – a stable queue managed by Replication Server to spool messages to a replicate database. See also *inbound queue*, *replicate database*, and *stable queue*.
- **primary data** – the data source used for replication. Primary data is stored and managed by the primary database. See also *primary database*.



- **primary database** – the database that contains the data to be replicated to another database (the replicate database) through a replication system. The primary database is the source of replicated data in a replication system. Sometimes called the active database. Contrast with *replicate database*. See also *primary data*.
- **primary key** – a column or set of columns that uniquely identifies each row in a table.
- **primary site** – the location or facility at which primary data servers and primary databases are deployed to support normal business operations. Sometimes called the active site or main site. See also *primary database* and *replicate site*.
- **primary table** – a table used as a source for replication. Primary tables are defined in the primary database schema. See also *primary data* and *primary database*.
- **primary transaction** – a transaction that is committed in the primary database and recorded in the primary database transaction log. See also *primary database*, *replicated transaction*, and *transaction log*.
- **quiesce** – to cause a system to go into a state in which further data changes are not allowed. See also *quiescent*.
- **quiescent** – in a replication system, a state in which all updates have been propagated to their destinations. Some Replication Agent and Replication Server commands require that you first quiesce the replication system.

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

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

- **RASD** – an abbreviation for Replication Agent System Database. Information in the RASD is used by the primary database to recognize database structure or schema objects in the transaction log.
- **RCL** – an abbreviation for Replication Command Language, the command language used to manage Replication Server. See also *Replication Server*.
- **RDBMS** – an abbreviation for relational database management system, an application that manages and controls relational databases. Compare with *DBMS*. See also *relational database*.
- **relational database** – a collection of data in which data is viewed as being stored in tables, which consist of columns (data items) and rows (units of information). Relational databases can be accessed by SQL requests. Compare with *database*. See also *SQL*.
- **replicate data** – A set of data that is replicated from a primary database to a replicate database by a replication system. See also *primary database*, *replication system*, and *replicate database*.
- **replicate database** – a database that contains data replicated from another database (the primary database) through a replication system. The replicate database is the database that receives replicated data in a replication system. Contrast with *primary database*. See also *replicate data*, *replicated transaction*, and *replication system*.

- **replicated transaction** – a primary transaction that is replicated from a primary database to a replicate database by a transaction replication system. See also *primary database*, *primary transaction*, *replicate database*, and *transaction replication*.
- **replicate site** – the location or facility at which replicate data servers and replicate databases are deployed to support normal business operations during scheduled downtime at the primary site. Contrast with *primary site*. See also *replicate database*.
- **Replication Agent** – an application that reads a primary database transaction log to acquire information about data-changing transactions in the primary database, processes the log information, and then sends it to a Replication Server for distribution to a replicate database. See also *primary database* and *Replication Server*.
- **replication definition** – a description of a table or stored procedure in a primary database, for which subscriptions can be created. The replication definition, maintained by Replication Server, includes information about the columns to be replicated and the location of the primary table or stored procedure. See also *Replication Server* and *subscription*.
- **Replication Server** – a Sybase software product that provides the infrastructure for a transaction replication system. See also *Replication Agent*.
- **replication system** – a data processing system that replicates data from one location to another. Data can be replicated between separate systems at a single site, or from one or more local systems to one or more remote systems. See also *transaction replication*.
- **rollback** – an instruction to a database to back out of the changes requested in a unit of work (called a transaction). Contrast with *commit*. See also *transaction*.
- **route** – A one-way message stream from a primary Replication Server to a replicate Replication Server. Routes carry data-changing commands (including those for RSSDs) and replicated functions (database procedures) between separate Replication Servers. See also *Replication Server*.
- **RSSD** – an abbreviation for Replication Server System Database, which manages replication system information for a Replication Server. See also *Replication Server*.
- **SQL** – an abbreviation for Structured Query Language, a nonprocedural programming language used to process data in a relational database. ANSI SQL is an industry standard. See also *transaction*.
- **stable queue** – a disk device-based, store-and-forward queue managed by Replication Server. Messages written into the stable queue remain there until they can be delivered to the appropriate process or replicate database. Replication Server provides a stable queue for both incoming messages (the inbound queue) and outgoing messages (the outbound queue). See also *database connection*, *Replication Server*, and *route*.
- **stored procedure** – a data server object that represents an operation or set of operations. This term is often used interchangeably with *function*.
- **subscription** – a request for Replication Server to maintain a replicated copy of a table, or a set of rows from a table, in a replicate database at a specified location. See also *replicate database*, *replication definition*, and *Replication Server*.

- **table** – in a relational DBMS, a two-dimensional array of data or a named data object that contains a specific number of unordered rows composed of a group of columns that are specific for the table. See also *database*.
- **transaction** – a unit of work in a database that can include zero, one, or many operations (including **insert**, **update**, and **delete** operations), and that is either applied or rejected as a whole. Each SQL statement that modifies data can be treated as a separate transaction, if the database is so configured. See also *SQL*.
- **transactional consistency** – A condition in which all transactions in the primary database are applied in the replicate database, and in the same order that they were applied in the primary database.
- **transaction log** – generally, the log of transactions that affect the data managed by a data server. Replication Agent reads the transaction log to identify and acquire the transactions to be replicated from the primary database. See also *Replication Agent*, *primary database*, and *Replication Server*.
- **transaction replication** – a data replication method that copies data-changing operations from a primary database to a replicate database. See also *data replication*.
- **UDB** – IBM DB2 Universal Database (formerly IBM DB2 for Linux, UNIX, and Windows).
- **WAN** – an abbreviation for “wide area network,” a system of local-area networks (LANs) connected together with data communication lines. Contrast with *LAN*.



# Index

## C

- CLASSPATH environment variable 6
- commands
  - pdb\_setreproc 36
  - pdb\_setrepseq 39
- communications
  - JDBC driver 6
- configuration parameters
  - pdb\_dflt\_object\_repl 37
- conventions
  - style 1
  - syntax 1

## D

- deferred updates 3

## J

- JDBC driver
  - Oracle 6

## L

- LTM locator
  - origin queue ID 20

## M

- marked procedures 69
- marker shadow tables 68
- marking a sequence 36
- multiple Replication Agents instances table 66
- multiple Replication Agents marked-procedures table 68
- multiple Replication Agents marked-tables table 67

## O

- Oracle database server
  - JDBC driver 6
  - origin queue ID 20

- primary database 3
  - TNS Listener Service 6
- Oracle partitioned tables 44
- origin commit time
  - Oracle 20
- origin queue ID
  - Oracle 20

## P

- partitioned tables 44
- pdb\_dflt\_object\_repl configuration parameter 37
- pdb\_setreproc command 36
- pdb\_setrepseq command 39
- primary databases
  - Oracle database server 3
- procedure-active table 66

## R

- Replication Agent
  - origin queue ID 20
  - transaction log 65
- Replication Agent for Oracle 3
  - JDBC driver 6
  - Running Oracle Server and Replication Agent on different machines 40
  - transaction log 65

## S

- sequence
  - marking 36, 37
  - unmarking 37
- sequences 68
  - marking 35
  - unmarking 35
- shadow table 69
- shadow tables
  - marker 68

## T

- TNS Listener Service, Oracle 6

## Index

transaction logs

    Replication Agent for Oracle 65

    shadow tables 68

    truncating 69

truncate partition command

    replicating 44

## U

unmarking a sequence 37