SYBASE®

Reference Manual

Sybase Replication Agent™

15.0

[Linux, Microsoft Windows, and UNIX]

DOCUMENT ID: DC00268-01-1500-01

LAST REVISED: October 2006

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

Contents

About This Book

Sybase® Replication Agent™ version 15.0 extends the capabilities of Replication Server® to support the following non-Sybase primary data servers in a Sybase replication system:

- DB2 Universal Database (on UNIX and Microsoft Windows platforms)
- Microsoft SQL Server
- Oracle Database Server

Audience

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

- Database Administrators
- Network Administrators
- System Administrators

How to use this book

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

This book is organized as follows:

Chapter 1, “Command Reference,” describes all Replication Agent commands, including syntax, options, examples, and detailed command usage notes.

Chapter 2, “Configuration Parameters,” describes the Replication Agent configuration file, and provides a configuration parameter reference.

Related documents

Sybase Replication Agent Refer to the following documents to learn more about the Sybase Replication Agent:

- Sybase Replication Agent *Administration Guide* – for an overview of the Sybase Replication Agent, information about configuring and administering Replication Agent instances, and information about configuring the other components in a Sybase replication system.

- Sybase Replication Agent *Primary Database Guide* – for detailed, database-specific information about each non-Sybase database that is supported by the Sybase Replication Agent.
- Sybase Replication Agent *Installation Guide* – for information about installing the Sybase Replication Agent software.
- The Sybase Replication Agent *Release Bulletin* – for last-minute information that was too late to be included in the books.

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

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

- Replication Server *Design Guide* – for an introduction to basic transaction replication concepts and Sybase replication technology.
- Replication Server *Heterogeneous Replication Guide* – for detailed information about configuring Replication Server and implementing a Sybase replication system with non-Sybase databases.

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

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

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

Other sources of information

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

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

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

  Refer to the *SyBooks Installation Guide* on the Getting Started CD, or the *README.txt* file on the SyBooks CD for instructions on installing and starting SyBooks.
- The Sybase Product Manuals Web site is an online version of the SyBooks CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Maintenance, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

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

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

v **Finding the latest information on product certifications**

1 Point your Web browser to Technical Documents at http://www.sybase.com/support/techdocs/.

2 Click Certification Report.

3 In the Certification Report filter select a product, platform, and timeframe and then click Go.

4 Click a Certification Report title to display the report.

v **Finding the latest information on component certifications**

1 Point your Web browser to Availability and Certification Reports at http://certification.sybase.com/.

2 Either select the product family and product under Search by Base Product; or select the platform and product under Search by Platform.

3 Select Search to display the availability and certification report for the selection.

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

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

1 Point your Web browser to Technical Documents at http://www.sybase.com/support/techdocs/.

2 Click MySybase and create a MySybase profile.

Sybase EBFs and software maintenance

v **Finding the latest information on EBFs and software maintenance**

1 Point your Web browser to the Sybase Support Page at http://www.sybase.com/support.

2 Select EBFs/Maintenance. If prompted, enter your MySybase user name and password.

3 Select a product.

4 Specify a time frame and click Go. A list of EBF/Maintenance releases is displayed.

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

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

Style conventions

The following style conventions are used in this book:

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

```
pdb_xlog
```

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

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

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

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

  where *pds* and *pdb* are the variables you should replace.

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

  Use the pdb_xlog command to initialize the primary database.

- In the regular text of this document, names of database objects (tables, columns, stored procedures, etc.) appear like this:

  Check the price column in the widgets table.

- In the regular text of this document, names of datatypes appear like this:

  Use the date or datetime datatype.

- In the regular text of this document, names of files and directories appear like this:

  Log files are located in the *$SYBASE/RAX-15_0/inst_name/log* directory.

Syntax conventions

The following syntax conventions are used in this book:

Table 1: Syntax conventions

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

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

ra_config [*param* [, *value*]]

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

Character case conventions

The following character case conventions are used in this book:

- All command syntax and command examples are shown in lowercase. However, Sybase Replication Agent command names are *not* case sensitive. For example, RA_CONFIG, Ra_Config, and ra_config are equivalent.
- Names of configuration parameters are case sensitive. For example, Scan_Sleep_Max is not the same as scan_sleep_max, and the former would be interpreted as an invalid parameter name.
- Database object names are *not* case sensitive in Replication Agent commands. However, if you need to use a mixed-case object name in a Replication Agent command (to match a mixed-case object name in the primary database), you must delimit the object name with quote characters. For example:

```
pdb_get_tables "TableName"
```

Accessibility features

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

Sybase Replication Agent version 15.0 and the HTML documentation have been tested for compliance with U.S. government Section 508 Accessibility requirements. Documents that comply with Section 508 generally also meet non-U.S. accessibility guidelines, such as the World Wide Web Consortium (W3C) guidelines for Web sites.

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

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

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

For a Section 508 compliance statement for Sybase Replication Agent version 15.0, see Sybase Accessibility at http://www.sybase.com/detail_list?id=52484.

If you need help

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

CHAPTER 1

Command Reference

This chapter describes all of the Sybase Replication Agent commands shown in Table 1-1.

Table 1-1: Replication Agent commands

Command name	Description	Page
log_system_name	Returns the path to the Sybase Replication Agent system log file.	5
pdb_capabilities	Returns a list of the Sybase Replication Agent capabilities.	6
pdb_date	Returns the current date and time from the primary data server.	6
pdb_execute_sql	Executes the specified SQL statement in the current database.	7
pdb_gen_id	Returns the current value of the database generation ID; updates the value of the database generation ID.	8
pdb_get_columns	Returns a list of all the columns in the specified table.	9
pdb_get_databases	Returns a list of all the databases in the primary data server.	11
pdb_get_primary_keys	Returns a list of all the columns that make up the primary keys in the specified table.	11
pdb_get_procedure_parms	Returns a list of the parameters for the specified procedure.	12
pdb_get_procedures	Returns a list of all the procedures in the specified database.	14
pdb_get_sql_database	Returns the name of the database specified for SQL statement execution.	15
pdb_get_tables	Returns a list of all the tables in the specified database.	16
pdb_ownerfilter	Returns a list of owners whose objects will be filtered for initialization; adds, removes owners to the list.	18

Command name	Description	Page
pdb_set_sql_database	Specifies the database to be used for SQL statement execution.	19
pdb_setrepcol	Returns replication marking status; enables or disables replication for all marked columns or a specified column.	20
pdb_setrepddl	Returns DDL replication status; enables or disables replication for DDL statements.	25
pdb_setrepproc	Returns stored procedure replication marking status; marks a specified procedure for replication; unmarks all marked procedures or a specified procedure; enables or disables replication for all marked procedures or a specified procedure.	27
pdb_setrepseq (For Oracle only)	Returns the sequence replication marking status; marks specified sequence for replication; unmarks all marked sequences or a specified sequence; enables or disables replication for all marked sequences or a specified sequence.	36
pdb_setreptable	Returns table replication marking status; marks all tables or a specified table for replication; unmarks all marked tables for replication; enables or disables replication for all marked tables or a specified table.	38
pdb_skip_op	returns, adds or removes record identifiers from a list of records to skip in processing.	49
pdb_truncate_xlog	Truncates the Sybase Replication Agent transaction log.	51
pdb_version	Returns the type and version of the primary data server.	53
pdb_xlog	Returns names of transaction log objects; creates transaction log base objects in the primary database; removes transaction log base objects from the primary database.	53

Command name	Description	Page
suspend	Immediately stops all Log Reader activity, drops connections, and puts Sybase Replication Agent in *Admin* state.	95
test_connection	Tests Replication Agent connections.	96
trace	Returns current trace flag settings; changes a specified trace flag.	99

The remaining sections in this chapter describe each Sybase Replication Agent command in detail.

log_system_name

Description

Returns the full path of the Sybase Replication Agent instance log file.

Syntax

log_system_name

Usage

- When you create a Sybase Replication Agent instance, a log directory is created automatically as part of the instance directory structure. The default value of the log_directory parameter points to that directory.
- The default path of the Sybase Replication Agent log directory is:

  ```
  %SYBASE%\RAX-15_0\inst_name\log\
  ```

  where:

  - *%SYBASE%* is the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Sybase Replication Agent instance.
- If you specify a valid directory path as the value of the log_directory parameter, the Sybase Replication Agent instance places its system log file in the directory you specify.

  If you change the value of the log_directory parameter with the ra_config command, the new value is recorded in the configuration file immediately, but you must shut down and restart the Sybase Replication Agent instance to make the new value take effect.

  See Chapter 2, “Configuration Parameters,” for more information.

- The log_system_name command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also: ra_config, trace

pdb_capabilities

Description: Returns a list of Sybase Replication Agent capabilities, which is used by the Replication Server Manager.

Syntax: pdb_capabilities [*true, false*]

Usage:

- When pdb_capabilities is invoked, it returns a list of the capabilities of the Sybase Replication Agent instance.
- The purpose of the pdb_capabilities command is to support the Replication Server Manager.
- The pdb_capabilities command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

pdb_date

Description: Returns the current date and time from the primary data server.

Syntax: pdb_date

Usage:

- When pdb_date is invoked, it returns the current date and time from the primary data server in the form of a Sybase datetime datatype, as follows:

```
Current PDB Date
 ------------------------------
     Jan 30 2004 12:09:47.310
(1 row affected)
```

- The pdb_date command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also: ra_date

pdb_execute_sql

Description — Executes a SQL statement in the current database at the primary data server.

Syntax — pdb_execute_sql *statement*

Parameters — *statement*

A string in the form of a SQL statement enclosed in double quotes.

Usage

- The Sybase Replication Agent instance executes the specified SQL statement against the "current" database.

  The current database is either:

  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command (to which the Sybase Replication Agent instance is currently connected).
- To set or change the current database, use the pdb_set_sql_database command.
- To find the name of the current database, use the pdb_get_sql_database command.

  Note If the pdb_set_sql_database command has not been invoked to set or change the current database, the pdb_get_sql_database command returns the name of the default current database.
- The SQL statement specified in the pdb_execute_sql command must be a single SQL command enclosed in double quotes. For example:

  ```
  pdb_execute_sql "select * from Authors"
  ```

  The string is passed directly to the database for execution. No command to terminate is required and no syntax or other validation is performed.
- Any results returned from execution of the SQL statement are passed to the Sybase Replication Agent administrative client, by way of the Sybase Replication Agent administration port.
- The pdb_execute_sql command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also — pdb_get_sql_database, pdb_set_sql_database

pdb_gen_id

Description

Returns the current value of the database generation ID, or updates the value of the database generation ID.

Syntax

pdb_gen_id [*number*]

Parameters

number

The value of the new database generation ID to be used when the database generation ID is updated.

Examples

Example 1

```
pdb_gen_id
```

This command returns the current value of the database generation ID.

Example 2

```
pdb_gen_id 10
```

This command updates the database generation ID to the value 10.

Usage

- When pdb_gen_id is invoked with no option, it returns the current value of the database generation ID stored in the Sybase Replication Agent transaction log system table (DB2 UDB, and Microsoft SQL Server) or in the RASD (Oracle).
- When pdb_gen_id is invoked with the number option, it updates the value of the database generation ID in the Sybase Replication Agent transaction log system table (DB2 UDB, and Microsoft SQL Server) or in the RASD (Oracle). Changing the database generation ID takes effect immediately.
- The database generation ID is the first 2 bytes of the origin queue ID. The database generation ID is used by Replication Server to support recovery operations, which may require the Sybase Replication Agent to re-send transactions.

  During recovery, if the Sybase Replication Agent must re-send operations that the Replication Server has already processed, you can change the database generation ID to prevent the Replication Server from recognizing the operations as already processed.
- For more information about the origin queue ID, see ra_helplocator on page 69, or refer to the chapter for your specific primary data server in the Sybase Replication Agent *Primary Database Guide*.
- If the Sybase Replication Agent transaction log (DB2 UDB, and Microsoft SQL Server) or the RASD (Oracle) does not exist, the pdb_gen_id command returns an error.

- The pdb_gen_id command with no parameters is valid when the Sybase Replication Agent instance is in the *Admin* state.

See also ra_locator, pdb_truncate_xlog

pdb_get_columns

Description Returns a list of columns in tables in the current database at the primary data server.

Syntax pdb_get_columns [*ownername*, *tablename*[, *colname*]]

Parameters

ownername
The user name of the owner of the table specified in the *tablename* option. This option can be delimited with quote characters to specify character case.

tablename
The name of the table in the current database for which information is returned. This option can be delimited with quote characters to specify character case.

colname
The name of the column for which information is returned. This option can be delimited with quote characters to specify character case.

Examples

Example 1

```
pdb_get_columns
```

This command returns a list of all of the columns in all of the user tables in the current database.

Example 2

```
pdb_get_columns bob, authors
```

This command returns a list of all of the columns in the table authors, owned by the user “bob” in the current database.

Example 3

```
pdb_get_columns bob, authors, au_fname
```

This command returns information about the column au_fname in the table authors, owned by the user “bob” in the current database.

Usage

- When pdb_get_columns is invoked with no option, it returns a result set that lists all of the columns in all of the user tables in the current database.

- When pdb_get_columns is invoked with the *ownername* and *tablename* options, it returns a result set that lists all of the columns in the specified table with the specified owner in the current database.
- When pdb_get_columns is invoked with the *ownername*, *tablename*, and *colname* options, it returns a result set with information about the specified column in the specified table with the specified owner in the current database.
- The pdb_get_columns command accepts the % wildcard character in the *ownername*, *tablename*, and *colname* options.
- The current database is either:
  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command (to which the Sybase Replication Agent instance is currently connected).
- To set or change the current database, use the pdb_set_sql_database command.

Note If the pdb_set_sql_database command has not been invoked to set or change the current database, the pdb_get_columns command returns information from the current database.

- To find the name of the current database, use the pdb_get_sql_database command.
- The pdb_get_columns command returns 0 rows if the specified table (with the specified owner) does not exist in the current database or if the specified column does not exist in the specified table.
- The pdb_get_columns command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

pdb_get_databases, pdb_get_primary_keys, pdb_get_procedure_parms, pdb_get_procedures, pdb_get_tables

pdb_get_databases

Description

Returns a list of all user databases in the primary data server.

Note The Oracle data server does not support multiple user databases. The pdb_get_databases command returns the name of the database instance.

Syntax

pdb_get_databases

Usage

- When pdb_get_databases is invoked, it returns a result set that lists all of the user databases in the primary data server.

  Note System databases may or may not be returned by some primary data servers. See the chapter for your specific primary data server in the Sybase Replication Agent *Database Guide* for more information.

- The pdb_get_databases command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

pdb_get_columns, pdb_get_primary_keys, pdb_get_procedure_parms, pdb_get_procedures, pdb_get_tables

pdb_get_primary_keys

Description

Returns a list of primary key columns in a specified table in the current database at the primary data server.

Syntax

pdb_get_primary_keys *ownername*, *tablename*

Parameters

ownername

The user name of the owner of the table specified in *tablename*. This option can be delimited with quote characters to specify character case.

tablename

The name of the table in the current database for which primary key column information is returned. This option can be delimited with quote characters to specify character case.

Usage

- When pdb_get_primary_keys is invoked, it returns a result set that lists all of the columns that are defined as primary keys in the specified table with the specified owner in the current database.

- The pdb_get_primary_keys command accepts the % wildcard character in the *ownername* option, but not in the *tablename* option.
- The current database is either:
  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command to which the Sybase Replication Agent instance is currently connected. (This is not valid for Oracle.)
- To set or change the current database, use the pdb_set_sql_database command.

Note In Oracle, you cannot change the current database.

- To find the name of the current database, use the pdb_get_sql_database command.
- The pdb_get_primary_keys command returns 0 rows if the specified table with the specified owner does not exist in the current database.
- The pdb_get_primary_keys command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also pdb_get_columns, pdb_get_databases, pdb_get_procedure_parms, pdb_get_procedures, pdb_get_tables

pdb_get_procedure_parms

Description Returns a list of input parameters for procedures in the current database at the primary data server.

Syntax pdb_get_procedure_parms [*ownername*, *procname* [, *paramname*]]

Parameters

ownername
The user name of the owner of the procedure specified in *procname*. This option can be delimited with quote characters to specify character case.

procname
The name of the procedure in the current database for which information is returned. This option can be delimited with quote characters to specify character case.

paramname
The name of the input parameter for which information is returned. This option can be delimited with quote characters to specify character case.

Examples

Example 1

```
pdb_get_procedure_parms
```

This command returns a list of all of the input parameters for all of the procedures in the current database.

Example 2

```
pdb_get_procedure_parms bob, sp_foo
```

This command returns a list of all of the input parameters for the procedure named sp_foo, owned by the user "bob" in the current database.

Example 3

```
pdb_get_procedure_parms bob, sp_foo, foo_count
```

This command returns information about the input parameter foo_count for the procedure sp_foo, owned by the user "bob" in the current database.

Usage

- When pdb_get_procedure_parms is invoked with no option, it returns a result set that lists all of the input parameters for all the procedures in the current database.
- When pdb_get_procedure_parms is invoked with the *ownername* and *procname* options, it returns a result set that lists all of the input parameters for the specified procedure with the specified owner in the current database.
- When pdb_get_procedure_parms is invoked with the *ownername*, *procname*, and *paramname* options, it returns a result set with information about the specified input parameter for the specified procedure with the specified owner in the current database.
- The pdb_get_procedure_parms command accepts the % wildcard character in both the *ownername* and *procname* options.
- The current database is either:
  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command to which the Sybase Replication Agent instance is currently connected. (This is not valid for Oracle.)

- To set or change the current database, use the pdb_set_sql_database command.

  Note In Oracle, you cannot change the current database.

- To find the name of the current database, use the pdb_get_sql_database command.
- The pdb_get_procedure_parms command returns 0 rows if the specified procedure (with the specified owner) does not exist in the current database.
- The pdb_get_procedure_parms command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

pdb_get_columns, pdb_get_databases, pdb_get_primary_keys, pdb_get_procedures, pdb_get_tables

pdb_get_procedures

Description

Returns a list of procedures in the current database at the primary data server.

Syntax

pdb_get_procedures [*ownername*, *procname*]

Parameters

ownername

The user name of the owner of the procedure specified in *procname*. This option can be delimited with quote characters to specify character case.

procname

The name of the procedure in the current database for which information is returned. This option can be delimited with quote characters to specify character case.

Examples

Example 1

```
pdb_get_procedures
```

This command returns a list of all of the procedures in the current database.

Example 2

```
pdb_get_procedures bob, sp_foo
```

This command returns information about the procedure named sp_foo, owned by the user “bob” in the current database.

Usage

- When pdb_get_procedures is invoked with no option, it returns a result set that lists all of the procedures in the current database.
- When pdb_get_procedures is invoked with the *ownername* and *procname* options, it returns a result set with information about the specified procedure with the specified owner in the current database.
- The pdb_get_procedures command accepts the % wildcard character in both the *ownername* and *procname* options.
- The current database is either:
  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command to which the Sybase Replication Agent instance is currently connected. (This is not valid for Oracle.)
- To set or change the current database, use the pdb_set_sql_database command.

Note In Oracle, you cannot change the current database.

- To find the name of the current database, use the pdb_get_sql_database command.
- The pdb_get_procedures command returns 0 rows if the specified procedure (with the specified owner) does not exist in the current database.
- The pdb_get_procedures command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

pdb_get_columns, pdb_get_databases, pdb_get_primary_keys, pdb_get_procedure_parms, pdb_get_tables

pdb_get_sql_database

Description

Returns the name of the current database, if any.

Syntax

pdb_get_sql_database

Usage

- When pdb_get_sql_database is invoked, it returns the name of the current database.

- If the pdb_set_sql_database command has not been invoked to set the current database, the pdb_get_sql_database command returns the default current database.
- The current database is either:
  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command to which the Sybase Replication Agent instance is currently connected. (This is not valid for Oracle.)
- To set or change the current database, use the pdb_set_sql_database command.

  Note In Oracle, you cannot change the current database.

- The pdb_get_sql_database command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also pdb_execute_sql, pdb_set_sql_database

pdb_get_tables

Description Returns a list of user tables in the current database at the primary data server.

Syntax pdb_get_tables [*ownername*, *tablename*]

Parameters *ownername*
The user name of the owner of the table specified in *tablename*. This option can be delimited with quote characters to specify character case.

tablename
The name of the table in the current database for which information is returned. This option can be delimited with quote characters to specify character case.

Examples **Example 1**

```
pdb_get_tables
```

This command returns a list of all of the user tables in the current database.

Example 2

```
pdb_get_tables bob, authors
```

This command returns information about the table authors, owned by the user "bob" in the current database.

Usage

- When pdb_get_tables is invoked with no option, it returns a result set that lists all of the user tables in the current database.

  Note System tables may or may not be returned by some primary data servers when the pdb_get_tables command is invoked.

- When pdb_get_tables is invoked with the *ownername* and *tablename* options, it returns a result set with information about the specified table with the specified owner in the current database.
- The pdb_get_tables command accepts the % wildcard character in the both the *ownername* and *tablename* options.
- The current database is either:
  - The *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter, or
  - The database specified in the pdb_set_sql_database command to which the Sybase Replication Agent instance is currently connected. (This is not valid for Oracle.)
- To set or change the current database, use the pdb_set_sql_database command.

  Note In Oracle, you cannot change the current database.

- To find the name of the current database, use the pdb_get_sql_database command.
- The pdb_get_tables command returns 0 rows if the specified table (with the specified owner) does not exist in the current database.
- The pdb_get_tables command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

pdb_get_columns, pdb_get_databases, pdb_get_primary_keys, pdb_get_procedure_parms, pdb_get_procedures

pdb_ownerfilter

Description

Note This command is available only for Oracle.

Returns a list of the owners whose objects will be filtered for initialization; adds, removes owners to the list.

Syntax

To return a list of the owners whose objects will be filtered for initialization:

```
pdb_ownerfilter
```

To add or remove an owner whose objects will be filtered for initialization:

```
pdb_ownerfilter [add | remove], owner
```

Parameters

add

The add keyword filters out any objects that are owned by the owner you specify. Any objects that are owned by this owner cannot be marked for initialization.

remove

The remove keyword removes the filter for the owner you specify. Any objects that are owned by this owner can be marked for initialization. You cannot remove the "SYS" owner.

owner

The name of the owner that is used for filtering.

The *owner* option can be delimited with quote characters to specify the character case.

If mixed case (uppercase and lowercase) is required, the name must be delimited. This parameter can be delimited with quotes to specify the character case. If mixed case is required, the name must be delimited. For example:

```
"Owner", "oWnEr"
```

Examples

Example 1

```
pdb_ownerfilter
```

This command returns a list of all owners whose objects can be replicated.

Example 2

```
pdb_ownerfilter add, SYSTEM
```

This command adds the "system" user to the list of owners whose objects will be filtered for replication.

Example 3

```
pdb_ownerfilter remove, SYSTEM
```

This command removes the "system" user from the list of owners whose objects will be filtered for replication.

Usage

- When pdb_ownerfilter is invoked, its function is determined by the keywords and options you specify.
- When multiple keywords and options are specified, each must be separated by a comma. Blank space before or after a comma is optional. For example:

  ```
  pdb_ownerfilter add, system
  ```

- You cannot remove the "SYS" owner.
- After initialization you can replicate any object with pdb_setreptable and pdb_setrepproc, *except* for the following objects which *cannot* be replicated at any time:
  - Objects that are owned by "SYS" owner
  - Any system table whose name begins with V$
  - Any system procedure or package whose name begins with DBMS
- When pdb_ownerfilter is invoked with no keyword, it returns a list of users whose objects will be filtered.
- The pdb_ownerfilter command is valid only when the Sybase Replication Agent instance is in *Admin* state.

See also

pdb_setrepproc, pdb_setreptable, ra_config

pdb_set_sql_database

Description

Sets the current database to be used for SQL statement execution.

Syntax

pdb_set_sql_database *database*

Parameters

database

The name of the database in the primary data server against which the Sybase Replication Agent can execute SQL statements (queries). This parameter can be delimited with quote characters to specify character case.

Usage

- When pdb_set_sql_database is invoked, it sets the "current" database, in which the Sybase Replication Agent can execute SQL queries.

  Note The pdb_set_sql_database command has no affect for Oracle, but it is included to provide continuity with other Replication Agents that support database servers with multiple databases.

- The Sybase Replication Agent does *not* validate the database name you specify with the pdb_set_sql_database command.

  If you specify an invalid database name, no error is returned until one of the following Sybase Replication Agent commands is invoked:

  - pdb_execute_sql
  - pdb_get_columns
  - pdb_get_primary_keys
  - pdb_get_procedure_parms
  - pdb_get_procedures
  - pdb_get_tables

- To find the name of the current database, use the pdb_get_sql_database command.

  Note If the pdb_set_sql_database command has not been invoked to set the current database, the pdb_get_sql_database command returns the *default* current database, which is the primary database specified in the Sybase Replication Agent pds_database_name configuration parameter.

- The pdb_set_sql_database command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

pdb_execute_sql, pdb_get_sql_database

pdb_setrepcol

Description

Returns LOB column replication status; enables or disables replication for LOB columns within marked tables.

Syntax

To return replication status of all columns in all tables or all columns in a specific table:

pdb_setrepcol [*tablename*|enable|disable]

To return replication status of a specific column in a specific table:

pdb_setrepcol *tablename, colname*

To enable or disable all LOB columns in all marked tables:

pdb_setrepcol all, {enable|disable [, force]}

To enable or disable replication for a specified LOB column:

pdb_setrepcol *tablename, colname*, {enable|disable [, force]}

Parameters

tablename

The name of the user table in the primary database that contains the column specified in the *colname* option.

The *tablename* option can be owner-qualified (include the owner name), with each element separated by a period. For example:

```
owner.table
```

The *tablename* option can be delimited with quote characters to specify the character case.

If mixed case (uppercase and lowercase) is required, the name must be delimited. For example:

```
"Owner".table
"Owner"."Table"
```

Each mixed-case element of the *tablename* option must be delimited separately, as shown in the previous example.

Note If you must use an object name case that does not match the value of the ltl_character_case parameter, the object name must be delimited.

If an object name contains any non-alphanumeric characters, such as spaces or periods, it must be delimited with quote characters. For example:

```
"table name"
owner."table name"
```

If an object name contains a period, it must be *both* owner-qualified and delimited with quote characters. For example:

```
owner."table.name"
"table.owner"."table.name"
```

colname

The name of a LOB column in the user table specified in the *tablename* option.

The *colname* option can be delimited with quote characters to specify the character case.

If mixed character case (both uppercase and lowercase) is required, the name must be delimited. For example:

```
"Colname"
"COLname"
```

Note If you must use a column name case that does not match the value of the ltl_character_case parameter, the column name must be delimited. See "ltl_character_case" on page 121 for more information.

---

all

A keyword that refers to all LOB columns in marked tables in the primary database. By using the all keyword, you can apply an enable or disable operation to all LOB columns in marked tables.

enable

A keyword that refers to enabling replication for LOB columns.

disable

A keyword that refers to disabling replication for LOB columns.

force

A keyword that refers to forcing replication to be disabled for LOB columns.

When the force keyword follows the disable keyword, the pdb_setrepcol command immediately disables replication for the specified LOB column, without first checking for pending operations in the transaction log. When the force keyword follows the disable keyword and the all keyword, the pdb_setrepcol command immediately disables replication for all marked LOB columns in marked tables in the primary database, regardless of any pending operations in the transaction log.

Examples

Example 1

```
pdb_setrepcol
```

This command returns replication information for all enabled LOB columns in marked tables in the primary database.

Example 2

```
pdb_setrepcol authors, picture
```

This command returns replication information for the column called “picture” in the table authors in the primary database.

Example 3

```
pdb_setrepcol authors, picture, enable
```

This command enables replication for the column “picture” in the table “authors” in the primary database.

Example 4

```
pdb_setrepcol all, disable
```

This command disables replication for all LOB columns in all marked tables in the primary database.

Usage

- When pdb_setrepcol is invoked, its function is determined by the keywords and options you specify.
- When multiple keywords or options are specified, each must be separated by a comma. Blank space before or after a comma is optional. For example:

  ```
  pdb_setrepcol all, disable
  ```

- When you specify a column name in the pdb_setrepcol command, you must use the name of a valid LOB column.
- You cannot specify the following items as a table name in the pdb_setrepcol command:
  - Primary database system tables
  - Aliases or synonyms
  - Views
  - Sybase Replication Agent transaction log objects
- If a column name in the primary database is the same as a keyword, it can be identified by adding the string col= to the beginning of the column name. For example:

  ```
  pdb_setrepcol tablename, col=enable, disable
  ```

- If you enable LOB column replication with the pdb_setrepcol command do not configure the Sybase Replication Agent to convert date or time datatypes in the primary database. See “pdb_convert_datetime” on page 128 for more information.

- When pdb_setrepcol is invoked with either no option or a single option, it returns information about the enabled status of LOB columns in the primary database.
  - If pdb_setrepcol is invoked with no option, it returns a list of all LOB columns for which replication is enabled in the primary database.

    Note Invoking the pdb_setrepcol command with no option produces the same result as invoking the pdb_setrepcol command with the enable keyword.

  - If pdb_setrepcol is invoked with a table name, it returns information about the enabled status of all the LOB columns in the specified primary table.
  - If pdb_setrepcol is invoked with the enable keyword, it returns a list of all LOB columns for which replication is enabled in the primary database.
  - If pdb_setrepcol is invoked with the disable keyword, it returns a list of all LOB columns for which replication is disabled in the primary database.

  LOB columns for which replication is enabled are listed in the marked objects table.

  Note Any LOB columns in Sybase Replication Agent transaction log tables and shadow tables are not included in the list of LOB columns for which replication is disabled. Also not included are any synonyms, views, or aliases of these database objects.

  For LOB columns listed as disabled, transactions will not be captured for replication.
- When pdb_setrepcol is invoked with a valid primary table name and valid LOB column name, with no keywords, it returns information about the enabled status of the specified LOB column in the specified table in the primary database.
- When pdb_setrepcol is invoked with the all keyword, the operation specified by the following keyword (enable or disable) is applied to all LOB columns in marked tables in the primary database.

  - If pdb_setrepcol is invoked with the all keyword and the enable keyword, it enables replication for all LOB columns in marked tables in the primary database.
  - If pdb_setrepcol is invoked with the all keyword and the disable keyword, it disables replication for all LOB columns in marked tables in the primary database.
- When pdb_setrepcol is invoked with a valid primary table name and valid LOB column name followed by one or more keywords, the operation specified by the keyword (enable or disable) is applied to the specified LOB column in the specified primary table.
  - If pdb_setrepcol is invoked with a table name and LOB column name and the enable keyword, it enables replication for the specified LOB column in the primary database.
  - If pdb_setrepcol is invoked with a table name and LOB column name and the disable keyword, it disables replication for the specified LOB column in the primary database.

  If the table name and LOB column name combination you specify does not exist in the primary database, the pdb_setrepcol command returns an error.
- If the Sybase Replication Agent transaction log does not exist in the primary database (DB2 UDB or Microsoft SQL Server) or the RASD is not initialized (Oracle), the pdb_setrepcol command returns an error.
- If pdb_setrepcol is invoked with a table containing a "DATE" column, the primary key in the primary table must *not* include the "DATE" column. This is true for all Sybase Replication Agents except Sybase Replication Agent for Oracle.

See also pdb_setrepproc, pdb_setreptable, ra_config

pdb_setrepddl

Description

Note This command is available only for Oracle.

Returns DDL replication status; enables or disables replication for DDL statements.

Syntax

To return replication status of DDL:

```
pdb_setrepddl
```

To enable or disable DDL replication:

```
pdb_setrepddl {enable|disable}
```

The default setting is disable.

Examples

Example 1

```
pdb_setrepddl
```

This command returns the current DDL replication status for the primary database.

Example 2

```
pdb_setrepddl enable
```

This command enables replication of DDL commands issued into the primary database after this point in time.

Example 3

```
pdb_setrepddl disable
```

This command disables replication of DDL commands issued into the primary database after this point in time.

Usage

- This is the flag that turns DDL replication on or off. By default, it is set to off (disable).
- In addition to setting enabling DDL replication using pdb_setrepddl command, you must set the Sybase Replication Agent ddl_username and ddl_password parameters to replication DDL in Oracle.

  If the Sybase Replication Agent has not been initialized, the pdb_setrepddl command returns an error.
- The pdb_setrepddl command with enable/disable option is valid when the instance is in *Admin* state.
- The pdb_setrepddl command with no option is valid only when the instance is in *Admin* or *Replicating* state.

Note Some DDL commands are filtered even when DDL replication is enabled. See the Replication Agent *Primary Database Guide* for more information.

See also

ddl_password, ddl_username

pdb_setrepproc

Description

Returns stored procedure replication marking status; unmarks all marked procedures or a specified procedure; enables or disables replication for all marked procedures or a specified procedure.

Note Sybase Replication Agent does not support procedure replication from DB2 Universal Database.

Syntax

To return stored procedure replication marking status:

pdb_setrepproc [*procname*|mark|unmark|enable|disable]

To unmark, enable, or disable all marked stored procedures:

pdb_setrepproc all, {unmark[, force]|enable|disable}

To mark, unmark, enable, or disable a specified stored procedure:

pdb_setrepproc *procname*, {mark|unmark[, force]|enable|disable}

Parameters

procname

The name of a user stored procedure in the primary database.

The *procname* option can be delimited with quote characters to specify the character case.

If mixed character case (both uppercase and lowercase) is required, the name must be delimited. For example:

```
"Proc"
```

Note If you must use an object name case that does not match the setting of the ltl_character_case parameter, the object name must be delimited.

If an object name contains any non-alphanumeric characters, such as spaces, periods, and so forth, it must be delimited with quote characters. For example:

```
"proc name"
"proc.name"
```

If an object name contains a period, it must be *both* owner-qualified and delimited with quote characters. For example:

```
owner."proc.name"
"proc.owner"."proc.name"
```

repname

The *replicated name* of the stored procedure specified in a function replication definition for the primary stored procedure.

The *repname* option can be delimited with quote characters to specify character case. See the previous description of the *procname* option for details.

By specifying a replicated name, stored procedure invocations can be replicated to a stored procedure invocation in the replicate database that has a different stored procedure name from the primary database.

Note The replicated name you specify with the pdb_setrepproc command must match the name of a Replication Server function replication definition for the primary database connection. The Sybase Replication Agent cannot validate the function replication definition, but if it does not exist, function replication from the primary database will fail.

---

all

A keyword that refers to all user stored procedures in the primary database. By using the all keyword, you can mark all user stored procedures, or apply an unmark, enable, or disable operation to all *marked* stored procedures.

mark

A keyword that refers to marking user stored procedures for replication.

unmark

A keyword that refers to unmarking marked stored procedures.

Note For trigger-based instances for Microsoft SQL, Sybase Replication Agent must be in *Admin* state in order to unmark.

---

force
A keyword that refers to the unmark operation.

When the force keyword follows the unmark keyword, the pdb_setrepproc command immediately unmarks the specified stored procedure in the primary database, without first checking the enable status of the stored procedure or checking for pending operations in the transaction log. When the force keyword follows the unmark keyword and the all keyword, the pdb_setrepproc command immediately removes replication marking from all marked stored procedures in the primary database, regardless of their enable status or pending operations in the transaction log.

The force keyword also forces complete execution of the unmarking script, even if errors occur during the unmarking process. Normally, when errors occur during script execution, the script terminates immediately without completing. The force keyword can be useful when a previous script execution failed and left the unmarking operation incomplete.

When errors occur during a forced script execution, the pdb_setrepproc command returns the following message:

```
Errors were encountered and ignored during FORCED script
execution. See error log for details.
```

enable
A keyword that refers to enabling replication for marked stored procedures.

disable
A keyword that refers to disabling replication for marked stored procedures.

Examples

Example 1

```
pdb_setrepproc
```

This command returns replication marking information for all marked stored procedures in the primary database.

Example 2

```
pdb_setrepproc authors
```

This command returns replication marking information for the user stored procedure named “authors” in the primary database.

Example 3

```
pdb_setrepproc authors, mark
```

This command marks the user stored procedure named “authors” in the primary database.

Example 4

```
pdb_setrepproc authors, enable
```

This command enables replication for the marked stored procedure named "authors" in the primary database.

Example 5

```
pdb_setrepproc all, unmark
```

This command unmarks all marked stored procedures in the primary database.

Usage

- For Oracle, you must disable DDL replication before marking or unmarking a procedure, and re-enable it after marking or unmarking.
- When pdb_setrepproc is invoked, its function is determined by the keywords and options you specify.
- When multiple keywords and options are specified, each must be separated by a comma. Blank space before or after a comma is optional. For example:

  ```
  pdb_setrepproc all, unmark, force
  ```

- When you specify a stored procedure name in the pdb_setrepproc command, you must use the name of a valid user stored procedure.
- You cannot specify the following items as a stored procedure name in the pdb_setrepproc command:
  - System procedures
  - Sybase Replication Agent transaction log procedures
- If a stored procedure name in the primary database is the same as a keyword, it can be identified by adding the string `proc=` to the beginning of the stored procedure name. For example:

  ```
  pdb_setrepproc proc=unmark, mark
  ```

- Marking and unmarking a stored procedure for replication requires that the Sybase Replication Agent drop, and then re-create the procedure. However, Sybase Replication Agent sets all the same privileges on the re-created procedure as those defined on the original procedure.

  Note Do *not* remove or alter the Sybase Replication Agent comments in a marked stored procedure.

- When pdb_setrepproc is invoked to mark a procedure for replication, Sybase Replication Agent does the following:
  - Modifies the user procedure to add code that captures input parameter values and generates Sybase Replication Agent transaction log records.
  - Generates a SQL script that creates the procedures required for the Sybase Replication Agent transaction log in the primary database.
  - Saves the generated script in a file called *mark.sql* in the *RAX-15_0\inst_name\scripts\procname* directory, where *inst_name* is the name of the Sybase Replication Agent instance, and *procname* is the name of the stored procedure being marked.

    Note If the value of the pdb_auto_run_scripts configuration parameter is false, the *mark.sql* script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

  - Executes the script to mark the stored procedure and create the transaction log objects in the primary database (if the value of the pdb_auto_run_scripts configuration parameter is true).
  - After the script completes successfully, moves the *partmark.sql* file to the *RAX-15_0\inst_name\scripts\procname\installed* directory.
  - If the mark script fails, it is stored in a file (*partmark.sql*) in the *RAX-15_0\inst_name\scripts\procname* directory, the stored procedure is not marked, and transaction log objects are not created. You can examine the script by viewing the *mark.sql* file.
- When pdb_setrepproc is invoked to unmark a marked stored procedure, Sybase Replication Agent does the following:

- Modifies the user procedure to remove Sybase Replication Agent code that captures input parameter values and generates transaction log records.
- Generates a SQL script that removes the tables and procedures required for the transaction log in the primary database.
- Saves the generated script in a file called *unmark.sql* in the *RAX-15_0\inst_name\scripts\procname* directory, where *inst_name* is the name of the Sybase Replication Agent instance and *procname* is the name of the stored procedure being unmarked. For Oracle, the script is named *partmark.sql* because it can *not* be manually executed—it is for informational purposes only.

  Note If the value of the pdb_auto_run_scripts configuration parameter is false, the *partmark.sql* script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

- Executes the script to unmark the stored procedure and remove the transaction log objects in the primary database (if the value of the pdb_auto_run_scripts configuration parameter is true).
- After the script completes successfully, moves the *unmark.sql* file to the *RAX-15_0\inst_name\scripts\procname\installed* directory.

  Note For Oracle, the file is called *partunmark.sql*.

- If the *unmark* script fails, it is stored in a file (*unmark.sql*) in the *RAX-15_0\inst_name\procname\scripts* directory and the stored procedure is not unmarked and the transaction log objects are not removed. You can examine the script by viewing the *unmark.sql* file.

  When the *unmark* script execution encounters a fatal error on any database object, the pdb_setrepproc command returns the following message:

  ```
  Could not unmark the following objects: ...
  See error log for details.
  ```

- When you use the unmark keyword to remove replication marking from a stored procedure, the Sybase Replication Agent verifies that replication is disabled for that stored procedure and there are no pending (unprocessed) operations for that stored procedure in the transaction log. If replication is not disabled for that procedure, or if there is a pending operation for that procedure in the transaction log, pdb_setrepproc returns an error.
- When pdb_setrepproc is invoked with either no option or a single option, it returns marking information about the stored procedures in the primary database.
  - If pdb_setrepproc is invoked with no option, it returns a list of all marked procedures in the primary database.

    Note Invoking the pdb_setrepproc command with no option produces the same result as invoking the pdb_setrepproc command with only the mark keyword.

  - If pdb_setrepproc is invoked with a procedure name, it returns complete marking information about the specified procedure.
  - If pdb_setrepproc is invoked with the mark keyword, it returns a list of all marked procedures in the primary database.
  - If pdb_setrepproc is invoked with the unmark keyword, it returns a list of all unmarked procedures in the primary database.
  - If pdb_setrepproc is invoked with the enable keyword, it returns a list of all marked procedures in the primary database, for which replication is currently enabled.
  - If pdb_setrepproc is invoked with the disable keyword, it returns a list of all marked procedures in the primary database, for which replication is currently disabled.

  Stored procedures marked for replication are listed in the marked objects table. For Oracle, the marking information is recorded in the RASD. All other user procedures are considered unmarked.

  Note The Sybase Replication Agent system procedures are not included in the list of unmarked procedures. Also not included are any synonyms or aliases of these procedures.

  For procedures listed as unmarked or disabled, their invocations will not be captured for replication.

- When pdb_setrepproc is invoked with the all keyword and an action keyword (unmark, enable, or disable), the action specified is applied to either all user stored procedures in the primary database, or to all marked procedures in the primary database.
  - If pdb_setrepproc is invoked with the all and unmark keywords, it removes replication marking from all marked procedures in the primary database.

    You can specify the force keyword after the unmark keyword to force immediate unmarking of all marked procedures, including procedures for which replication is still enabled or pending operations remain in the transaction log.
  - If pdb_setrepproc is invoked with the all and enable keywords, it enables replication for all marked procedures in the primary database.
  - If pdb_setrepproc is invoked with the all and disable keywords, it disables replication for all marked procedures in the primary database.
- When pdb_setrepproc is invoked with a valid user stored procedure name and followed by an action keyword (mark, unmark, enable, or disable), the action specified is applied to the specified procedure.
  - If pdb_setrepproc is invoked with a procedure name and the mark keyword, it marks the specified procedure in the primary database for replication.
  - If pdb_setrepproc is invoked with a procedure name and the unmark keyword, it removes replication marking from the specified procedure in the primary database.

    You can specify the force keyword after the unmark keyword to force immediate unmarking of the specified procedure, unmark a procedure for which replication is still enabled or pending operations remain in the transaction log, or force the script execution to ignore errors and continue an unmarking operation that failed previously.

    If the *unmark* script execution encounters a fatal error on any database object, the pdb_setrepproc command returns the following message:

    ```
    Could not unmark the following objects: ...
    See error log for details.
    ```
  - If pdb_setrepproc is invoked with a procedure name and the enable keyword, it enables replication for the specified marked procedure in the primary database.

If the *enable* script execution encounters a fatal error on any database object, the pdb_setrepproc command returns the following message:

```
Could not enable the following objects: ...
See error log for details.
```

- If pdb_setrepproc is invoked with a procedure name and the disable keyword, it disables replication for the specified marked procedure in the primary database.

  If the *disable* script execution encounters a fatal error on any database object, the pdb_setrepproc command returns the following message:

  ```
  Could not disable the following objects: ...
  See error log for details.
  ```

- If you specify a stored procedure name that does not exist in the primary database, the pdb_setrepproc command returns an error.
- When pdb_setrepproc is invoked with a procedure name and a replicated name, followed by the mark keyword, the primary procedure is marked for replication with the specified replicated name.

  If the primary procedure name you specify does not exist in the primary database, the pdb_setrepproc command returns an error.

  By specifying a replicated name, procedure invocations can be replicated to a procedure in the replicate database that has a different name from the primary procedure.

  Note The replicated name you specify with the pdb_setrepproc command must match the name of a Replication Server function replication definition for the primary database connection. The Sybase Replication Agent cannot validate the function replication definition, but if it does not exist, function replication from the primary database will fail.

- If the Sybase Replication Agent transaction log does not exist in the primary database (DB2 UDB, or Microsoft SQL Server) or the RASD is not initialized (Oracle), the pdb_setrepproc command returns an error.

See also pdb_setrepcol, pdb_setreptable, ra_config

pdb_setrepseq

Description

Note This command is available only for Oracle.

Returns the sequence replication marking status; marks specified sequence for replication; unmarks all marked sequences or a specified sequence; enables or disables replication for all marked sequences or a specified sequence. This command is available for Oracle only.

Syntax

To return sequence replication marking status:

```
pdb_setrepseq [sequence_name|mark|unmark|enable|disable]
```

To unmark, enable, or disable all marked sequences:

```
pdb_setrepseq all, {unmark[, force] |enable|disable}
```

To mark, unmark, enable, or disable a specified sequence:

```
pdb_setrepseq sequence_name, {mark|unmark[, force] |enable|disable}
```

To mark a specified sequence for replication with a replicated name:

```
pdb_setrepseq sequence_name, repname, mark
```

Parameters

sequence_name

The name of a user sequence in the primary database. The *sequence_name* option can be delimited with quote characters to specify the character case. If mixed character case (both uppercase and lowercase) is required, the name must be delimited. For example:

```
"Sequence"
```

The *sequence_name* parameter can be owner-qualified to include the primary sequence owner name, with each element separated by a period. For example:

```
owner.sequence
```

Note If you must use an object name case that does not match the setting of the ltl_character_case parameter, the object name must be delimited. If an object name contains any non-alphanumeric characters, such as spaces and periods, it must be delimited with quote characters. For example, "sequence name" or owner."sequence name."

repname
The replicated name of the sequence to be updated at the replicate site, if desired to be different than the sequence name at the primary site. The *repname* option can be delimited with quote characters to specify character case. See the previous description of the *sequence_name* parameter for details. By specifying a replicated name, sequence updates can be replicated to a sequence in the replicate database that has a different sequence name from the primary database.

The *repname* option can be owner-qualified to include the replicate sequence owner name, with each element separated by a period. For example:

```
repowner.repname
```

all
A keyword that refers to all user sequences in the primary database. By using the *all* keyword, you can unmark all user sequences, or apply an enable or disable operation to all marked sequences.

mark
A keyword that refers to marking user sequences for replication.

unmark
A keyword that refers to unmarking user sequences for replication.

force
A keyword that refers to the unmark operation. When the *force* keyword follows the *unmark* keyword, the pdb_setrepseq command immediately unmarks the specified sequence in the primary database, without first checking the enable status of the sequence. When the *force* keyword follows the *unmark* keyword and the all keyword, the pdb_setrepseq command immediately removes replication marking from all marked sequences in the primary database, regardless of their enable status

enable
A keyword that refers to enabling replication for marked sequences.

disable
A keyword that refers to disabling replication for marked sequences.

Usage

- When pdb_setrepseq is invoked, its function is determined by the keywords and options you specify.
- When multiple keywords and options are specified, each must be separated by a comma. Blank space before or after a comma is optional. For example:

```
pdb_setrepseq all, unmark, force
```

- When you specify a sequence in the pdb_setrepseq command, you must use the name of a valid user sequence.

pdb_setreptable

Description

Returns replication marking status; marks all user tables or a specified table for replication; unmarks all marked tables or a specified table; enables or disables replication for all marked tables or a specified table.

Syntax

To return replication marking status:

pdb_setreptable [*tablename*|mark|unmark|enable|disable]

To mark all user tables (available only for Oracle):

pdb_setreptable all, mark

To unmark, enable, or disable all marked tables:

pdb_setreptable all, {unmark[, force]|enable|disable}

To mark, unmark, enable, or disable a specified table:

pdb_setreptable *tablename*, {mark[, owner]|
unmark[, force] |enable|disable}

To mark a specified table for replication with a replicated name:

pdb_setreptable *tablename*, *repname*, mark[, owner]

Parameters

tablename

The name of a user table in the primary database.

The *tablename* parameter can be owner-qualified to include the primary table owner name, with each element separated by a period. For example:

```
owner.table
```

This parameter can be delimited with quote characters to specify the character case.

If mixed character case (both uppercase and lowercase) is required, the name must be delimited. For example:

```
"Owner".table
```

```
"Owner"."Table"
```

Each mixed case element of the *tablename* option must be delimited separately, as shown in the previous example.

If an object name contains any non-alphanumeric characters, such as spaces or periods, it must be delimited with quote characters. For example:

```
"table name"
owner."table name"
```

If an object name contains a period, it must be *both* owner-qualified and delimited with quote characters. For example:

```
owner."table.name"
"table.owner"."table.name"
```

repname

The name of the table specified in the replication definition for a primary table.

Note The replicated name you specify with the pdb_setreptable command must match a table name specified by a with primary table named clause in a Replication Server replication definition for the primary database connection. The Sybase Replication Agent cannot validate the replication definition, but if it does not exist, or if the with primary table named clause does not match the replicated name specified with pdb_setreptable, replication from the primary table will fail.

The *repname* option can be owner-qualified to include the replicate table owner name, with each element separated by a period. For example:

```
repowner.reptable
```

Note If you want to use an owner-qualified replicate table name with the replicate owner's name, use the owner keyword with the pdb_setreptable command. If you specify an unqualified replicate table name, the primary table owner name is sent with the replicate table name in the LTL.

The *repname* option can also be delimited with quote characters to specify the character case. See the previous description of the *tablename* option for details.

Note If the replicate table name contains a period (for example, table.name), without owner qualification, you must set the value of the Sybase Replication Agent use_rssd parameter to true.

all
: A keyword that refers to all tables in the primary database. By using the all keyword, you can mark all user tables, or apply an unmark, enable, or disable operation to all marked tables.

mark
: A keyword that refers to replication marking. For trigger-based instances for Microsoft SQL, Sybase Replication Agent must be in *Admin* state in order to unmark a table.

owner
: A keyword that refers to the mark operation.

  When the optional owner keyword follows the mark keyword, the pdb_setreptable command marks the specified table in the primary database so that when operations against that table are replicated, the owner name is included with the table name in the form owner.tablename in the LTL sent to the Replication Server.

Note If you want to use an owner-qualified replicate table name with the replicate owner's name, use the owner keyword with the pdb_setreptable command. If you specify an unqualified replicate table name, the primary table owner name is sent with the replicate table name in the LTL.

unmark
: A keyword that refers to unmarking a marked table.

force
A keyword that refers to the unmark operation.

When the force keyword follows the unmark keyword, the pdb_setreptable command immediately removes replication marking for the specified table in the primary database, without first checking the enable status of the table or checking for pending operations in the transaction log. When the force keyword follows the unmark keyword and the all keyword, the pdb_setreptable command immediately removes replication marking from all marked tables in the primary database, regardless of their enable status or any pending operations in the transaction log.

The force keyword also forces complete execution of the unmarking script, even if errors occur during the unmarking process (Microsoft only). Normally, when errors occur during script execution, the script terminates immediately without completing. The force keyword can be useful when a previous script execution failed and left the unmarking operation incomplete.

When errors occur during a forced script execution, the pdb_setreptable command returns the following message:

```
Errors were encountered and ignored during FORCEd script
execution. See error log for details.
```

enable
A keyword that refers to enabling replication for marked tables.

disable
A keyword that refers to disabling replication for marked tables.

Examples

Example 1

```
pdb_setreptable authors
```

This command returns replication marking information for the table named "authors" in the primary database.

Example 2

```
pdb_setreptable mark
```

This command returns replication marking information for all marked tables in the primary database.

Example 3

```
pdb_setreptable disable
```

This command returns replication marking information for all marked tables for which replication has been disabled in the primary database.

Example 4

```
pdb_setreptable all, unmark, force
```

This command forces unmarking for all marked tables in the primary database.

Example 5

```
pdb_setreptable all, enable
```

This command enables replication for all marked tables in the primary database.

Example 6

```
pdb_setreptable authors, mark
```

This command marks for replication the table named “authors” in the primary database.

Example 7

```
pdb_setreptable authors, mark, owner
```

This command marks for replication the table named “authors” in the primary database so that the name of the primary table owner will be passed along with the table name in the LTL.

Example 8

```
pdb_setreptable authors, auth_name, mark
```

This command marks for replication the table named “authors” in the primary database with a replicate name “auth_name.”

Example 9

```
pdb_setreptable authors, auth_name, mark, owner
```

This command marks for replication the table named “authors” in the primary database with a replicate name “auth_name” so that the name of the primary table owner will be passed along with the replicate name in the LTL.

Example 10

```
pdb_setreptable authors, bob.auth_name, mark, owner
```

This command marks for replication the table named “authors” in the primary database with a replicate name “auth_name” so that the name of the replicate table owner “bob” will be passed along with the replicate name in the LTL.

Example 11

```
pdb_setreptable authors, enable
```

This command enables replication for the marked table “authors” in the primary database.

Example 12

```
pdb_setreptable table=mark, enable
```

This command enables replication for the marked table named “mark” in the primary database.

Example 13

```
pdb_setreptable authors, unmark, force
```

This command forces unmarking for the marked table “authors” in the primary database.

Usage

- When pdb_setreptable is invoked, its function is determined by the keywords and options you specify.
- When multiple keywords and options are specified, each must be separated by a comma. Blank space before or after a comma is optional. For example:

  ```
  pdb_setreptable all, unmark, force
  ```

- When you specify a primary table in the pdb_setreptable command, you must use the name of a valid user table.
- You cannot specify the following items as a primary table in the pdb_setreptable command:
  - System tables
  - Views
  - Sybase Replication Agent transaction log tables
- If you specify an alias or synonym as a primary table in the pdb_setreptable command, the actual table that the alias or synonym refers to is acted upon. The actual table name is the table name sent to the primary Replication Server.
- If a table name in the primary database is the same as a keyword, it can be identified by adding the table=*string* to the beginning of the name. For example:

  ```
  pdb_setreptable table=unmark, mark
  ```

This is true for both primary table names and replicated names.

- When pdb_setreptable is invoked to mark a table for replication, trigger-based Sybase Replication Agent (for Microsoft SQL Server only) do the following:
  - Generate a SQL script that creates the tables, procedures, and triggers required for the transaction log in the primary database.
  - Save the generated script in a file called *mark.sql* in the *RAX-15_0\inst_name\scripts\tablename* directory, where *inst_name* is the name of the Sybase Replication Agent instance, and *tablename* is the name of the table being marked.

    Note If the value of the pdb_auto_run_scripts configuration parameter is false, the *mark.sql* script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

  - Execute the script to mark the table and create the transaction log objects in the primary database (if the value of pdb_auto_run_scripts is true).
  - After the script completes successfully, move the *mark.sql* file to the *RAX-15_0\inst_name\scripts\tablename\installed* directory.
  - If a user trigger exists on the table, modify the user trigger to add code that captures data and generates Sybase Replication Agent transaction log records.
- If the *mark* script fails, it is stored in a file (*mark.sql*) in the *RAX-15_0\inst_name\scripts\tablename* directory, the table is not marked, and transaction log objects are not created. You can examine the script by viewing the *mark.sql* file.
- When pdb_setreptable is invoked to unmark a marked primary table, trigger-based Sybase Replication Agents (for Microsoft SQL Server) do the following:
  - Generate a SQL script that removes the tables, procedures, and triggers required for the transaction log in the primary database.

- Save the generated script in a file called *unmark.sql* in the *RAX-15_0\inst_name\scripts\tablename* directory, where *inst_name* is the name of the Sybase Replication Agent instance and *tablename* is the name of the table being unmarked.

  Note If the value of the pdb_auto_run_scripts configuration parameter is false, the *mark.sql* script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

- Execute the script to unmark the table and remove the transaction log objects in the primary database (if the value of pdb_auto_run_scripts is true).
- After the script completes successfully, move the *unmark.sql* file to the *RAX-15_0\inst_name\scripts\tablename\installed* directory.
- If a user trigger existed on the table when it was marked, modify the user trigger to remove Sybase Replication Agent code that captures data and generates transaction log records.

- If the *unmark* script fails, it is stored in a file (*unmark.sql*) in the *RAX-15_0\inst_name\tablename\scripts* directory, the table is not unmarked, and the transaction log objects are not removed. You can examine the script by viewing the *unmark.sql* file. When the *unmark* script execution encounters a fatal error on any database object, the pdb_setreptable command returns the following message:

  ```
  Could not unmark the following objects: ...
  See error log for details.
  ```

- When you use the unmark keyword to remove replication marking from a primary table, the Sybase Replication Agent verifies that replication is disabled for that table and checks to make sure that there are no pending (unprocessed) operations for that table in the transaction log. If replication is not disabled, or there is a pending operation for that table in the transaction log, pdb_setreptable returns an error.
- When you use the unmark keyword to remove replication marking from primary tables, you can also specify the force keyword to immediately remove replication marking from primary tables, without regard to whether replication is disabled or pending operations exist in the transaction log.

The force keyword also ignores script execution errors. If the *unmark* script execution encounters a fatal error on any database object, the pdb_setreptable command returns the following message:

```
Could not unmark the following objects: ...
See error log for details.
```

- When pdb_setreptable is invoked with either no option or a single option, it returns marking information about the user tables in the primary database:
  - If pdb_setreptable is invoked with no option, it returns a list of all marked tables in the primary database.

    Note Invoking the pdb_setreptable command with no option produces the same result as invoking the pdb_setreptable with the mark keyword.

  - If pdb_setreptable is invoked with a table name, it returns complete marking information about the specified primary table.
  - If pdb_setreptable is invoked with the mark keyword, it returns a list of all marked tables in the primary database.
  - If pdb_setreptable is invoked with the unmark keyword, it returns a list of all unmarked tables in the primary database.
  - If pdb_setreptable is invoked with the enable keyword, it returns a list of all marked tables in the primary database for which replication is enabled.
  - If pdb_setreptable is invoked with the disable keyword, it returns a list of all marked tables in the primary database for which replication is disabled.

  Tables marked for replication are listed in the marked objects table. All other user tables are considered unmarked.

  Note The Sybase Replication Agent transaction log tables and shadow tables are not included in the list of unmarked tables. Also not included are any synonyms, views, or aliases of these database objects.

  For tables listed as unmarked or disabled, transactions will not be captured for replication.

- When pdb_setreptable is invoked with the all keyword and an action keyword (mark, unmark, enable, or disable), the action specified is applied to either all user tables in the primary database, or all marked tables in the primary database.
  - If pdb_setreptable is invoked with the all and mark keywords, all user tables in the primary database are marked for replication.
  - If pdb_setreptable is invoked with the all and unmark keywords, it removes replication marking from all marked tables in the primary database.

    You can specify the force keyword after the unmark keyword to force immediate unmarking of all marked tables, or to unmark tables for which replication is still enabled or pending operations remain in the transaction log, or to force the script execution to ignore errors and continue an unmarking operation that failed previously.
  - If pdb_setreptable is invoked with the all and enable keywords, it enables replication for all marked tables in the primary database.
  - If pdb_setreptable is invoked with the all and disable keywords, it disables replication for all marked tables in the primary database.
- When pdb_setreptable is invoked with a valid user table name, followed by an action keyword (mark, unmark, enable, or disable), the action specified is applied to the specified table.
  - If pdb_setreptable is invoked with a table name and the mark keyword, it marks the specified table in the primary database for replication.

    You can specify the owner keyword after the mark keyword so that when operations against the table are replicated, the primary table object owner name will be attached to the table name in the form owner.tablename in the LTL sent to the Replication Server.

    Note If you want to use owner-qualified table names for either primary tables or replicate tables, you must set the value of the Sybase Replication Agent use_rssd parameter to true.
  - If pdb_setreptable is invoked with a table name and the unmark keyword, it removes replication marking from the specified table in the primary database.

You can specify the force keyword after the unmark keyword to force immediate unmarking of the specified table, to unmark a table for which replication is still enabled or pending operations remain in the transaction log, or to force the script execution to ignore errors and continue an unmarking operation that failed previously.

- If pdb_setreptable is invoked with a table name and the enable keyword, it enables replication for the specified marked table in the primary database.

  If the *enable* script execution encounters a fatal error on any database object, the pdb_setreptable command returns the following message:

  ```
  Could not enable the following objects: ...
  See error log for details.
  ```

- If pdb_setreptable is invoked with a table name and the disable keyword, it disables replication for the specified marked table in the primary database.

  If the *disable* script execution encounters a fatal error on any database object, the pdb_setreptable command returns the following message:

  ```
  Could not disable the following objects: ...
  See error log for details.
  ```

- If the table name you specify does not exist in the primary database, the pdb_setreptable command returns an error.
- When pdb_setreptable is invoked with a primary table name and a replicated name, followed by the mark keyword, the primary table is marked for replication with the specified replicated name.

  If the primary table name you specify does not exist in the primary database, the pdb_setreptable command returns an error.

By specifying a replicated name, transactions can be replicated to a table in the replicate database that has a different name from the primary table.

Note The replicated name you specify with the pdb_setreptable command must match a table name specified by a with all tables named clause in a Replication Server replication definition for the primary database connection. The Sybase Replication Agent cannot validate the replication definition, but if it does not exist, or if the with all tables named clause does not match the replicated name specified with pdb_setreptable, replication from the primary table will fail.

You can also specify the owner keyword after the mark keyword so that when operations against the primary table are replicated, the primary table owner name will be attached to the replicate table name in the form owner.tablename.

Note If you want to use an owner-qualified replicate table name with the replicate owner's name, use the owner keyword with the pdb_setreptable command. If you specify an unqualified replicate table name, the primary table owner name is sent with the replicate table name in the LTL.

- If the Sybase Replication Agent transaction log does not exist in the primary database (DB2 UDB or Microsoft SQL Server) or the RASD is not initialized (Oracle), the pdb_setreptable command returns an error.
- To replicate a table that contains column names that have spaces, you must set structured_tokens to true.

See also pdb_setrepcol, pdb_setrepproc, ra_config

pdb_skip_op

Description

Note This command is available only for Oracle.

Maintains a list of record identifiers. Log records matching these identifiers will not be processed. The format of the identifier is database-specific.

Syntax pdb_skip_op [{add, id} | { remove, { all | id } }]

Parameters

add
: The add keyword adds a specified ID to the list of identifiers of records to skip.

remove
: The remove keyword removes a specified ID from the list of identifiers to skip.

id
: The identifier of the record that you want to add or remove from the list of identifiers to skip.

all
: The all keyword allows you to add or remove all IDs in the list of identifiers to skip.

Examples

Example 1

```
pdb_skip_op
```

This command with no parameters returns a list of the identifiers for the records you want to skip.

Example 2

```
pdb_skip_op add, id
```

This command adds an ID to the list of identifiers you want to skip.

Example 3

```
pdb_skip_op remove, id
```

To remove an ID to the list of identifiers you want to skip:

Example 4

```
pdb_skip_op remove, all
```

This command removes all the IDs on the list of identifiers you want to skip.

Usage

- The pdb_skip_op command allows you to skip problem records, thereby avoid having to reinitialize the Sybase Replication Agent.
- Skipped records are written to the system log as a warning message.
- The pdb_skip_op command is valid when the Sybase Replication Agent is in *Admin* state.
- For Oracle, the format of the identifier is a rollback address (RBA). An RBA has the following form:

  ```
  LSN.BLKNUM.BLKOFFSET
  ```

where:

- LSN is the log sequence number.
- BLKNUM is the block number.
- BLKOFFSET is the offset into the block where this record resides.

The three values must be specified in the pdb_skip_op command together, enclosed in quotes, each item separated by a period.

For example:

```
'0012.0000444.0000123'
```

An RBA value can be described as a hexadecimal by prefixing the identifier with an '0x' as follows:

```
'0x000c.00001bc.000007b'
```

See also ra_helplocator, ra_locator

pdb_truncate_xlog

Description Truncates the Sybase Replication Agent transaction log on demand. Based on the target database the transaction log is processed differently:

- For Oracle, the behavior of this command changes based on the value of configuration parameter pdb_include_archives:
  - When pdb_include_archives is *false*, this command triggers the Oracle archive process to archive any on-line redo logs that have already been processed by Sybase Replication Agent for Oracle.
  - When pdb_include_archives is true, removes old archive redo log files from the path specified by pdb_archive_path.

  .

  Note For Oracle, truncation of the old archive log files from the *pdb* archive path directory is performed only if the pdb_archive_remove property is set to true.

- For trigger-based Replication Agent for Microsoft SQL Server, the *tran log* table and *shadow* tables are truncated.

- For Replication Agent for UDB, the *primary database* log files will be deleted up to, but not including, the log file that contain the lsn found in the current truncation point.

  Warning! For UDB, the Replication Agent deletes the *primary database* log files that it no longer needs. For more information, see the Sybase Replication Agent *Primary Database Guide*

For more information on how Sybase Replication Agent affects each type of database when pdb_truncate is executed, see the Sybase Replication Agent *Primary Database Guide*

Syntax

pdb_truncate_xlog

Usage

- When pdb_truncate_xlog is invoked, Sybase Replication Agent immediately truncates the transaction log based on the most recent truncation point received from the primary Replication Server. The truncation point is part of the information contained in the LTM Locator.
- To update the LTM Locator from the primary Replication Server, use the ra_locator command.
- The pdb_truncate_xlog command is asynchronous and it does not return success or failure (unless an immediate error occurs). You must examine the Replication Agent system log to determine success or failure of the pdb_truncate_xlog command.
- You can use the ra_statistics command to view the "Number of transactions truncated" both before and after you use the pdb_truncate_xlog command.
- If the Sybase Replication Agent transaction log does not exist or if a connection failure occurs, the pdb_truncate_xlog command returns an error message.
- You can use the ra_config command to specify the type of automatic truncation you want. You can use the pdb_truncate_xlog command to truncate the transaction log if automatic truncation is not sufficient to manage the size of the transaction log.
- The pdb_truncate_xlog command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

ra_config, ra_locator

pdb_version

Description — Returns the type and version of the primary data server and the JDBC driver.

Syntax — pdb_version

Usage

- When pdb_version is invoked, it returns the primary data server vendor name and version, and the JDBC driver name and version for the primary database.
- The actual results returned vary depending on the type of primary data server.
- If the primary database connection is down or not configured, the pdb_version command returns an error.
- The pdb_version command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also — ra_version, ra_version_all

pdb_xlog

Description — Returns the names of transaction log base objects; creates transaction log base objects in the primary database; or removes transaction log base objects from the primary database.

For Oracle, this command verifies permissions are valid for the Sybase Replication Agent to obtain system data from the primary database. It also checks the condition of the primary database to determine if archiving is turned on or off, and then loads the RASD with system data from the primary database.

Note Because the pdb_xlog create command is deprecated, Sybase recommends that you use pdb_xlog init.

Syntax

```
pdb_xlog [[init |create | remove]
[, force]] | move_truncpt]
```

Parameters

init

The keyword for moving the truncation point to the end of the transaction log.

create
The keyword for creating a transaction log.

remove
The keyword for removing a transaction log.

force
A keyword that refers to the remove or init operation.

move_truncpt
A keyword that moves the truncation point for Oracle only.

Examples

```
pdb_xlog init 2048, 4096
```

This command creates the xlog base components, and specifies in the transaction log table an initial extent size of 2048 bytes and a next extent size of 4096 bytes.

Usage

- When pdb_xlog is invoked with no option, it returns the actual names (not synonyms or aliases) of all Sybase Replication Agent transaction log base objects in the primary database. For Oracle, if you have initialized the Sybase Replication Agent, it returns the name of the component and the Oracle database instance name.

  Note See the chapter for your specific primary data server in the Sybase Replication Agent *Primary Database Guide* for more information on Sybase Replication Agent object names.

- If pdb_xlog is invoked with no option, and the Sybase Replication Agent transaction log base objects do not exist in the primary database, or the RASD has not been initialized (for Oracle), the command returns no information.
- If pdb_xlog is invoked with the init keyword the truncation point is moved to the end of the log.
- If pdb_xlog is invoked with the init, force keywords the truncation point is moved to the end of the log if the Replication Agent is not already initialized. However, if the Replication Agent is already initialized, the truncation point is not moved.

- If pdb_xlog is invoked with the init, move_truncpt keyword the truncation point is moved to the end of the log (the current online Oracle redo log). The move_truncpt option has no effect if the Replication Agent has not been inialized.

  Note To prevent Replication Server from requesting a log starting point which occurs earlier in the log than the location established by the *move_truncpt* option, the Replication Server's LTM locator value for the primary connection must be zeroed. Execute Replication Server System Database command rs_zeroltm against the primary database connection to zero the LTM locator.

- When pdb_xlog is invoked with the init keyword, Sybase Replication Agent does the following (not supported in Oracle):
  - Generates a SQL script that creates the tables and procedures required for the transaction log base objects in the primary database.
  - Saves the generated script in a file called *create.sql* in the *RAX-15_0\inst_name\scripts\xlog* directory, where *inst_name* is the name of the Sybase Replication Agent instance.

    Note If the value of the pdb_auto_run_scripts configuration parameter is false, the mark.sql script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

  - Executes the script to create the Sybase Replication Agent transaction log base objects in the primary database (if the value of the pdb_auto_run_scripts configuration parameter is true).
  - After the script completes successfully, moves the *create.sql* file to the *RAX-15_0\inst_name\scripts\xlog\installed* directory.

  Note For Oracle, the file is called *partinit.sql* and is for informational purposes only. It can *not* be manually executed.

- If the *create* script fails, it is stored in a file (*create.sql*) in the *RAX-15_0\inst_name\scripts\xlog* directory and the transaction log is not created. You can examine the script by viewing the *create.sql* file.

- If pdb_xlog is invoked with the init keyword, and transaction log base objects already exist in the primary database (using the prefix string specified by the pdb_xlog_prefix configuration parameter) or the RASD has been initialized (for Oracle), then pdb_xlog returns an error message.
- When pdb_xlog is invoked with the remove keyword, Sybase Replication Agent does the following (not supported in Oracle):
  - Generates a SQL script that deletes the tables and procedures required for the transaction log base objects in the primary database.
  - Saves the generated script in a file called *remove.sql* in the *RAX-15_0\inst_name\scripts\xlog* directory, where *inst_name* is the name of the Sybase Replication Agent instance.

    Note If the value of the pdb_auto_run_scripts configuration parameter is false, the *mark.sql* script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

    Note If the value of the pdb_auto_run_scripts configuration parameter is false, the mark.sql script will be saved but not executed automatically. To manually re-run the script, you must first set pdb_auto_run_scripts to true and then re-run the command.

  - Executes the script to delete the transaction log base objects in the primary database (if the value of the pdb_auto_run_scripts configuration parameter is true).
  - After the script completes successfully, moves the *remove.sql* file to the *RAX-15_0\inst_name\scripts\xlog\installed* directory.

    Note For Oracle, the file is called *partdeinit.sql* and is for informational purposes only. It can *not* be manually executed.

- If the *remove* script fails, it is stored in a file (*remove.sql*) in the *RAX-15_0\inst_name\scripts\xlog* directory and the transaction log is not deleted. You can examine the script by viewing the *remove.sql* file.
- When pdb_xlog is invoked with the remove keyword followed by the force keyword, the *remove.sql* script continues executing, even if errors occur. The force keyword may be useful when a previous remove operation failed and the *remove.sql* script terminated with an error.

- If pdb_xlog is invoked with the remove keyword, and transaction log base objects do not exist in the primary database (using the prefix string specified by the pdb_xlog_prefix configuration parameter) or the RASD has not been initialized (for Oracle), then pdb_xlog returns an error message.
- If pdb_xlog is invoked with the remove keyword, and any objects in the primary database are still marked for replication, then pdb_xlog returns an error message.

  You can use the pdb_setrepproc and pdb_setreptable commands to determine which stored procedures and tables in the primary database are still marked. You also can use the pdb_setrepddl command to determine if DDL is enabled.

  Even if objects are marked in the primary database, you can use the pdb_xlog command with the remove keyword followed by the force keyword to unmark any marked objects, and then remove the transaction log objects.
- If pdb_xlog is invoked with no option, the command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.
- If pdb_xlog is invoked with either the init or remove keyword, the command is valid only when the Sybase Replication Agent instance is in the *Admin* state.
- For more information about the Sybase Replication Agent transaction log, see the chapter for your specific primary data server in the Sybase Replication Agent *Primary Database Guide*.

See also pdb_setrepcol, pdb_setrepproc, pdb_setreptable, ra_config

quiesce

Description Stops all Sybase Replication Agent processing in *Replicating* state, and puts the Sybase Replication Agent instance in *Admin* state.

Syntax quiesce

Usage
- When the quiesce command is invoked, it stops all current replication processing in the Sybase Replication Agent instance:

  - The Log Reader component stops reading operations from the transaction log as soon as the current scan is complete. It continues to send change-set data to the Log Transfer Interface component until it finishes processing the last operation scanned from the log.
  - The Log Transfer Interface component stops sending LTL commands to the Replication Server as soon as it finishes processing the last change set it receives from the Log Reader.
  - When the Log Transfer Interface component is finished processing its input queue and sending the resulting LTL, the Sybase Replication Agent instance releases all of its connections to the primary database, and drops its connection to the primary Replication Server (and RSSD, if connected).
  - The Sybase Replication Agent instance goes from *Replicating* state to *Admin* state.
- If the Sybase Replication Agent internal queues are full when the quiesce command is invoked, the quiesce processing may take a while to complete, and there may be a delay before the Sybase Replication Agent instance completes its transition to *Admin* state.
- Before moving the Sybase Replication Agent to the *Admin* state, the quiesce command waits until all data in the primary log has been read and sent to the Replication server.
- If the Sybase Replication Agent instance is in *Admin* state, the quiesce command returns an error.
- The quiesce command is valid only when the Sybase Replication Agent instance is in *Replicating* state.

Note The action of the suspend command is similar to that of the quiesce command, except that the suspend command stops Sybase Replication Agent processing immediately and flushes all data in the internal queues.

See also ra_status, resume, shutdown, suspend

ra_config

Description Returns help information for Sybase Replication Agent configuration parameters, or sets the value of a specified configuration parameter.

Syntax

ra_config [*param* [, *value*]]

Parameters

param

The name of a Sybase Replication Agent configuration parameter.

value

The value to be assigned to the configuration parameter specified in the *param* option. You can use the keyword default to set the specified parameter to its default value.

Examples

Example 1

```
ra_config use_rssd
```

This command returns the current value of the use_rssd configuration parameter.

Example 2

```
ra_config scan_sleep_max, 60
```

This command changes the value of the scan_sleep_max parameter to 60.

Usage

- If ra_config is invoked with no option, it returns a list of all Sybase Replication Agent configuration parameters.
- If ra_config is invoked with the *param* option, it returns information only for the specified configuration parameter.
- If ra_config is invoked with the *param* and *value* options, it changes the setting of the specified configuration parameter to the value specified in the *value* option.
- You can use the keyword default in place of the *value* option to reset a configuration parameter to its default value. For example:

  ```
  ra_config use_rssd, default
  ```

- The following information is returned for each configuration parameter:
  - Parameter name – the name of the parameter.
  - Parameter type – the datatype of the parameter's value (for example, string, numeric, or Boolean).
  - Current value – the value of the parameter in effect at the time ra_config is invoked.
  - Pending value – if different from the current value, the value to which the parameter was set by a previous invocation of the ra_config command, but which has not yet taken effect.

  - Default value – the value of the parameter when the Sybase Replication Agent instance configuration file is created.
  - Legal values – the values that are allowed for the parameter, for example, a range of numbers or a list of specific strings.
  - Category – refers to the Sybase Replication Agent component affected by the value of the parameter.
  - Restart – refers to parameters that require the Sybase Replication Agent instance to be shut down and restarted before a change in value takes effect.

- When ra_config is invoked with either no option, or only the *param* option, the command is valid when the Sybase Replication Agent instance is in 4either *Admin* or *Replicating* state.
- If ra_config is invoked when the Sybase Replication Agent instance is in *Replicating* state, with the *param* and *value* options for a parameter that can be changed only in *Admin* state, it returns an error.
- When ra_config is invoked with the *param* and *value* options, the command is always valid when the Sybase Replication Agent instance is in *Admin* state.
- When pdb_setrepproc is invoked to mark a procedure or procedures, a replication definition is created at Replication Server for each procedure that gets marked for replication, if this property is set to *true*.
- When pdb_setreptable is invoked to unmark a table or tables, the replication definition is dropped at Replication Server for each table that gets unmarked for replication, if this property is set to *true*.
- When pdb_setrepproc is invoked to unmark a procedure or procedures, a replication definition is dropped at Replication Server for each procedure that gets unmarked for replication if this property is set to *true*.
- When pdb_xlog is initialized and table auto marking is enabled, a replication definition is created for each table that is marked for replication if this property is set to *true*.
- Replication definition names for tables always begin with the prefix "*ra$*," followed by a unique alphanumeric identifier (maximum of 8 characters), and ending with a table or object name. For example, for a replicate name of "My Table," the resulting *repdef* name is "ra$0x7952_mytable." For an especially long replicate name of "mytable890123456789012345678 90" (30 characters), the resulting *repdef* name is "ra$0x7952_mytable8901234567890" (30 characters maximum).

- Replication definition names for procedures are the same name as the procedure.
- See Chapter 2, “Configuration Parameters,” for more information.

See also ra_help, ra_set_login

ra_date

Description Returns the current date and time from the Sybase Replication Agent instance.

Syntax ra_date

Usage

- When ra_date is invoked, it returns the current date and time from the Sybase Replication Agent instance in the form of a Sybase datetime datatype, as follows:

```
Current RA Date
 ------------------------------
     Apr 30 2005 12:09:47.310
(1 row affected)
```

- The ra_date command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also pdb_date, ra_config

ra_devicepath

Description **Note** This command is available only for Oracle.

Changes the disk device path for a log device recorded in the RASD.

Syntax ra_devicepath *device*, *dev_path*

Parameters

device
The device ID is the Oracle redo log “Group number.”

dev_path
The path that points to the disk log device for the device specified in the *device* option.

Examples Example 1

```
ra_devicepath 3,
d:\software\oracle\devices\redo001.log
```

This command specifies the disk device path to the log device ID "3" as:

```
d:\software\oracle\devices\redo001.log
```

Usage

- When ra_devicepath is invoked, Sybase Replication Agent records the specified disk device path for the specified log device in the RASD.
- To get information about log devices stored in the RASD, use the ra_helpdevice command.
- If you invoke ra_updatedevices after you set a device path using ra_devicepath, you must use ra_devicepath again to re-set the path if you need to alter the default path for a log device. The default device path is the device path returned by the primary data server.
- If you invoke ra_devicepath when the Sybase Replication Agent instance is in *Replicating* state, it returns an error.
- The ra_devicepath command is valid only when the Sybase Replication Agent instance is in *Admin* state.

See also

ra_helpdevice, pdb_xlog, ra_updatedevices

ra_dump

Description

Emulates the Replication Server rs_dumpdb and rs_dumptran system functions.

Syntax

ra_dump [database|transaction,] *dbname*, *dump_label*

Parameters

database
: A keyword that causes the primary Replication Server to apply the function string associated with the rs_dumpdb system function.

transaction
: A keyword that causes the primary Replication Server to apply the function string associated with the rs_dumptran system function.

dbname
: The name of the database to be dumped.

dump_label
: A varchar(30) value that contains information to identify the database dump.

Usage

- When ra_dump is invoked, Sybase Replication Agent places a dump marker in the Sybase Replication Agent transaction log to facilitate coordinated dumps.
- The ra_dump command returns an error message if the transaction log does not exist.
- The ra_dump command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.
- For more information about the Replication Server rs_dumpdb and rs_dumptran system functions, refer to the Sybase Replication Agent *Administration Guide* and Sybase Replication Agent *Reference Manual.*

See also ra_config, ra_migrate

ra_help

Description Returns help information for Sybase Replication Agent commands.

Syntax ra_help [*command*]

Parameters *command*

The name of a Sybase Replication Agent command for which you want to view help information.

Examples **Example 1**

```
ra_help
```

This command returns help for all Sybase Replication Agent commands.

Example 2

```
ra_help pdb_gen_id
```

This command returns help for the pdb_gen_id command.

Usage

- If ra_help is invoked with no option, it returns help information for all Sybase Replication Agent commands.
- If ra_help is invoked with the *command* option, it returns help information only for the specified command.
- The ra_help command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_config

ra_helparticle

Description

Note This command is available only for Oracle.

Returns information about primary database articles from the RASD.

Syntax

ra_helparticle [*article*, [*version*]]

Parameters

article

The name or object ID of an article (table or procedure) in the primary database. Article names can be qualified with an owner name in the following form:

```
owner.article
```

Owner qualification of article names is optional.

version

A hexadecimal locator value that identifies the version of the article specified in the *article* option.

Examples

Example 1

```
ra_helparticle
```

This command returns information about all versions of all articles in the RASD.

Example 2

```
ra_helparticle table1
```

This command returns information about the current version of the article named “table1” in the RASD.

Example 3

```
ra_helparticle table1,
0000000000210a40000333400070000333400699940000d413c50000000000
```

This command returns information about version 0000000000210a40000333400070000333400699940000d413c50000000000 of the article named “table1” in the RASD.

Usage

- The ra_helparticle command returns the following information for articles (tables and procedures):
  - Article object ID
  - Primary database name
  - Article owner name or alias

- Article name
- Article type (table or procedure)
- Article status (Current, Archived, or Dropped)
- Article version number

All information except the article type, article status, and article version number are the values returned by the primary database when the Sybase Replication Agent is initialized with the ra_init command.

- If ra_helparticle is invoked with no option, it returns information for all versions of all articles (tables and procedures) in the RASD.
- If ra_helparticle is invoked with the *article* option, it returns information only for the current version of the specified article in the RASD.
- If ra_helparticle is invoked with the *article* and *version* options, it returns information only for the specified version of the specified article in the RASD.
- The ra_helparticle command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_helpdb, ra_helpfield, ra_helplocator, ra_helpuser

ra_helpdb

Description Returns information about the primary database from the RASD.

Syntax ra_helpdb

Usage

- When ra_helpdb is invoked, it returns the following information about the primary database:
  - Database object ID
  - Database name

  The database ID and database name are the values returned by the primary database when the Sybase Replication Agent is initialized with the ra_init command.
- The ra_helpdb command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

- The ra_helpdb command is valid only after the RASD has been initialized, that is, only after you have executed ra_init.

See also

ra_devicepath, ra_helparticle, ra_helpdevice, ra_helpfield, ra_helplocator, ra_helpuser, ra_updatedevices

ra_helpdevice

Description

Note This command is available only for Oracle.

Returns information about primary database log devices from the RASD log device repository.

Syntax

ra_helpdevice [*device*]

Parameters

device

The device ID of the primary database log device.

Examples

Example 1

```
ra_helpdevice
```

This command returns information about all primary database log devices recorded in the log device repository.

Example 2

```
ra_helpdevice 1
```

This command returns information about the primary database log device ID "1" in the log device repository.

Usage

- The ra_helpdevice command returns the following information for each primary database log device recorded in the RASD:
  - Device ID – the log device ID defined by the primary data server.

    Note For Oracle, the ID is the value of the Oracle Redo Log Group to which this file belongs.

  - Database name – the name of the primary database associated with the log device.
  - Device name – the logical name of the log device defined by the primary data server.

  - Server device path – the path to a multiplexed version of the log device.
  - Disk device path – the path to the log device (at the standby site).
  - Disk device status – the current status of the server device path (ACCESSIBLE, NOT_VALID, or OPEN).
- The log device ID, primary database name, log device name, and server log device path are values returned by the primary data server when the Sybase Replication Agent is initialized with the pdb_xlog init command, or when the log device repository is updated with the ra_updatedevices command.

Note The ra_helpdevice command does *not* return information about software devices created with the Adaptive Server disk command. See the Adaptive Server documentation for more information.

- The disk device path is the current value recorded in the RASD. Sybase Replication Agent uses the disk device path recorded in its RASD to find each log device.

  For each log device recorded in the RASD, you can set or change the disk device path with the ra_devicepath command.

  If you do not specify a disk device path (using ra_devicepath), the value recorded for the disk device path is DEFAULT, and Sybase Replication Agent uses the value recorded for the server device path to find the log device.
- The disk device status is updated by the Log Reader component each time you invoke the ra_helpdevice command.
- If ra_helpdevice is invoked with no option, it returns information for all log devices recorded in the RASD log device repository.
- If ra_helpdevice is invoked with the *device* option, it returns information only for the specified log device.
- The ra_helpdevice command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_devicepath, ra_helpdb, ra_updatedevices

ra_helpfield

Description

Note This command is available only for Oracle.

Returns information about primary database fields (columns in tables, or input parameters in stored procedures) from the RASD.

Syntax

ra_helpfield *article* [, *version* [, *field*]]

Parameters

article

The name or object ID of an article (table or procedure) in the primary database. Article names can be qualified with an owner name in the following form:

```
owner.article
```

Owner qualification of article names is optional.

version

A hexadecimal locator value that identifies the version of the specified article.

field

The name or object ID of a field (column or input parameter) in the specified article.

Examples

Example 1

```
ra_helpfield authors
```

This command returns information about all fields in the current version of the article named authors in the RASD.

Example 2

```
ra_helpfield authors,
00000000000210a400003334000700003334000699940000d413c50000000000
```

This command returns information about all fields in version 00000000000210a400003334000700003334000699940000d413c50000000000 of the article named “authors” in the RASD.

Example 3

```
ra_helpfield authors,
00000000000210a400003334000700003334000699940000d413c50000000000, au_fname
```

This command returns information about the field named au_fname in version 00000000000210a400003334000700003334000699940000d413c50000000000 of the article named “authors” in the RASD.

Usage

- The ra_helpfield command returns the following information for fields:
  - Field (column or input parameter) object ID
  - Field name
  - Field type ID
  - Field datatype (with precision, length, and scale)
  - Field NULL mode
  - Field IDENTITY status
  - Field primary key status

  All field information items are the values returned by the primary database when the Sybase Replication Agent is initialized with the pdb_xlog init command.
- If ra_helpfield is invoked with the *article* option, it returns information for all fields in the current version of the specified article in the RASD.
- If ra_helpfield is invoked with the *article* and *version* options, it returns information for all fields in the specified version of the specified article in the RASD.
- If ra_helpfield is invoked with the *article*, *version*, and *field* options, it returns information for the specified field in the specified version of the specified article in the RASD.
- The ra_helpfield command is valid when the Sybase Replication Agent is in either *Admin* or *Replicating* state.
- No results are returned by this command if the RASD has not yet been initialized with the pdb_xlog init command.

See also

ra_config, ra_help, ra_helparticle, ra_helpdb, ra_helpdevice, ra_helplocator, ra_helpuser

ra_helplocator

Description

Note This command is available only for Oracle.

Returns information about fields in the LTM Locator value.

Syntax

ra_helplocator [*locator_value*]

Parameters

locator_value
The hexadecimal string value of an LTM Locator.

Examples

Example 1

```
ra_helplocator
```

This command returns information about fields in the current LTM Locator value.

Example 2

```
ra_helplocator locator_value
```

This command returns information about fields in the specified LTM Locator value.

Usage

- The ra_helplocator command returns the following information about the LTM Locator value:
  - Locator field names
  - Locator field hexadecimal values
  - Locator field decimal values
- If ra_helplocator is invoked with no option, it returns information about fields in the current LTM Locator value.
- If ra_helplocator is invoked with the *locator_value* option, it returns information about fields in the specified LTM Locator value.
- The ra_helplocator command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

ra_config, ra_help, ra_locator

See the Replication Agent *Primary Database Guide* for more information about locator fields and contents.

ra_helpuser

Description

Note This command is available only for Oracle.

Returns information about primary database users from the RASD.

Syntax

ra_helpuser [*user* [, *version*]]

Parameters

user
The name or user ID of a user in the primary database.

version
The version number of the database user in the RASD.

Examples

Example 1

```
ra_helpuser
```

This command returns information about all users in the RASD.

Example 2

```
ra_helpuser bob
```

This command returns information about all versions of the database user named “bob” in the RASD.

Example 3

```
ra_helpuser bob,
00000000000210a400003334000700003334000699940000d413c50000000000
```

This command returns information about version 00000000000210a400003334000700003334000699940000d413c50000000000 of the database user named “bob” in the RASD.

Usage

- The ra_helpuser command returns the following information about primary database users:
  - User ID
  - User name
  - User status (Current, Archived, or Dropped)
  - Primary database version (locator value)

  The user ID and user name are the values returned by the primary database when the Sybase Replication Agent is initialized with the ra_init command.
- If ra_helpuser is invoked with no option, it returns information about all users in the RASD.
- If ra_helpuser is invoked with the *user* option, it returns information about the specified user in all versions of the primary database in the RASD.
- If ra_helpuser is invoked with the *user* and *version* options, it returns information about the specified user in the specified version of the primary database in the RASD.

- The ra_helpuser command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.
- No results are returned by this command if the RASD has not been initialized by the ra_init command.

See also

ra_config, ra_help, ra_helparticle, ra_helpdb, ra_helpdevice, ra_helpfield, ra_helplocator

ra_locator

Description

Returns the current value of the LTM Locator maintained by the Sybase Replication Agent, requests an LTM Locator value from the primary Replication Server, or sets the value of the LTM Locator maintained by the Sybase Replication Agent to zero.

Syntax

ra_locator [update|zero]

Parameters

update
: The optional keyword to request a new LTM Locator value from the primary Replication Server.

zero
: The optional keyword to set the value of the LTM Locator stored in the Sybase Replication Agent transaction log to zero.

Examples

Example 1

```
ra_locator
```

This command returns the current value of the LTM Locator maintained by the Sybase Replication Agent, as shown:

```
Locator
 ------------------------------------------------
 0000000052000000000000527FFFFFFFFFFFFFF0022FB3B
 (1 row affected)
```

Example 2

```
ra_locator update
```

This command requests a new LTM Locator value from the primary Replication Server.

Example 3

```
ra_locator zero
```

This command sets the value of the LTM Locator maintained by the Sybase Replication Agent to all zeros.

Usage

- When ra_locator is invoked with no option, it returns the current value of the LTM Locator maintained by the Sybase Replication Agent instance. The Sybase Replication Agent stores the value of the LTM Locator in a table in the primary database; for Oracle this value is stored in the RASD.

  Note The value of the LTM Locator that is maintained by the Sybase Replication Agent is also known as the *origin queue ID*.

- When ra_locator is invoked with the update keyword, it requests a new LTM Locator value from the primary Replication Server, and the Sybase Replication Agent saves the value.

  Note When the ra_locator command is invoked with the update keyword, the change takes effect only if the Sybase Replication Agent instance is in *Replicating* state.

- When ra_locator is invoked with the zero keyword, it sets the value of the LTM Locator maintained by the Sybase Replication Agent to zero.
- The LTM Locator contains information that the Sybase Replication Agent uses to determine where to start reading the transaction log.

  Upon start-up or recovery from a connection failure, the Sybase Replication Agent automatically requests an LTM Locator value from the primary Replication Server.

  - If the value of the LTM Locator returned from the primary Replication Server is zero, then Sybase Replication Agent uses the LTM Locator value stored in the transaction log system table.
  - If the value of the LTM Locator in the transaction log system table is zero, then Sybase Replication Agent starts reading the transaction log from either the current beginning of the log, or from the end of the log for UDB.

- For more information about the format of the origin queue ID, see the chapter for your specific primary data server in the Sybase Replication Agent *Primary Database Guide*.
- If the Sybase Replication Agent transaction log does not exist, the ra_locator command returns an error message.

- The ra_locator command with the zero keyword is valid only when the Sybase Replication Agent instance is in *Admin* state.
- Without the zero keyword, the ra_locator command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also pdb_gen_id, pdb_truncate_xlog

ra_maintid

Description Returns the login name of the primary database Maintenance User currently stored in the Sybase Replication Agent transaction log.

For Oracle, the value is stored in the RASD.

Syntax ra_maintid

Usage

- Replication Server requires a Maintenance User login name for each database connection. The Maintenance User login name for a database connection is specified with the Replication Server create connection or alter connection command.

  When the primary database Maintenance User login name is changed in the Replication Server (using the alter connection command), Replication Server automatically sends the new Maintenance User login name to the Sybase Replication Agent, if the Sybase Replication Agent is in *Replicating* state.

  Each time the Sybase Replication Agent goes into *Replicating* state, it automatically retrieves the primary database Maintenance User login name from the primary Replication Server, and stores it in the Sybase Replication Agent transaction log.

- When ra_maintid is invoked, it returns the login name of the primary database Maintenance User stored in the Sybase Replication Agent transaction log, as follows:

```
Maintenance User
 ----------------
 SYS
(1 row affected)
```

- If ra_maintid is invoked when the Sybase Replication Agent is in *Replicating* state, it always returns the correct Maintenance User login name.

  If ra_maintid is invoked when the Sybase Replication Agent is in *Admin* state, it may not return the correct Maintenance User login name, because the Maintenance User login name could have changed in the Replication Server after the last time the Sybase Replication Agent retrieved the value and stored it.

- The filter_maint_userid configuration parameter is provided to support bidirectional replication, wherein the primary database also acts as a replicate database that has transactions applied to it by a Replication Server.

  If the value of the filter_maint_userid parameter is true, database operations applied by the Maintenance User are *not* replicated from the primary database. When it reads the transaction log, the Sybase Replication Agent Log Reader component filters out data-changing operations applied by the Maintenance User.

- If ra_maintid is invoked when the primary database connection is down, it returns an error.

- The ra_maintid command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also

ra_config, ra_statistics

ra_marker

Description

Emulates the Replication Server ra_marker system function, placing a marker object in the Oracle redo log.

Syntax

ra_marker *command_tag*

Parameters

command_tag

A varchar(255) value that contains information used for subscription materialization.

Examples

```
ra_marker 'activate subscription 309 0 with suspension'
```

This command places a marker object in the Sybase Replication Agent transaction log that invokes the Replication Server activate subscription command.

Usage

- When ra_marker is invoked, the Sybase Replication Agent inserts a row in a shadow table marked for replication. The replicated transaction is sent as a marker object to the primary Replication Server.
- The ra_marker command returns an error message if the Sybase Replication Agent transaction log does not exist.
- The ra_marker command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.
- For more information about the Replication Server ra_marker system function, refer to the Replication Server *Administration Guide* and Sybase Replication Agent *Reference Manual*.

See also ra_dump

ra_migrate

Description Performs any migration task (as necessary) between releases of Replication Agent.

Syntax ra_migrate

Parameters *None*

Usage

- Minimum migration for each release will at least update the xlog table with the latest build version.
- When you received a new drop of the Replication Agent, you must first run this command to update to the latest version of the Replication Agent.
- The ra_migrate command is valid when the Sybase Replication Agent instance is in *Admin* state.

ra_set_login

Description Sets the Sybase Replication Agent administrator login and password.

Syntax ra_set_login *username*, *password*

Parameters *username*

The login name of the Sybase Replication Agent administrator.

password
The password of the Sybase Replication Agent administrator.

Examples

```
ra_set_login bob3, bug3wag
```

This command sets the Sybase Replication Agent administrator login to "bob3" and the password to "bug3wag."

Usage

- The Replication Agent administrator login has permission to log in to the Sybase Replication Agent instance through the administration port.
- Only one Replication Agent administrator login name is valid at any time.
- Any change in the Replication Agent administrator login or password takes place immediately, and you must use the new login and password the next time you log in to the Sybase Replication Agent instance.
- The password specified for the administrator login is encrypted in the Sybase Replication Agent configuration file.
- The ra_set_login command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_config

ra_statistics

Description Returns performance-related statistics for Replication Agent components and the Java Virtual Machine (Java VM), or resets the statistics counters.

Note The statistics counters may vary by primary database.

Syntax ra_statistics [*component*|reset]

Parameters

component
The optional keyword that identifies a Replication Agent component or the Java VM. Valid *component* keywords are:

LR – Log Reader
LTI – Log Transfer Interface
LTM – Log Transfer Manager
VM – Java Virtual Machine

reset
The optional keyword that resets the statistics counters.

Examples

Example 1

```
ra_statistics
```

This command returns performance statistics for the Sybase Replication Agent instance and the Java VM.

Example 2

```
ra_statistics reset
```

This command resets the statistics counters for the Sybase Replication Agent instance.

Usage

- If you invoke ra_statistics with no option, it returns statistics for all Replication Agent components and the Java VM.
- If you invoke ra_statistics with a *component* option, it returns statistics for the Log Transfer Manager component, as well as the component (or Java VM) you specify.
- Table 1-2 lists the statistics returned for the Java VM.

Table 1-2: Java VM statistics

Component	Statistic	Description
VM	VM maximum memory	Maximum memory (in bytes) available to the Java VM
VM	VM total memory allocated	Total memory (in bytes) allocated to the Java VM at start-up
VM	VM free memory	Memory (in bytes) allocated but not used by the Java VM
VM	VM memory usage	Memory (in bytes) allocated and in use by the Java VM
VM	VM % max memory used	Percentage of the maximum memory available to the Java VM, currently in use by the Java VM

- Table 1-3 lists the statistics returned for the Log Transfer Manager component.

Table 1-3: Log Transfer Manager statistics

Component	Statistic	Description
LTM	Time statistics obtained	Day, date, and time ra_statistics was invoked and information returned
LTM	Time replication last started	Day, date, and time that *Replicating* state was entered

Component	Statistic	Description
LTM	Time statistics last reset	Day, date, and time that statistics counters were reset
LTM	Items held in Global LRUCache	Number of object references in the internal Least Recently Used cache

- Table 1-4 lists the statistics returned for the Log Reader component.

Table 1-4: Log Reader statistics for Oracle

Component	Statistic	Description
LR	Total operations scanned	Number of operations read from log devices since last reset
LR	Total operations processed	Number of operations read from log devices and passed to LTI since last reset
LR	Total operations skipped	Number of operations read from log devices and not processed for any reason since last reset
LR	Total maintenance user operations filtered	Number of Maintenance User operations read from log devices and skipped since last reset
LR	Avg operation processing time	Average Log Reader operation processing time (in milliseconds) since last reset
LR	Total transactions processed	Number of transactions read from log devices since last reset
LR	Total transactions skipped	Number of transactions read from log devices and not processed for any reason since last reset
LR	Total transactions opened	Number of begin transaction commands read from log devices since last reset
LR	Total transactions closed	Number of commit and rollback commands read from log devices since last reset
LR	Total transactions committed	Number of commit commands read from log devices since last reset
LR	Total transactions aborted (rolled back)	Number of rollback commands read from log devices since last reset
LR	Total system transactions skipped	Number of system transactions read from log devices and skipped since last reset
LR	Avg operations per transaction	Average number of operations in each transaction read from log devices since last reset

Component	Statistic	Description
LR	Current scan buffer size	Current size (in bytes) of the Log Reader scan buffer
LR	Current operation queue size	Current size (in bytes) of the Log Reader input queue
LR	Current session cache size	Current size (in bytes) of the session cache
LR	Log reposition point locator	Locator value of reposition point in log device
LR	Last processed operation locator	Locator value of most recently processed operation read from log devices
LR	Avg xlog operation wait time (ms)	Average time (in milliseconds) that Log Reader had to wait for each new operation to appear in the log since last reset
LR	Avg sender operation processing time (ms)	Average time (in milliseconds) that Log Reader sender took to process each operation since last reset
LR	Avg sender operation wait time (ms)	Average time (in milliseconds) that Log Reader sender had to wait to send each processed operation to the LTI input queue since last reset
LR	Avg ChangeSet send time (ms)	Average time (in milliseconds) that Log Reader sender took to send each processed operation to the LTI input queue since last reset
LR	Total sender operations processed	Number of operations that Log Reader sender processed since last reset
LR	Marked objects cache size	Current marked objects cache size

- Table 1-5 lists the Log Reader statistics for Microsoft SQL Server and UDB.

Table 1-5: Log Reader statistics for Microsoft SQL Server and UDB

Component	Statistic	Description
LR	Number of xlogs scanned	Number of operations read from log devices
LR	Average unprocessed operations per XLog scan	Average number of unprocessed operations for each XLog scan
LR	Average XLog scan time	Average XLog scan time for operations read from log devices

Component	Statistic	Description
LR	Number of operations replicated	Number of operations that were successfully replicated
LR	Number of transactions replicated	Number of transactions that were successfully replicated
LR	Number of XLog operations skipped (maint_user, unmarked tables)	Number of XLog operations that were skipped
LR	Average wait time on empty XLog	Average time that the XLog was not in use
LR	Average PDB Service Time/Operations	Average service and operations time for each database
LR	Operation Queue Size	The queue size used for the operations
LR	Operation Data Hash Size	The data hash size for the operations
LR	Number of transactions truncated	Number of transactions that were truncated

- Table 1-6 lists the statistics returned for the Log Transfer Interface component.

Table 1-6: Log Transfer Interface statistics

Component	Statistic	Description
LTI	Number of LTL commands sent	Total number of LTL commands sent to Replication Server since last reset
LTI	Avg LTL command size	Average size (in bytes) of each LTL command sent to Replication Server since last reset
LTI	Avg LTL commands/sec	Average number of LTL commands sent per second to Replication Server since last reset
LTI	Total bytes sent	Number of bytes sent to Replication Server since last reset
LTI	Avg Bytes/second during transmission	Average bytes per second sent over connection to Replication Server since last reset
LTI	Avg Replication Server turnaround time	Average time (in milliseconds) it takes Replication Server to acknowledge each LTL command buffer sent since last reset

Component	Statistic	Description
LTI	Avg data arrival time	Average time (in milliseconds) LTI waits between receiving change sets from Log Reader since last reset
LTI	Avg time to create distributes	Average time (in milliseconds) LTI takes to convert a change-set into LTL since last reset
LTI	Avg LTL buffer cache time	Average time (in milliseconds) it takes between placing the LTL commands into the LTL buffer to the time it is actually sent to the Replication Server
LTI	Avg LTL buffer size	Average size (in bytes) of each LTL buffer sent to Replication Server since last reset
LTI	Avg LTM buffer utilization (%)	Average utilization (in percentage of LTL buffer size) of each LTL buffer sent to Replication Server since last reset
LTI	Avg LTL commands/buffer	Average number of LTL commands per buffer sent to Replication Server since last reset
LTI	Input queue size	Current number of change sets in the LTI input queue
LTI	Output queue size	Current number of distributes in the LTI output queue
LTI	Last QID sent	Hex value of most recent origin queue ID sent to Replication Server
LTI	Last transaction sent	Hex value of most recent transaction ID sent to Replication Server

- Statistics counters are reset automatically each time the Sybase Replication Agent instance goes into *Replicating* state.
- If you invoke ra_statistics with the reset keyword, Sybase Replication Agent immediately resets all of the statistics, except the following:
  - Time statistics obtained (LTM)
  - Time replication last started (LTM)
  - Time statistics last reset (LTM)
  - Last QID sent (LTI)
  - Last transaction ID sent (LTI)

- All Java VM statistics

Note All Java VM statistics are refreshed each time you invoke ra_statistics.

- The ra_statistics command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_status

ra_status

Description Returns the current state of the Sybase Replication Agent instance.

Syntax ra_status

Usage

- When ra_status is invoked, it returns the current state of the Sybase Replication Agent instance, and a brief description of the current state, as follows:

```
State  Action
 ------ ----------------------------
 ADMIN  Waiting for operator command
(1 row affected)
```

Note If the first word in the description is "Transitioning," the Sybase Replication Agent instance is in transition between states. Some commands are not valid when the Sybase Replication Agent instance is in state transition.

- Sybase Replication Agent states are:
  - *Admin* – in this state, the Sybase Replication Agent instance is running, but no connections are up. You can change any configuration parameter when the Sybase Replication Agent instance is in *Admin* state.
  - *Replicating* – in this state, the Log Reader component is scanning the transaction log for operations to replicate from the primary database. If there are operations to be replicated, the Log Transfer Interface component is sending LTL commands to the Replication Server.

- *Replicating (Waiting at end of log)* – in this state, the Log Reader component has reached the end of the transaction log, the Sybase Replication Agent has finished processing all operations in the transaction log, and the Log Transfer Interface component has successfully sent LTL commands for all replicated operations to the Replication Server.

  If the primary database is not quiesced or is otherwise inactive, transactions could arrive in the log immediately after the state is returned, so even though the state is returned as *Replicating (Waiting at end of log)*, the Sybase Replication Agent could actually be in *Replicating* state and processing log records.

See the Replication Agent *Administration Guide* for more information about Sybase Replication Agent states.

- The ra_status command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also quiesce, ra_statistics, resume, shutdown, suspend

ra_truncatearticles

Description

Note This command is available only for Oracle.

Truncates older versions of primary database articles in the RASD.

Syntax ra_truncatearticles *locator*

Parameters *locator*

The log locator value (LTM Locator) that identifies the cutoff point for truncating older versions of articles from the system data repository.

Usage

- When ra_truncatearticles is invoked, it truncates all non-current versions of all primary database articles in the system data repository older than the version identified by the *locator* value.

  If the current (most recent) version of an article is older than the version identified by the *locator* value, it is not truncated.

- Most common DDL commands and stored procedures executed in the primary database (such as alter table) are recorded in the transaction log, and replicated to the standby database. When it processes those DDL transactions for replication, Sybase Replication Agent updates its RASD automatically, creating a new version of the affected primary database articles.

  Use ra_truncatearticles as part of a periodic maintenance procedure to prevent the RASD from growing indefinitely. See the Sybase Replication Agent *Administration Guide* for more information.

  Note Be sure to back up the RASD using rasd_backup before you truncate it.

- The ra_truncatearticles command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_truncateusers

ra_truncateusers

Description **Note** This command is available only for Oracle.

Truncates older versions of primary database users in the RASD.

Syntax ra_truncateusers *locator*

Parameters *locator*
The log locator value (LTM Locator) that identifies the cutoff point for truncating older versions of database users from the system data repository.

Usage
- When ra_truncateusers is invoked, it truncates all non-current versions of all primary database users in the system data repository older than the version identified by the *locator* value.

  If the current (most recent) version of a user is older than the version identified by the *locator* value, it is not truncated.

- The ra_truncateusers command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_truncatearticles

ra_updatedevices

Description

Note This command is available only for Oracle.

Updates information about primary database log devices in the RASD.

Syntax

ra_updatedevices

Usage

- When ra_updatedevices is invoked, Sybase Replication Agent does the following:
  - Refreshes the archive log information
  - Deletes all of the data in its log device repository.

    Note If the device location is set, it is not overwritten.

  - Queries the primary database for information about all of its log devices
  - Re-populates the log device repository in the RASD with current information about primary database log devices returned by the primary database
- If any log device associated with the primary database is added, dropped, extended, or moved at the primary data server, you must:
  - Stop replication (using quiesce or suspend) to put the Sybase Replication Agent instance in *Admin* state
  - Invoke ra_updatedevices to update the log device repository in the RASD

  See the Sybase Replication Agent *Administration Guide* for more information.

  Note The primary database need not be quiesced when you update the log device repository.

- If the primary data server writes to a new (or altered) log device before you update the log device repository, the Sybase Replication Agent instance will stop replication processing and go to *Admin* state.

  Sybase recommends that you coordinate all log device changes at the primary database with updating the Sybase Replication Agent log device repository.

- Because Sybase Replication Agent re-creates the entire log device repository when you invoke ra_updatedevices, any log device path that you modified previously (using ra_devicepath) is overwritten with the current log device information from the primary database.

  For example:

  ```
  ID=1 serverpath=/dev1=/dev1a
  ```

  becomes the following when you change the `server path` to "`dev44`.":

  ```
  ID=1 serverpath=/dev44=/dev1a
  ```

  Note If you need to alter the "default" path for a log device (that is, the log device path returned by the primary database), you must use the ra_devicepath command *after* you invoke ra_updatedevices.

- For each log device recorded in the RASD, you can set or change the disk device path with the ra_devicepath command.

  If you do not specify a disk device path (using ra_devicepath), the value recorded for the disk device path is DEFAULT, and Sybase Replication Agent uses the value recorded for the server device path to find the log device.

- The ra_updatedevices command is valid only when the Sybase Replication Agent instance is in *Admin* state.

See also ra_devicepath, ra_helpdevice

ra_version

Description Returns the version of the Sybase Replication Agent instance, the host operating system version, and the JRE version.

Syntax ra_version

Usage

- When ra_version is invoked, it returns the Sybase Replication Agent version string in a row, as follows:

  ```
  Sybase Replication Agent for Unix &
  Windows/15.0.0.5400/P/generic/JDK
  1.4.2/main/5400/VM: Sun Microsystems Inc.
  1.4.2_12/OPT/Mon Oct 30 23:34:07 MST 2006
  ```

- The ra_version command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also pdb_version, ra_status, ra_version_all

ra_version_all

Description Returns the name, type, and version of the Sybase Replication Agent instance, and version information for the primary data server, primary Replication Server, and communications drivers.

Syntax ra_version_all

Usage

- When ra_version_all is invoked, it returns the following information:

```
Component              Version
-------------------------------------------------------------------------

Instance:              rao2 - Oracle

RepAgent:              Sybase Replication Agent for Unix &
                       Windows/15.0.0.5400/P/generic/JDK
                       1.4.2/main/5400/VM: Sun Microsystems Inc.
                       1.4.2_12/OPT/Mon Oct 30 23:34:07 MST 2006

JRE:                   Sun Microsystems Inc. Java(TM) 2 Runtime
                       Environment, Standard Edition/1.4.2_12-b03/Windows
                       XP 5.1/x86/32

RASD:                  Sybase Adaptive Server
                       Anywhere/9.0.2.3302/WindowsXP

Primary Data Server:Oracle Oracle9i Enterprise Edition Release
                       9.2.0.1.0 - Production With the Partitioning, OLAP
                       and ORACLE Data Mining options JServer
                       Release 9.2.0.1.0 - Production

PDS JDBC Driver:       Oracle JDBC driver 10.1.0.2.0

RepServer:             Replication Server/12.6/EBF 12739
                       ESD#5/NT (IX86)/Windows 2000
                       /1/OPT/Fri Jul 01 14:01:20 2005

RSSD:                  Sybase Adaptive ServerAnywhere/8.0.2.4322/WindowsXP


Sybase JDBC Driver:jConnect (TM) for JDBC(TM)/6.05(Build
                       25828)/P/EBF13044/JDK14/Fri Sep 30 1:05:16 2005
```

- The ra_version_all command returns information about the primary data server and communications driver, and the primary Replication Server and communications driver, only if the connections to those servers are properly configured in the Sybase Replication Agent.

  If a connection is not configured, or it is configured incorrectly, ra_version_all returns `<unknown>` for that connection.

  Note The ra_version_all command always returns information about the Sybase Replication Agent version and instance.

- The ra_version_all command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also pdb_version, ra_status, ra_version

rasd_backup

Description

Note This command is available only for Oracle.

Backs up the Replication Agent System Database (RASD).

Syntax rasd_backup

Usage

- When rasd_backup is invoked, it starts the database backup process for the RASD.

  Note Sybase recommends that you always back up the RASD before you truncate using ra_truncatearticles or ra_truncateusers.

- Sybase Replication Agent places RASD backup files in the directory identified by the rasd_backup_dir configuration parameter.

  When you create a Sybase Replication Agent instance, a RASD backup directory is created automatically as part of the instance directory structure. The default value of the rasd_backup_dir parameter points to that directory.

- The rasd_backup command is valid when the Sybase Replication Agent instance is in *Admin* or *Replicating* state.

See also rasd_restore, ra_truncatearticles, ra_truncateusers

rasd_restore

Description

Note This command is available only for Oracle.

Restores the RASD.

Syntax

rasd_restore

Usage

- When rasd_restore is invoked, it starts the restore process for the RASD.
- Sybase Replication Agent looks for the most recent RASD backup files in the directory identified by the rasd_backup_dir configuration parameter.

  When you create a Sybase Replication Agent instance, a RASD backup directory is created automatically as part of the instance directory structure. The default value of the rasd_backup_dir parameter points to that directory.
- If you invoke rasd_restore when the Sybase Replication Agent instance is in *Replicating* state, it returns an error.
- The rasd_restore command is valid only when the Sybase Replication Agent instance is in *Admin* state.

See also

rasd_backup

resume

Description

Starts replication processing in the Sybase Replication Agent instance.

Syntax

resume

Usage

- When resume is invoked, the Sybase Replication Agent instance attempts to go to *Replicating* state and start replication operations, as follows:
  - Sybase Replication Agent attempts to open network connections to the primary database, primary Replication Server, and RSSD.

    If it fails to establish a connection, the Sybase Replication Agent logs a warning message in its system log, and it attempts to retry the connection, based on its configuration parameters for the connection.
  - If the Sybase Replication Agent cannot establish a connection to the primary database after exhausting its configured retry attempts, it aborts all subsequent resume processing, returns to *Admin* state, and logs the error.

    - Sybase Replication Agent requests the current LTM Locator value from the primary Replication Server, and it stores the value in the Sybase Replication Agent transaction log.
    - The Log Reader component begins scanning the transaction log, looking for operations to be replicated. Log Reader begins scanning the log at the point identified by the LTM Locator value.
    - When it finds transactions to replicate, Log Reader passes them (as change-set data) to the input queue of the Log Transfer Interface component.
    - The Log Transfer Interface component reads the change-set data from its input queue, generates LTL commands, and places the LTL commands in its output queue for transmission to the Replication Server.
- If any start-up operation fails, the Sybase Replication Agent instance returns to *Admin* state, and it logs the error.
- If the resume command is successful, the Sybase Replication Agent instance goes to *Replicating* state. To determine the current state of the Sybase Replication Agent, use the ra_status command.
- The resume command returns an error under any of the following conditions:
    - The Sybase Replication Agent instance is already in *Replicating* state.
    - The Sybase Replication Agent transaction log does not exist (DB2 UDB, or Microsoft SQL Server.)
    - The system data repository in the RASD does not exist or is not initialized (Oracle).
    - The Sybase Replication Agent connection configuration parameters are not set correctly, or it fails otherwise to connect with the primary database or the primary Replication Server.
    - The database connection for the primary database is not defined correctly in the primary Replication Server.
- If the resume command is successful, the Sybase Replication Agent instance goes into *Replicating* state.
- The resume command is valid only when the Sybase Replication Agent instance is in *Admin* state.

See also quiesce, ra_status, shutdown, suspend

rs_create_repdef

Description

Note This command is available only for Oracle.

Creates a replication definition at Replication Server for a marked table and procedure, or for all marked tables and procedures.

Note Replication Agent 15.0 is pre-configured to match replication definition datatypes available in Replication Server 15.0. If replication definitions are to be generated against an earlier version of Replication Server, this configuration needs to be changed. Contact Sybase Technical Support for assistance in making this adjustment.

Syntax

```
rs_create_repdef [all, TABLE_NAME]
```

Parameters

all
: A replication definition is created for all tables and procedures that are marked for replication.

TABLE_NAME
: A replication definition is created for that table or procedure.

Usage

- When rs_create_repdef is invoked and the parameter "all" or "ALL" is entered, a replication definition is created for all tables or procedures that are marked for replication.
- When rs_create_repdef is invoked and the name of a table or procedure that is marked for replication is entered, a replication definition is created for that table or procedure.
- For each table or procedure for which a replication definition create is attempted, a result set is returned. The result set contains the replication definition name and status of the create. If the replication definition was created, the status will be "created." If an error occurred, an error message from Replication Server will be returned.
- The character case of the object names in the replication definition will be set according to the ltl_character_case setting.

- Replication definition names for tables always begin with the prefix "*ra$*," followed by a unique alphanumeric identifier (maximum of 8 characters), and ending with a table or object name. For example, for a replicate name of "My Table," the resulting repdef name is "ra$0x7952_mytable." For an especially long replicate name of "mytable89012345678901234567890" (30 characters), the resulting repdef name is "ra$0x7952_mytable8901234567890" (30 characters maximum).
- Replication definition names for procedures are the same name as the procedure.

See also rs_drop_repdef

rs_drop_repdef

Description A replication definition at the configured Replication Server for a table and procedure is dropped.

Syntax rs_drop_repdef [TABLE_NAME]

Parameters TABLE_NAME
A replication definition is dropped for that table or procedure.

Usage
- When rs_drop_repdef is invoked, a replication definition for that table is dropped at the Replication Server.
- When rs_drop_repdef is invoked and the name of a table or procedure that is marked for replication is entered, a replication definition is created for that table or procedure.
- For each table or procedure for which a replication definition create is dropped, a result set is returned. The result set contains the table name and status of the create. If the replication definition was created, the status will be "dropped." If an error occurred, an error message from Replication Server will be returned.
- The character case of the object names in the replication definition will be set according to the ltl_character_case setting.

- Replication definition names for tables always begin with the prefix "*ra$*," followed by a unique alphanumeric identifier (maximum of 8 characters), and ending with a table or object name. For example, for a replicate name of "My Table," the resulting *repdef* name is "ra$0x7952_mytable." For an especially long replicate name of "mytable890123456789012345678 90" (30 characters), the resulting *repdef* name is "ra$0x7952_mytable8901234567890" (30 characters maximum).

See also rs_create_repdef

shutdown

Description Shuts down the Sybase Replication Agent instance, terminating its process.

Syntax shutdown [immediate]

Parameters immediate

The optional keyword that shuts down the Sybase Replication Agent instance immediately.

Usage

- When shutdown is invoked with no option, the Sybase Replication Agent starts a normal (graceful) shutdown.

  In a normal shutdown , the Sybase Replication Agent first quiesces, and then the process terminates. See quiesce on page 57 for more information about quiescing the Sybase Replication Agent.

- When shutdown is invoked with the immediate keyword, the Sybase Replication Agent starts an immediate shutdown.

  In an immediate shutdown, the Sybase Replication Agent:

  - Stops all of its replication processing, without regard to transactions in process or in transit
  - Drops all of its connections
  - Terminates the application process

- The shutdown command with the immediate keyword is valid at any time, when the Sybase Replication Agent instance is in any state, including transition between states.
- The shutdown command with no keyword (normal shutdown) is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state, but not in state transition.

See also quiesce, ra_status, resume, suspend

suspend

Description Stops all current replication processing and puts the Sybase Replication Agent instance into *Admin* state.

Syntax suspend

Usage

- When suspend is invoked, it stops all current replication processing in the Sybase Replication Agent instance.
  - The Log Reader component stops scanning the transaction log immediately, and the Log Transfer Interface component stops sending LTL to the Replication Server immediately.
  - Any data in the Sybase Replication Agent internal queues (input and output queues of the Log Reader and Log Transfer Interface components) is flushed without further processing.
  - The Sybase Replication Agent instance immediately releases all of its connections to the primary database, and drops its connection to the primary Replication Server (and RSSD, if connected).
  - The Sybase Replication Agent instance goes from *Replicating* state to *Admin* state.

  Note The action of the quiesce command is similar to that of the suspend command, except that quiesce allows pending transactions in the Sybase Replication Agent internal queues to be processed first, before putting the Sybase Replication Agent instance in *Admin* state.

- If the Sybase Replication Agent instance is in *Admin* state, the suspend command returns an error.
- The suspend command is valid only when the Sybase Replication Agent instance is in *Replicating* state.

See also quiesce, ra_status, resume, shutdown

test_connection

Description

Tests Sybase Replication Agent connection configurations and network connectivity.

Syntax

test_connection [*conn_name*]

Parameters

conn_name

The keyword for a Sybase Replication Agent connection to be tested. Valid keywords are:

- PDS – primary data server
- RS – primary Replication Server (and RSSD, if so configured)

Note If the value of the use_rssd configuration parameter is true, the test_connection command tests Sybase Replication Agent connectivity to the RSSD when it tests connectivity to the Replication Server. If the value of the use_rssd configuration parameter is false, the test_connection command does *not* test Sybase Replication Agent connectivity to the RSSD.

Examples

Example 1

```
test_connection
```

This command tests all Sybase Replication Agent connections, including the primary data server connection, the primary Replication Server connection, and the RSSD connection (if so configured).

Example 2

```
test_connection PDS
```

This command tests only the Sybase Replication Agent connection for the primary data server.

Usage

- When test_connection is invoked with no option, Sybase Replication Agent tests all of its connections by attempting to log in to the corresponding server for each connection, using the connection parameters stored in its configuration file.
- When test_connection is invoked with either the RS or PDS keyword, Sybase Replication Agent tests the specified connection.
- The test_connection command verifies both network connectivity and the following Sybase Replication Agent connection configuration parameters for the primary database:

- connection type (connectivity driver and protocol) – pds_connection_type
- database name – pds_database_name
- data server name – pds_server_name
- Data source name (ODBC drivers only) – pds_datasource_name
- host machine name – pds_host_name
- port number – pds_port_number
- user login access – pds_password and pds_username

Note The test_connection command does *not* validate Sybase Replication Agent user login permissions in the primary database. It verifies only that the user login and password specified in the pds_username and pds_password parameters can log in to the primary data server.

- The test_connection command verifies both network connectivity and the following Sybase Replication Agent connection configuration parameters for the primary Replication Server (and RSSD, if so configured):
  - Character set – rs_charset (and rssd_charset)
  - Database name – rssd_database_name (RSSD only)
  - Data server name – rssd_server_name (RSSD only)
  - Replication Server data source (as specified in the Replication Server primary database connection) – rs_source_db and rs_source_ds (Replication Server only)
  - Host machine name – rs_host_name (and rssd_host_name)
  - Network packet size – rs_packet_size (Replication Server only)
  - Port number – rs_port_number (and rssd_port_number)

- User login access – rs_password and rs_username (rssd_password and rssd_username)

Note The test_connection command verifies that the Sybase Replication Agent user login (specified in the rs_username and rs_password parameters) has connect source permission in the primary Replication Server.

- The test_connection command returns the connection type and its status, as follows:

```
Type Connection
---- ----------
PDS  succeeded
RS   succeeded

(2 rows affected)
```

  If the connection status is failed, it indicates one of the following:

  - The Sybase Replication Agent connection configuration parameters are not set correctly.
  - A network failure or communication error prevents the connection.
  - The server associated with the connection is down.

- If the connection status is failed, check the Sybase Replication Agent system log to determine the cause of the failure.

Note You may also need to check the system log of the server associated with the connection to determine the cause of the failure.

- See the Sybase Replication Agent *Administration Guide* for information about setting up Sybase Replication Agent connection configuration parameters.
- See Chapter 2, “Configuration Parameters,” for information about specific connection configuration parameters.
- The test_connection command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also ra_config, ra_statistics, ra_status

trace

Description

Returns current trace flag settings, or changes trace flag settings for the Sybase Replication Agent instance.

Syntax

trace [*flag*|all, *switch*]

Parameters

flag

The name of the trace flag to change the setting for.

all

A keyword that allows you to apply a switch value to all of the trace flags at once.

switch

A Boolean (true or false) value that enables or disables tracing for the trace point identified in the *flag* option.

Usage

- The trace command is intended for use by Sybase Technical Support engineers when troubleshooting Sybase Replication Agent.
- When trace is invoked with no option, it returns the current settings for all Sybase Replication Agent trace flags.
- When trace is invoked with the *flag* and *switch* options, it changes the setting of the trace flag identified, and it returns the current (new) setting for the trace flag.
- When trace is invoked with the all keyword and a *switch* option, it sets all Sybase Replication Agent trace flags to the value specified in the *switch* option, and it returns the current (new) setting for all of the trace flags.
- Changes made with the trace command take effect immediately.
- When a trace flag is set to true, tracing is enabled for the trace points identified by the flag. When set to false, tracing is disabled for the trace points.
- Output from all trace points (except *LTITRACELTL*) is sent to the Sybase Replication Agent system log file. Use the log_system_name command to find the name and path of the Replication Agent system log file.

- Output from the *LTITRACELTL* trace point is sent to a separate trace output file named *LTITRACELTL.log*. To view the contents of the *LTITRACELTL.log* file, your file viewer must be capable of handling very long lines.

  Note The *LTITRACELTL.log* file contains a human-readable representation of the LTL, not the actual LTL commands sent to the primary Replication Server.

- Table 1-7 lists Sybase Replication Agent trace flags:

Table 1-7: Sybase Replication Agent trace flags

Trace flag	Description
BMGRTRACE	When set to “true,” this flag enables Bean Management event tracing.
CACHETRC	When set to "true," this flag enables tracing of internal cache events.
DBCONTEXT	When set to “true,” this flag turns on tracing of database context events.
LATRC	When set to “true,” this flag traces general Log Administrator operations.
LATRCSQL	When set to “true,” this flag traces SQL conversations between Log Administrator and the primary database.
LICTRACE	When set to “true,” this flag traces feature license check-in/checkout events.
LOGREADTRC	When set to “true”, turns on trace of database log reading.
LRTRACE	When set to “true,” this flag traces general execution of the Log Reader component.
LTITRACE	When set to “true,” this trace flag enables tracing operations of the Log Transfer Interface component.
LTITRACELTL	When set to “true,” this trace flag enables LTL statement tracing in the *LTITRACELTL.log* file.
LTMCI	When set to “true,” causes tracing of LTM component interface invocations and LTM invocations of other components' interfaces.
LTLFMTTRC	When set to "true" this trace flag enables tracing of the LTL formatter.
LTMHL	When set to “true,” causes highlights in the LTM execution path to be noted.
LTMSC	When set to “true,” causes tracing of all Sybase Replication Agent state changes.
RACONTRC	When set to “true,” causes tracing of connection and query execution.
RACONTRCSQL	When set to “true,” causes tracing of SQL statements to be executed.
RASDTRC	(For Oracle only) When set to "true", turns on tracing of Sybase Replication Agent System Data Repository events.
RATRACE	When set to “true,” causes tracing of Sybase Replication Agent events.
STMTRACE	When set to “true,” causes tracing of LTM state monitor events.
THREADTRC	When set to "true", logs ThreadPool trace events.

- You *cannot* change the settings of SYSTEM trace flags.

  Table 1-8 lists Sybase Replication Agent SYSTEM trace flags:

Table 1-8: Sybase Replication Agent SYSTEM trace flags

Trace flag	Description
CONFIG	Configuration change event logged.
ERROR	Serious error; manual intervention may be needed to recover.
FATAL	Critical error; application shut down; manual intervention required to recover.
INFORMATION	Information only; no action required.
WARNING	Minor error; operation not affected, or problem is recoverable.

- The trace command is valid when the Sybase Replication Agent instance is in either *Admin* or *Replicating* state.

See also log_system_name

CHAPTER 2

Configuration Parameters

This chapter describes the Sybase Replication Agent configuration file and configuration parameters.

Configuration parameter overview

Configuration parameters record the user-configurable settings that control how a Replication Agent instance operates. The current values of all configuration parameters are stored in the *configuration file* of each Replication Agent instance.

Sybase Replication Agent configuration file

The configuration file is created automatically when you create a Sybase Replication Agent instance. It resides in the instance subdirectory, under the Sybase Replication Agent base directory.

The configuration file is named after the Replication Agent instance, with the extension *.cfg* (for example, if the instance is named "my_ra," the configuration file is *my_ra.cfg*).

Each time a Sybase Replication Agent instance starts up, it reads the configuration file to get the configuration information it needs to run. After start-up, the only time the Sybase Replication Agent accesses the configuration file is when the ra_config or ra_set_login command is invoked to change the value of a configuration parameter.

When the value of a configuration parameter is changed, Sybase Replication Agent saves the new value, overwriting the entire configuration file.

Configuration file format

The configuration file is a flat ASCII file that contains configuration information for a single Sybase Replication Agent instance.

The first two lines in the configuration file identify the file as a Replication Agent configuration file and record the time that the file was last modified. For example:

```
#RA Property File
#Wed Mar 16 07:33:18 MST 2005
```

Each configuration parameter name appears on a separate line, followed by the equal symbol (=) and the current value of the parameter. For example:

```
compress_ltl_syntax=true
```

If the Sybase Replication Agent instance is not running, you can view the configuration file to examine the current Sybase Replication Agent configuration.

Note Sybase recommends that you do *not* edit the configuration file, because Sybase Replication Agent overwrites the entire configuration file every time the ra_config or ra_set_login command is invoked to change a parameter value.

If the Replication Agent instance is running, use the ra_config command to view the current Replication Agent configuration.

Changing configuration parameters

To view, set, or change the current value of a Sybase Replication Agent configuration parameter, use the ra_config command.

To change the current Sybase Replication Agent administrator login (ltm_admin_user) or administrator password (ltm_admin_pw), you must use the ra_set_login command.

Note The ltm_admin_user and ltm_admin_pw parameters cannot be changed with the ra_config command, and they do not appear in the parameter list returned by ra_config.

See Chapter 1, “Command Reference,” for more information about using the ra_config and ra_set_login commands.

Copying a Sybase Replication Agent configuration

When you create a new Sybase Replication Agent instance with the ra_admin utility, you can specify the new instance to use the same configuration parameter values as an existing Sybase Replication Agent instance.

Note When you copy an existing configuration for a new Sybase Replication Agent instance, certain configuration parameter values are not copied to the new configuration. See the Sybase Replication Agent *Administration Guide* for more information.

If you do not copy an existing configuration when you create a new Replication Agent instance, the ra_admin utility creates a default configuration file, with default values for all configuration parameters.

Configuration parameter reference

Table 2-1 lists all of the Sybase Replication Agent configuration parameters, with a brief description of each parameter.

Table 2-1: Replication Agent configuration parameters

Parameter name	Description
admin_port	Port number for Replication Agent administrative port.
column_compression	Use minimal column information.
compress_ltl_syntax	Use abbreviated LTL syntax.
connect_to_rs	Enable/disable connection from LTI to Replication Server.
ddl_password	Password for ddl_username. This configuration is available for Oracle only.
ddl_username	The database user name included in LTL for replicating DDL commands to the standby database. This configuration is available for Oracle only.
dump_batch_timeout	Time to send incomplete buffer to Replication Server.
filter_maint_userid	Log Reader filters operations with maintenance user ID.
function_password	Password for user ID passed in LTL with replicated stored procedure invocations.
function_username	User ID passed in LTL with replicated stored procedure invocations.
log_backup_files	Determines the number of log backup files kept in the log directory.
log_directory	Directory where Replication Agent system log file is located.
log_trace_verbose	Switch on/off verbose mode in trace log file.
log_wrap	Number of 1KB blocks written to log file before wrapping.
lr_ntext_byte_order	Specifies which byte order to use when replicating NTEXT data.(For Microsoft SQL Server only)
lti_batch_mode	Switches on/off LTI batch mode.
lti_max_buffer_size	Maximum number of change sets stored in the LTI input buffer.
lti_update_trunc_point	Number of LTL commands sent before LTI requests new LTM Locator.
ltl_batch_size	Size of the LTL batch buffer.
ltl_character_case	Case of database object names sent to Replication Server.

Parameter name	Description
ltl_origin_time_required	Specifies whether to send origin_time command tag in LTL.
ltm_admin_pw	Password for Replication Agent administrative port.
ltm_admin_user	User ID for Replication Agent administrative port.
max_ops_per_scan	Maximum number of operations Log Reader will read in a single log scan.
pdb_archive_path	Identifies the directory path where the Replication Agent expects to find archived Oracle *redo log* files. This configuration is available for Oracle only.
pdb_archive_remove	Enables or disables the removal of archived Oracle *redo log* files from the path specified by pdb_archive_path. This configuration is available for Oracle only.
pdb_automark_tables	Determines if the Sybase Replication Agent automatically marks tables for replication during initialization or DDL replication (For Oracle only)
pdb_auto_create_repdefs	If set to true, when tables and procedures are marked for replication, a replication definition is automatically created at Replication Server for that table or procedure. (For Oracle only)
pdb_auto_run_scripts	Automatic execution of SQL scripts used to create/remove transaction log objects and mark/unmark primary database objects.
pdb_convert_datetime	Converts native date/time formats to Sybase datetime format.
pdb_dflt_column_repl	Enables replication for LOB columns by default when table is marked.
pdb_dflt_object_repl	Enables replication by default when object is marked.
pdb_exception_handling	Allows primary database to continue operation if Replication Agent triggers fail.
pdb_include_archives	Enables or disables the use of Oracle archive log files. (For Oracle only)
pdb_support_large_identifier	To support replication of large identifiers up to 255 characters in length with Replication Server 15.0.

Parameter name	Description
pdb_xlog_device	Name of the primary database device.
pdb_xlog_prefix	Character string prefix used to identify transaction log objects.
pdb_xlog_prefix_chars	Non-alphabetic characters allowed in pdb_xlog_prefix.
pds_connection_type	Type of connection to primary data server.
pds_database_name	Name of database replicated from primary data server.
pds_datasource_name	Data source name of database replicated from primary data server.
pds_host_name	Name of primary data server host machine.
pds_integrated_security	Determines if the Sybase Replication Agent should use Windows authentication when connecting to the primary SQL Server. (For Microsoft SQL Server only)
pds_password	Password for user ID Replication Agent uses to access primary data server.
pds_port_number	Port number for primary data server.
pds_retry_count	Number of times to retry connection to primary data server.
pds_retry_timeout	Number of seconds to wait between connection retry attempts.
pds_server_name	Server name of primary data server.
pds_username	User ID that Replication Agent uses to access primary data server.
ra_retry_count	Number of times LTM attempts to get back to *Replicating* state after a failure.
ra_retry_timeout	Number of seconds to wait between LTM attempts to get back to *Replicating* state.
rasd_backup_dir	The directory path for Replication Agent System Database (RASD) backup files. This configuration is available for Oracle only.
rasd_database	The directory path for the Replication Agent System Database (RASD) database file. This configuration is available for Oracle only.
rasd_mirror_tran_log	Enables or disables Replication Agent System Database (RASD) transaction log mirroring. This configuration is available for Oracle only.

Parameter name	Description
rasd_trace_log_dir	The directory path for the Replication Agent System Database (RASD) trace log file. This configuration is available for Oracle only.
rasd_tran_log	The directory path for the Replication Agent System Database (RASD) transaction log file. This configuration is available for Oracle only.
rasd_tran_log_mirror	The directory path for the Replication Agent System Database (RASD) transaction log file mirror. This configuration is available for Oracle only.
rs_charset	Character set used to communicate with Replication Server.
rs_host_name	Name of primary Replication Server host machine.
rs_packet_size	Network I/O packet size sent to Replication Server.
rs_password	Password for user ID Replication Agent uses to access Replication Server.
rs_port_number	Port number for primary Replication Server.
rs_retry_count	Number of times to retry connection to primary Replication Server.
rs_retry_timeout	Number of seconds to wait between connection retry attempts.
rs_source_db	Name of primary database identified to Replication Server.
rs_source_ds	Name of primary data server identified to Replication Server.
rs_username	User ID that Replication Agent uses to access primary Replication Server.
rssd_charset	Character set used to communicate with RSSD.
rssd_database_name	Name of RSSD database.
rssd_host_name	Name of RSSD host machine.
rssd_password	Password for user ID Replication Agent uses to access RSSD.
rssd_port_number	Port number for RSSD.
rssd_username	User ID that Replication Agent uses to access RSSD.

Parameter name	Description
scan_sleep_increment	Number of seconds to increase Log Reader wait before next scan after finding no operations to replicate.
scan_sleep_max	Maximum number of seconds for Log Reader to wait before next scan after finding no operations to replicate.
skip_ltl_errors	LTI ignores error messages returned by Replication Server.
structured_tokens	LTI uses structured tokens when generating LTL output.
truncation_interval	Number of minutes to wait between automatic log truncations.
truncation_type	Methods of log truncation allowed.
use_rssd	Switches on/off access to RSSD for replication definitions.

The following sections list all Sybase Replication Agent configuration parameters in alphabetical order.

admin_port

The client socket port number of the Replication Agent.

Default

"" (empty string)

Value

A valid port number on the Replication Agent host machine.

Comments

- When you create a Sybase Replication Agent instance, you must specify a client socket port number for the instance administration port. Client applications use this port number to connect to the Sybase Replication Agent instance.
- You must specify a port number that does not conflict with any port numbers already in use on the Replication Agent host machine.
- If you change the value of the admin_port parameter with the ra_config command, the new value is recorded in the configuration file immediately, but you must shut down and restart the Sybase Replication Agent instance to make the new port number take effect.
- After you change the value of the admin_port parameter with the ra_config command, the next time you log in to the Replication Agent administration port, you must use the new port number.

column_compression

Determines whether the Log Transfer Interface component sends all columns in row after images, or only the columns that changed in an update operation.

Default: true

Values: true – enables minimal column information (only changed columns in row after images) in Log Transfer Language (LTL) for update operations.

false – disables minimal column information in LTL for update operations.

Comments:

- When the column_compression parameter is set to false, the LTI component sends complete row after images in LTL, including columns in which no data changed as a result of an update operation.
- When the column_compression parameter is set to true, the LTI component sends minimal column information in the row after images in LTL, with only the columns that changed as a result of an update operation. Columns in which no data changed as a result of the update are not sent in LTL.
- In general, setting the value of the column_compression parameter to true provides better Sybase Replication Agent throughput.

compress_ltl_syntax

Determines whether the Log Transfer Interface component compresses Log Transfer Language (LTL) commands using abbreviated syntax.

Default: true

Values: true – enables LTL compression, using abbreviated LTL syntax.

false – disables LTL compression.

Comments:

- In general, setting the value of the compress_ltl_syntax parameter to true will provide better Replication Agent throughput.
- See the Replication Server *Administration Guide* and Sybase Replication Agent *Reference Manual* for more information about LTL commands and abbreviated LTL syntax.

connect_to_rs

Enables or disables the network connection to the primary Replication Server.

Default: true

Values: true – enables the network connection to the Replication Server.

false – disables the network connection to the Replication Server.

Comments

- When the value of the connect_to_rs parameter is false, the network connection from the Sybase Replication Agent to the Replication Server is disabled, and no replication can occur.
- When the network connection to the Replication Server is disabled by the connect_to_rs parameter, the Sybase Replication Agent instance can still go to *Replicating* state, with the following limitations:
  - A "dummy" connection in the Sybase Replication Agent emulates a real connection to the Replication Server.
  - The value of the LTM Locator stored in the Replication Agent transaction log is set to zero.
  - The Maintenance User name is set to an invalid user ID.

  Note Maintenance User operations cannot be filtered when the value of the connect_to_rs parameter is false.
- You can use the connect_to_rs parameter to temporarily disable the network connection to the Replication Server for testing.
- When the value of the connect_to_rs parameter is false, you can put the Sybase Replication Agent instance in *Replicating* state, set the value of the LTITRACELTL trace flag to true, and view a readable representation of the LTL that would have been sent to the Replication Server if the connection had not been disabled.
- During normal Replication Agent operation, the value of the connect_to_rs parameter must be true.

ddl_password

Note This parameter is available only for Oracle.

Updates the log device repository in the RASD.

Identifies the password for ddl_username.

Default "" (empty string)

Value A valid password.

Comments
- The value of the ddl_password parameter can be up to 30 characters.
- The value of the ddl_password parameter is the password for the database user name specified in the ddl_username parameter.
- The value of the ddl_password parameter is encrypted in the Sybase Replication Agent configuration file.

ddl_username

Note This parameter is available only for Oracle.

The database user name included in LTL for replicating DDL commands to the standby database.

This user must have permission to execute all replicated DDL commands at the standby database.

Default <not_configured>

Value A valid user name in the standby database.

Comments
- The value of the ddl_username parameter is sent in the LTL for all replicated DDL statements.
- The value of the ddl_password parameter is the password for the database user name specified in the ddl_username parameter.
- When DDL is replicated, Replication Server will connect to the standby database using the ddl_username and ddl_password.
- Replication Server then issues the following message:

  ```
  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 standby database. If the ddl_username does not have permission to issue `ALTER SESSION SET CURRENT_SCHEMA` or to execute the DDL command against the *user* schema, the command fails.

dump_batch_timeout

Specifies the time interval to wait before sending the contents of the Log Transfer Interface (LTI) buffer to the Replication Server, even though the buffer is not full.

Default: 5

Value: An integer from 1 to 60.

Comments:

- The value of the dump_batch_timeout parameter is the number of seconds from the time the previous LTI buffer was sent to the Replication Server until the next buffer will be sent.
- The dump_batch_timeout parameter has no effect if the value of the lti_batch_mode parameter is false.

filter_maint_userid

Determines whether operations applied by the Maintenance User are ignored.

Default: true

Values: true – enables the Log Reader to ignore Maintenance User operations.

false – disables the Log Reader filter to allow replicating Maintenance User operations.

Comments:

- The filter_maint_userid configuration parameter is provided to support bidirectional replication, in which the primary database also serves as a replicate database that has transactions applied to it by a Replication Server Maintenance User.
- If the value of the filter_maint_userid parameter is true, database operations applied by the Maintenance User are *not* replicated. The Log Reader component filters out (ignores) operations applied by the Maintenance User when it reads the transaction log.
- If the value of the filter_maint_userid parameter is false, database operations applied by the Maintenance User are replicated. The Log Reader component replicates all operations on marked objects, regardless of the user that applied the operation.
- The Maintenance User login is specified when the database connection for the primary database is created in the Replication Server.

function_password

The password included in Log Transfer Language for replication of "request" stored procedures.

Default: "" (empty string)

Values: A valid password.

Comments:

- The value of the function_password parameter can be up to 30 characters.
- The value of the function_password parameter is the password for the database user name specified in the function_username parameter.
- The value of the function_password parameter is encrypted in the Sybase Replication Agent configuration file.

function_username

The database user name included in Log Transfer Language (LTL) for replication of "request" stored procedures.

Default: sa

Values: A valid user name in the primary database.

Comments:

- The value of the function_username parameter is sent in the LTL for all replicated stored procedures in the primary database.
- The value of the function_password parameter is the password for the database user name specified in the function_username parameter.

log_backup_files

The number of backup log files kept in the Replication Agent *log* directory.

Default: 3

Values: An integer greater than or equal to 1.

Comments:

- When the system log wraps, Sybase Replication Agent copies the current log file to a backup file, with a generated number appended to the file's name.

  For example, if the system log file is named *my_ra.log*, the first backup file created when the system log wraps would be named *my_ra1.log*. The second backup file created would be named *my_ra2.log*, and so on.

- When the number of backup files exceeds the value of the log_backup_files parameter, the oldest backup file (that is, the one with the lowest generated number) is deleted from the *log* directory before the next backup file is created.

log_directory

The directory for Replication Agent system log files.

Default

The path to the *log* directory created when the Sybase Replication Agent instance was created. For example:

- On Microsoft Windows platforms:

  ```
  %SYBASE%\RAX-15_0\inst_name\log
  ```

  where:

  - *%SYBASE%* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.

- On UNIX platforms:

  ```
  $SYBASE/RAX-15_0/inst_name/log
  ```

  where:

  - *$SYBASE* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.

Value

A valid path on the Replication Agent host machine.

Comments

- When a Sybase Replication Agent instance is created, the *log* directory is created as part of the instance directories. The default value of the log_directory parameter points to that directory.
- If you specify any valid path as the value of the log_directory parameter, the Sybase Replication Agent instance places its system log files in the directory you specify the next time it is started.
- If you specify the default value of the log_directory parameter by using the default keyword in the ra_config command, then the next time it is started, Replication Agent will place its system log files in the *log* directory that was created when the Sybase Replication Agent instance was created.

- If you change the value of the log_directory parameter with the ra_config command, the new value is recorded in the configuration file immediately. However, you must shut down and restart the Sybase Replication Agent instance to make the new value take effect.

log_trace_verbose

Enables or disables additional diagnostic information in Replication Agent system log files.

Default: false

Values: true – enables detailed diagnostic information in log files.

false – disables detailed diagnostic information in log files.

Comment: Detailed diagnostic information is intended for troubleshooting only, with assistance from Sybase Technical Support.

log_wrap

The maximum size of the Replication Agent system log file before wrapping.

Default: 1000

Value: An integer greater than or equal to 1000.

Comments:

- The value of the log_wrap parameter is the number of 1KB blocks written by Sybase Replication Agent, before it wraps the system log file.
- Larger values for the log_wrap parameter allow more log history in each file. Smaller values produce smaller log files.
- When the log file wraps, Sybase Replication Agent copies the current log file to a backup file, with a generated number appended to the file's name.

  For example, if the system log file is named *my_ra.log*, the first backup file created when the system log wraps would be named *my_ra1.log*. The second backup file created would be named *my_ra2.log*, and so on.
- When the number of backup files exceeds the value of the log_backup_files parameter, the oldest backup file (that is, the one with the lowest generated number) is deleted from the *log* directory before the next backup file is created.

lr_ntext_byte_order

Note For Microsoft SQL Server only.

Specifies which byte order to use when replicating NTEXT data.

Default: big

Values: big – big endian

little – little endian

Comments: Big endian indicates a left-to-right byte order architecture; little endian indicates a right-to-left byte order architectures.

lti_batch_mode

Enables or disables the Log Transfer Interface component LTL batch mode.

Default: true

Values: true – enables LTL batch mode.

false – disables LTL batch mode.

Comments:

- If the value of the lti_batch_mode parameter is true, the LTI component sends LTL commands to the Replication Server in batches, instead of one command at a time:
  - The LTI component fits as many LTL commands as it can into its LTL batch mode buffer, before it sends any commands to the Replication Server.
  - When the time interval specified in the dump_batch_timeout parameter expires, the LTI component sends the current LTL batch mode buffer contents to the Replication Server, even if the buffer is not full.
- If the value of the lti_batch_mode parameter is false, the LTI component sends individual LTL commands to the Replication Server for each change set in its input queue.
- When Sybase Replication Agent connects to the Replication Server, it determines the version of the Replication Server:
  - If the Replication Server version is earlier than 12.5, the size of the LTL batch mode buffer is set to 16KB automatically.

    - If the Replication Server version is 12.5 or later, Sybase Replication Agent sets the size of the LTL batch mode buffer to the size specified by the ltl_batch_size parameter.
- If the Replication Server version is 12.5 or later, you can use the Sybase Replication Agent ltl_batch_size parameter to set the size of the LTI component's LTL batch mode buffer.

  Note Adjusting the size of the LTL batch mode buffer can help you optimize the performance of the replication system.

- When the Replication Server version is earlier than 12.5 and the value of the lti_batch_mode parameter is true, if any single LTL distribute command exceeds the 16K size of the LTL batch mode buffer, Replication Server returns an error and it drops its DSI connection with the Sybase Replication Agent.
- In general, setting the value of the lti_batch_mode parameter to true provides better Replication Agent throughput.

lti_max_buffer_size

The maximum size of the Log Transfer Interface (LTI) component's queues.

Default: 5000

Value: An integer in the range of 1000 to 100000.

Comments:

- The value of the lti_max_buffer_size parameter is the maximum number of operations that can be stored in the LTI component's inbound and outbound queues:
    - Operations in the inbound queue represent change sets received from the Log Reader component.
    - Operations in the outbound queue are the Log Transfer Language commands to be sent to the Replication Server.
- The LTI component's inbound queue is a bounded buffer that blocks the processing of the Log Reader component when it gets full.

lti_update_trunc_point

The number of Log Transfer Language (LTL) commands sent before requesting a new LTM Locator from the Replication Server.

Default: 1000

Value: An integer from 1 to 100000.

Comments:

- The value of the lti_update_trunc_point parameter is the number of LTL commands that Sybase Replication Agent sends to the Replication Server, before it requests a new LTM Locator (secondary truncation point) from the Replication Server.
- Lower numbers cause Sybase Replication Agent to request a new LTM Locator from the Replication Server more often.
- If the value of the truncation_type parameter is locator_update, setting the value of the lti_update_trunc_point parameter to a lower number causes automatic log truncation to occur more frequently.
- The value of the lti_update_trunc_point parameter is a trade-off between better system performance and longer recovery time:
  - Lower values reduce the time it takes to recover from a replication failure, but they may have an adverse affect on overall system throughput.
  - Higher values improve overall system throughput, but they may increase the time it takes to recover from a replication failure.
- If the Replication Agent is operating in an unreliable network environment, it may be prudent to set the lti_update_trunc_point parameter to a lower value to ensure faster recovery.

ltl_batch_size

The size of the Log Transfer Interface component's LTL batch mode buffer.

Default: 40000

Value: An integer from 1 to 10485760.

Comments:

- The value of the ltl_batch_size parameter is the size (in bytes) of the LTI component's LTL batch mode buffer.
- When Sybase Replication Agent connects to the Replication Server, it determines the version of the Replication Server:

- If the Replication Server version is earlier than 12.5, the size of the LTL batch mode buffer is set to 16K automatically, and the value of the ltl_batch_size parameter is ignored.
- If the Replication Server version is 12.5 or later, Sybase Replication Agent sets the size of the LTL batch mode buffer to the size specified by the ltl_batch_size parameter.

- The Log Transfer Interface component uses the LTL batch mode buffer only if the value of the lti_batch_mode parameter is true. If the value of the lti_batch_mode parameter is false, the LTL batch mode buffer is not used.

ltl_character_case

The character case used for database object names in Log Transfer Language (LTL) sent to the Replication Server.

Default

asis

Values

asis – Database object names are sent in the same character case as they are returned from the primary database, or (if the value of the use_rssd parameter is true) in the same character case as they are specified in replication definitions.

lower – Database object names are sent in *all lowercase*, regardless of how they are returned from the primary database, or specified in replication definitions.

upper – Database object names in LTL are sent in *all uppercase*, regardless of how they are returned from the primary database, or are specified in replication definitions.

Comments

- The ltl_character_case configuration parameter allows you to customize the handling of database object names in LTL to work with replication definitions that specify the object names differently than the way the primary database returns them.
- If the value of the ltl_character_case parameter is asis, and the value of the use_rssd parameter is true, database object names are sent in the same character case as they are specified in replication definitions.
- If the value of the ltl_character_case parameter is asis, and the value of the use_rssd parameter is false, database object names are sent in the same character case as they are returned from the primary database.
- If replication definitions specify database object names in all lowercase, set the value of the ltl_character_case parameter to lower.

- If replication definitions specify database object names in all uppercase, set the value of the ltl_character_case parameter to upper.
- If you want to send database object names with "mixed" character case (for example, MyTable), set the value of the ltl_character_case parameter to asis.

ltl_origin_time_required

Enables or disables the Log Transfer Language (LTL) origin_time command tag.

Default false

Values true – enables the origin_time command tag in LTL.

false – disables the origin_time command tag in LTL.

Comments

- If the value of the ltl_origin_time_required parameter is true, the Log Transfer Interface component includes the origin_time command tag in the LTL it generates.
- If a Replication Server function string checks for the origin_time command tag, set the value of the ltl_origin_time_required parameter to true.
- The datetime value placed in the LTL origin_time command tag is the time that the original primary database operation was recorded in the transaction log, not the time it was scanned and processed by the Log Reader component.
- Setting the value of the ltl_origin_time_required parameter to false provides better Sybase Replication Agent throughput.
- If you use Replication Server Manager to report latency, you must set the value of the ltl_origin_time_required parameter to true.

ltm_admin_pw

The Replication Agent administrator login password.

Default "" (empty string)

Value A valid password.

Comments

- The value of the ltm_admin_pw parameter is the password for the user name authorized to log in to the Sybase Replication Agent.

- The value of the ltm_admin_pw parameter is encrypted in the Sybase Replication Agent configuration file.
- To change the value of the ltm_admin_pw parameter, use ra_set_login.
- When you change the value of the ltm_admin_pw parameter with ra_set_login, the new value is recorded in the configuration file immediately. However, you must shut down and restart the Sybase Replication Agent instance to make the new password take effect.

  After you change the value of the ltm_admin_pw parameter with ra_set_login, you must use the new password next time you log in to the Replication Agent.

ltm_admin_user

The Replication Agent administrator login name.

Default: sa

Value: A valid user name on the Replication Agent host machine.

Comments:

- The value of the ltm_admin_user parameter is the user name authorized to log in to the Sybase Replication Agent.
- To change the value of the ltm_admin_user parameter, use the ra_set_login command.
- If you change the value of the ltm_admin_user parameter with the ra_set_login command, the new value is recorded in the configuration file immediately. However, you must shut down and restart the Sybase Replication Agent instance to make the new administrator name take effect.
- After you change the value of the ltm_admin_user parameter with ra_set_login, you must use the new administrator name the next time you log in to the Sybase Replication Agent.
- Only one administrator name is valid at any time.

max_ops_per_scan

Note This parameter is not available in Oracle.

The maximum number of operations the Log Reader component reads during each log scan operation.

Default: 1000

Values: An integer from 25 to 2147483647.

Comments:

- The value of the max_ops_per_scan parameter is the maximum number of database operations that can be read from the Replication Agent transaction log during each Log Reader scan operation (the size of the Log Reader operation queue).
- The Log Reader component always reads at least one transaction in each scan, regardless of how many operations are in the transaction.

  For example, if the value of the max_ops_per_scan parameter is 1000, and a transaction contains 1200 operations, the Log Reader component reads all 1200 operations in one scan when it reads that transaction.
- See the Sybase Replication Agent *Database Guide* for more information about how the max_ops_per_scan parameter affects Replication Agent performance.

pdb_archive_path

Identifies the directory path where the Replication Agent expects to find archived Oracle *redo log* files. This configuration is available for Oracle only.

Default: none

Values: A valid directory path on the machine hosting the Replication Agent that points to a location where Oracle places the archived *redo log* files.

Comments:

- Setting of the configuration parameter is required when configuration property pdb_include_archives is set to *true*, and must be set to a valid location before the Replication Agent can be placed in a replicating state.
- If the Replication Agent cannot find an expected log record in the Oracle online redo logs, the Replication Agent will search this directory for the archived log file containing the required record.
- See the pdb_archive_remove and pdb_include_archives configuration properties.

pdb_archive_remove

Enables or disables the removal of archived Oracle *redo log* files from the path specified by pdb_archive_path. This configuration is available for Oracle only.

Default: false

Values: true – Allows the removal of archived Oracle *redo log* files from the path specified by pdb_archive_path. Removal occurs based on the execution of command pdb_truncate_xlog, or the timing of automatic truncation based on parameters truncation_type and truncation_interval.

false – Disables the removal of archived Oracle redo log files.

Comments:

- Set this configuration to *true* when the path specified by pdb_archive_path is established solely for Replication support, and automatic removal of unneeded *archived log* files is desired.
- If the path specified by pdb_archive_path is shared by other processes, or the removal of *archived log* files is expected to be performed by processes other than the Replication Agent, this parameter should be false.
- Parameters truncation_type and truncation_interval, and command pdb_truncate_xlog have no impact when this configuration parameter is set to *false*.
- See the pdb_archive_path, truncation_type, truncation_interval configuration properties. Also, see the pdb_truncate_xlog command.

pdb_auto_create_repdefs

When used to mark a table or procedure, a replication definition is created at Replication Server for each table or procedure that is marked for replication.

Default: false

Values: true – Replication definitions are automatically created at the Replication Server when tables are marked.

false – No replication definitions are created when tables are marked.

Comments:

- When pdb_setrepproc is invoked to mark a procedure or procedures, a replication definition is created at Replication Server for each procedure that gets marked for replication, if this property is set to *true*.
- When pdb_setreptable is invoked to unmark a table or tables, the replication definition is dropped at Replication Server for each table that gets unmarked for replication, if this property is set to *true*.

- When pdb_setrepproc is invoked to unmark a procedure or procedures, a replication definition is dropped at Replication Server for each procedure that gets unmarked for replication if this property is set to *true*.
- When pdb_xlog is initialized and table auto marking is enabled, a replication definition is created for each table that is marked for replication if this property is set to *true*.
- Replication definition names for tables always begin with the prefix "*ra$*," followed by a unique alphanumeric identifier (maximum of 8 characters), and ending with a table or object name. For example, for a replicate name of "My Table," the resulting repdef name is "ra$0x7952_mytable." For an especially long replicate name of "mytable8901234567890 1234567890" (30 characters), the resulting repdef name is "ra$0x7952_mytable8901234567890" (30 characters maximum).
- Replication definition names for procedures are the same name as the procedure.

pdb_automark_tables

Determines if the Sybase Replication Agent automatically marks tables for replication during initialization or DDL replication for Oracle only.

Default: false

Values: true – User tables are automatically marked during initialization or DDL replication.

false – User tables are not automatically marked during initialization or DDL replication. They must always be marked using the pdb_setreptable command (default).

Comments

- The default value for pdb_automark_tables is set to *false* when an Sybase Replication Agent instance is created. In this default setting, tables are never automatically marked for replication. If automatic marking of tables is desired, this configuration parameter value should be changed to *true*. When set to *true*, all user tables (those whose owners are not contained in the owner_filter_list) will be marked for replication when the pdb_xlog command is executed with the init keyword. In addition, when replication of DDL commands is enabled (pdb_setrepddl setting is disabled by default) any create table command for a user table (those whose owners are not contained in the owner_filter_list) will automatically be marked for replication.

- Automatic marking of new tables (those created in the primary database with the create table command) will only occur when replication of DDL commands is enabled (pdb_setrepddl is set to enable) and the table is a user table (those whose owners are not contained in the owner_filter_list) and pdb_automark_tables is set to *true*.
- Tables are automatically unmarked for replication when a drop table command issued at the primary and is recorded in the *redo* log, regardless of the settings of pdb_setrepddl or pdb_automark_tables. This is due to the fact a dropped table cannot be replicated from.
- Automatic marking of user tables is independent of manual marking of tables using the pdb_setreptable command. Meaning, you can always mark or unmark individual or all tables for replication using the pdb_setreptable command, regardless of the setting of pdb_automark_tables.

pdb_auto_run_scripts

Determines whether Sybase Replication Agent automatically runs scripts (for transaction log creation and removal, and object marking and unmarking) at the primary database.

Default

true

Values

true – Sybase Replication Agent automatically runs scripts.

false – Sybase Replication Agent generates and saves the scripts, but it does not automatically run them at the primary database.

Comments

- When the pdb_xlog command is invoked to create or remove the transaction log, Replication Agent generates a script to create or remove the transaction log base objects.
- When either the pdb_setrepproc or pdb_setreptable command is invoked to mark or unmark an object in the primary database, Sybase Replication Agent generates a script to create or remove the transaction log objects necessary for object marking.
- Sybase Replication Agent always saves the scripts in a file. Log creation and removal scripts are saved in files named *partinit.sql* and *partdeinit.sql*. Object marking and unmarking scripts are saved in files named *partmark.sql* and *unmark.sql*.

- When the pdb_auto_run_scripts parameter is set to false, the scripts are created but no action is taken. This allows you to review the scripts to see what action will be taken before execution. You cannot execute the scripts. You must set pdb_auto_run_scripts parameter back to true and re-execute the command to have the desired action take place.
- As described above for the pdb_xlog and pdb_setrepproc commands, Oracle creates the *partinit*, *partdeinit*, *partmark* and *partunmark* scripts. For Oracle, these scripts can *not* be executed (since they do not update the RASD) and are for informational purposes only. This is different than the trigger-based Replication Agent script, which *can* be executed.
- For Oracle, this property must be set to true for initialization to occur.

pdb_convert_datetime

Determines whether Sybase Replication Agent converts non-Sybase temporal datatypes to the Sybase datetime format.

Default true

Values true – Sybase Replication Agent converts all data in the primary database native date/time datatypes to the Sybase datetime format.

false – Sybase Replication Agent replicates data in the primary database native datetime datatypes as character strings using the Sybase rs_oracle_datetime datatype.

Comments

- The pdb_convert_datetime parameter is provided for backward compatibility with previous versions of Sybase Replication Agents and Replication Server. If you use Replication Server version 12.0 or later, Sybase recommends that you use the Replication Server heterogeneous datatype support (HDS) feature for all datatype conversion and translation.
- Sybase Replication Agent checks the value of the pdb_convert_datetime parameter at the time an object is marked for replication. Transaction log objects that support replication of the marked object are constructed to provide the desired date format.

  If you change the value of the pdb_convert_datetime parameter *after* an object is marked, it has no effect on the marked object. To change the datetime datatype conversion for a marked object, you must unmark the object, change the value of the pdb_convert_datetime parameter, then re-mark the object.

- For trigger-based Replication Agents, the actual conversion of datatypes takes place when an operation is recorded in the Replication Agent transaction log. For log-based Replication Agents, the conversion takes place after the log records have been read and before LTL is generated to send to the Replication Server.
- Any missing component in the non-Sybase date/time datatype format is treated as an implied 0 (zero) when it is converted to the Sybase datetime format.
- When the value of the pdb_convert_datetime parameter is true, the replication definition for each table should specify that the declared datatype for all date/time columns is datetime.
- If the value of the pdb_convert_datetime parameter is false, the Replication Agent sends date/time data to the primary Replication Server as default-sized character strings. The default character string size varies by database and datatype:
  - DB2 Universal Database: DATE = char(10), TIME = char(8), TIMESTAMP = char(16)
  - Microsoft SQL Server: datetime or smalldatetime = char(23), timestamp = binary(8)
  - Oracle: DATE = char(8)

  If the non-Sybase date/time format requires a longer string to be replicated correctly, you can either set the value of the pdb_convert_datetime parameter to true, or modify the trigger-based Replication Agent table marking script (mark.sql) to create larger shadow table date/time columns.
- Set the value of the pdb_convert_datetime parameter to true if *all* date/time values replicated from the primary database will be replicated as the Sybase datetime datatype.
- pdb_convert_datetime must be false if a table containing replicated LOB columns has datetime datatype in the primary key.
- Replication Agent date/time datatype conversion does not work with LOB column replication, unless either of the following conditions exist (these conditions are *not* required for Oracle):
  - There are no date/time columns in the tables that have LOB column replication enabled, or
  - The primary keys in tables that have LOB column replication enabled do not contain date/time datatypes.

Otherwise, if you use the pdb_setrepcol command to enable LOB column replication, you must set the value of the pdb_convert_datetime parameter to false.

The Replication Agents, other than Oracle, query LOB data directly from the primary database (the LOB data is not captured by the triggers and is not logged for UDB). To successfully query a primary database table for a LOB column value, any date column value must retain the primary database format and structure for the date value to appear correctly in the query. The format and structure for the date value cannot be converted to the Sybase datetime format.

- Sybase recommends that you set the value of the pdb_convert_datetime parameter to false for better Sybase Replication Agent throughput performance and optimal datatype handling.

pdb_dflt_column_repl

Determines whether LOB column replication is enabled by default when tables are marked.

Default: true

Values: true – LOB column replication is enabled by default (automatically) when tables are marked.

false – LOB column replication is disabled by default when tables are marked.

Comments:

- If the value of the pdb_dflt_column_repl parameter is false when a table is marked for replication, no transactions that affect LOB columns in the table can be replicated until replication is explicitly enabled with the pdb_setrepcol command.
- You can use the pdb_setrepcol command to enable or disable replication for all LOB columns in all marked tables at once.
- When replication is disabled for a LOB column, any part of an operation that affects that column will not be recorded in the transaction log, even if the operation also affects other columns for which replication is enabled.

pdb_dflt_object_repl

Determines whether replication is enabled by default when objects (tables or stored procedures) are marked.

Default: true

Values: true – enables replication by default (automatically) when objects are marked.

false – disables replication by default when objects are marked.

Comments:

- If the value of the pdb_dflt_object_repl parameter is false when a table is marked for replication, no transactions can be replicated from that table until replication is explicitly enabled with the pdb_setreptable command.
- If the value of the pdb_dflt_object_repl parameter is false when a stored procedure is marked for replication, no invocations of that stored procedure can be replicated until replication is explicitly enabled with the pdb_setrepproc command.
- You can use the pdb_setrepproc or pdb_setreptable command to enable or disable replication for all marked stored procedures or tables at once.
- When replication is disabled for a table, no operations that affect that table will be recorded in the transaction log.
- When replication is disabled for a stored procedure, no invocations of that stored procedure will be recorded in the transaction log.

pdb_exception_handling

Determines how Replication Agent trigger errors are handled in the primary database.

Default: false

Value: true – If an error occurs during trigger execution, the error is logged in the exceptions table and the transaction in the primary database continues without being recorded in the Replication Agent transaction log.

false – If an error occurs during trigger execution, the transaction in the primary database fails, reporting the error to the entity that submitted the transaction.

Comments

- When the value of the pdb_exception_handling parameter is true, trigger errors are logged in the Replication Agent exceptions table, and transactions against the marked object are allowed to continue without being captured in the Replication Agent transaction log. Although replication fails for all marked objects, the operation of the primary database is not interrupted by the failure of a Replication Agent trigger.

  Note If the value of the pdb_exception_handling parameter is true and replication is interrupted by a trigger error, transactions will continue to occur in the primary database, but they will *not* be replicated. When transactions in the primary database are not replicated, the replicate database must be rematerialized before replication can resume.

- Once replication is interrupted by the failure of a Replication Agent trigger, it cannot be resumed until the cause of the trigger error is corrected.
- When the value of the pdb_exception_handling parameter is false, a trigger error will cause the transaction in the primary database to fail. The primary database reports the transaction failure to the entity that submitted the transaction. Although operation of the primary database is interrupted, the integrity of the replication system is not adversely affected (no transactions can occur without being replicated).
- If maintaining the availability of the primary database is a higher priority than maintaining integrity of the replicate database, set the value of the pdb_exception_handling parameter to true.
- If maintaining the integrity of the replicate database is a higher priority than maintaining the availability of the primary database, set the value of the pdb_exception_handling parameter to false.

pdb_include_archives

Enables or disables the use of Oracle archive log files. This configuration is available for Oracle only.

Default

false

Values

true – Allows reading of the archived Oracle redo log files from the path specified by pdb_archive_path. The configuration of Oracle automatic archiving is supported under this mode. Removal of old archives logs (no longer needed to support replication) may be provided using configuration property pdb_archive_remove.

false – Only on-line redo logs files are read. Oracle automatic archiving must be disabled. The RepAgent executes Oracle archive commands to archive the redo logs once they are no longer needed for replication.

Comments

- Set the configuration to true when use of archive logs is preferred or when Oracle must be configured to perform automatic archiving. Set this value to false if accessing only the on-line redo logs is preferred.
- Set this value to false if using only the online redo logs is preferred.
- See the pdb_archive_path, truncation_type, truncation_interval configuration properties.

pdb_support_large_identifier

To support replication of large identifiers up to 255 characters in length with Replication Server 15.0.

Default

false

Value

- true – Objects containing large identifiers may be marked for replication.
- false – Objects containing large identifiers may *not* be marked for replication.

Comments

- If pdb_support_large_identifier value is *false*, when an object (Table/Procedure/Function) is being marked for replication, the object is checked for any identifiers that are longer than 30 characters. An error is returned and the object is not marked for replication if the object has identifiers longer than 30 characters.
- This property may be set to true if the Replication Server being used is at version 15.0 or later and the replicate database must be able to support large identifiers.
- When pdb_support_large_identifier is set to true, objects being marked for replication are not checked for identifiers longer than 30 characters.

pdb_timezone_file

Specifies the file to read at Sybase Replication Agent initialization to obtain Oracle time zone information for Oracle only.

Default: <not configured>

Value: A valid path to the Oracle time zone file including the *time zone* file name.

Comments:

- If the value is not specified, it will default to the Oracle installation's *oracore/zoneinfo/timezone.dat* file. For example,

  ```
  $ORACLE_HOME/oracore/zoneinfo/timezone.dat
  ```

- The *timezone* file specified must be for the same release and platform as the primary Oracle database. For example, an Oracle 9i timezone file is not compatible with an Oracle 10g primary database, and a Windows *timezone* file is not compatible with UNIX.

pdb_xlog_device

The primary database device on which Sybase Replication Agent transaction log objects are created.

Default: NULL

Value: A valid primary database device name or NULL.

Comments:

- The value of the pdb_xlog_device parameter is the device specification of the primary database device to be used in SQL scripts generated by the Sybase Replication Agent to create transaction log objects.
- The pdb_xlog_device parameter allows you to specify a single device on which all Sybase Replication Agent transaction log objects will be created, even if the database uses multiple devices.
- If the value of the pdb_xlog_device parameter is NULL, no device is specified in the SQL create statements, and Sybase Replication Agent transaction log objects are placed on the primary data server's system-defined default device.

pdb_xlog_prefix

The prefix string used in database object names to identify Replication Agent transaction log objects.

Default: ra_

Value: A character string of 1 to 3 characters.

Comments:

- When Sybase Replication Agent generates database object names for transaction log components in the primary database, it uses the value of the pdb_xlog_prefix parameter as an object name prefix.
- Sybase Replication Agent uses the value of the pdb_xlog_prefix parameter to recognize its transaction log objects in the primary database. Therefore, if you change the value of the pdb_xlog_prefix parameter after the transaction log objects are created, Sybase Replication Agent will not be able to find its transaction log objects.
- The value of the pdb_xlog_prefix_chars parameter specifies the non-alphabetic characters that can be used in the prefix string.

pdb_xlog_prefix_chars

The non-alphabetic characters that are allowed in the database object name prefix string that identifies Sybase Replication Agent transaction log objects.

Default: _#@ (DB2 Universal Database)

_$#@ (Microsoft SQL Server)

_#$ (Oracle)

Value: A string of characters with no separators.

Comments:

- The default value of the pdb_xlog_prefix_chars parameter depends on the type of primary database that the Replication Agent instance was created for. The default value is based on the standard, non-alphabetic characters allowed by each non-Sybase database.
- When you set or change the value of the pdb_xlog_prefix_chars parameter, the new value replaces any existing value; it does not add or append the new value to a previous value.
- When you use the ra_config command to set the value of the pdb_xlog_prefix parameter, any non-alphabetic characters specified on the command line are validated against the value of the pdb_xlog_prefix_chars parameter.
- Alphabetic characters a-z are always valid in the pdb_xlog_prefix parameter, and they need not be specified.

- Sybase Replication Agent does not support delimited names for transaction log objects, so you cannot use a space character in the value of the pdb_xlog_prefix parameter.
- The value you specify for the pdb_xlog_prefix_chars parameter is not validated. There are no restrictions on the characters you can include.

Note The primary data server may restrict the characters used in certain positions in a database object name. Refer to the documentation for your primary data server for more information.

pds_connection_type

The type of connectivity driver used on the primary database connection.

Default

<not_configured> (One of the following values is set automatically when the Replication Agent instance is created.)

Values

MSSQLJDBC – Replication Agent uses the Microsoft SQL Server JDBC 2.0 driver to connect to the primary Microsoft SQL Server database.

ORAJDBC – Replication Agent uses the Oracle JDBC driver to connect to the primary Oracle database.

UDBODBC – Replication Agent uses the DB2 Universal Database JDBC driver to connect to the primary database in DB2 Universal Database.

Comments

- The value of the pds_connection_type parameter is set automatically at the time a Sybase Replication Agent instance is created. The specific value depends on the type of Sybase Replication Agent instance created.

  Note Sybase recommends that you do *not* change the default value of the pds_connection_type parameter.

- The value of the pds_connection_type parameter determines which of several other Sybase Replication Agent configuration parameters related to the primary database connection must also have values specified.
  - MSSQLDBC requires corresponding values for the following parameters:
    - pds_server_name
    - pds_port_number

- pds_database_name
- UDBODBC requires corresponding values for the following parameters:
  - pds_database_name
  - pds_datasource_name
- ORAJDBC requires corresponding values for the following parameters:
  - pds_host_name
  - pds_port_number
  - pds_database_name

Comment

The value of the pds_connection_type parameter is set automatically when a Replication Agent instance is created.

pds_database_name

The name of the primary database.

Default

<not_configured>

Value

A valid database name.

Comments

- The value of the pds_database_name parameter is the name of the primary database on the primary data server.

  Note Some primary data servers may not support multiple databases in a single instance of the data server. In that case, the value of the pds_database_name parameter should be the name of the data server instance.

- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

pds_datasource_name

The data source name (DSN) specified for the ODBC driver used on the primary database connection.

Default <not_configured>

Value A valid ODBC data source name.

Comments
- The value of the pds_datasource_name parameter is the data source name (DSN) of the ODBC driver on the primary database connection.
- If the value of the pds_connection_type parameter is UDBODBC, the value of the pds_datasource_name parameter must be the database alias of the primary database in the DB2 Universal Database server.
- This parameter is used only if the value of the pds_connection_type parameter is UDBODBC.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

pds_host_name

The name of the primary data server host machine.

Default <not_configured>

Value A valid host name.

Comments
- The value of the pds_host_name parameter is the network name of the host machine on which the primary data server resides.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

pds_integrated_security

Determines if the Sybase Replication Agent should use Windows authentication when connecting to the primary SQL Server. (For Microsoft SQL Server only)

Default false

Value
- true -- Specifies that Sybase Replication Agent should connect to the primary SQL Server using Windows authentication.
- false -- Specifies that Sybase Replication Agent should connect to the primary SQL Server using SQL Server authentication (default).

Comments

- The default value for pds_integrated_security is set to false when a Sybase Replication Agent instance is created. In this default setting, you must configure the pds_password parameter, and the primary SQL Server must be configured to allow SQL Server authentication.
- Set this value to true to have the Sybase Replication Agent connect to the primary SQL Server using Windows authentication. The Windows environment and the primary SQL Server must be configured to use Windows authentication. For more information, see the Microsoft SQL Server section in the Replication Agent *Primary Database Guide*.

pds_password

The password that the Replication Agent uses for primary data server access.

The password that Sybase Replication Agent uses for primary data server access.

Default

"" (empty string)

Value

A valid password.

Comments

- The value of the pds_password parameter is the password for the user login name that the Sybase Replication Agent uses to access the primary data server.
- The value of the pds_password parameter is encrypted in the Sybase Replication Agent instance configuration file.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

pds_port_number

The client port number for the primary data server.

Default

1111

Value

A valid port number on the primary data server host machine.

Comments

- The value of the pds_port_number parameter is the client port number for the primary data server.

- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

pds_retry_count

The number of times the Sybase Replication Agent tries to establish a connection to the primary database.

Default: 5

Value: An integer from 0 to 2,147,483,647.

Comments:

- The value of the pds_retry_count parameter is the number of times that Sybase Replication Agent will try to establish a network connection to the primary database after a connection failure.
- Sybase recommends a setting of 5 for this parameter.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

pds_retry_timeout

The number of seconds the Sybase Replication Agent waits between retry attempts to connect to the primary database.

Default: 10

Value: An integer from 0 to 3600.

Comments:

- The value of the pds_retry_timeout parameter is the number of seconds that the Sybase Replication Agent will wait between retry attempts to establish a network connection to the primary database after a connection failure.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

pds_server_name

The server name of the primary data server.

Default	<not_configured>
Value	A valid server name.
Comments	• The value of the pds_server_name parameter is the server name of the primary data server. • See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

pds_username

The user login name that the Replication Agent uses for primary data server access.

Default	<not_configured>
Value	A valid user name.
Comments	• The value of the pds_username parameter is the login name that the Sybase Replication Agent uses to log in to the primary data server. This login name must be defined in the primary data server, with appropriate privileges or authority in the primary database. • The Replication Agent uses this login to access primary database objects and to create, remove, and manage its transaction log objects in the primary database. • See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

ra_retry_count

The number of times the Sybase Replication Agent attempts to restart replication after a failure.

Default	2
Value	An integer greater than 0.

Comments

- The value of the ra_retry_count parameter is the number of times that the Log Transfer Manager component will try to get the Sybase Replication Agent instance back into *Replicating* state after a failure or error causes the instance to go to *Admin* state.
- When a network connection fails, the Sybase Replication Agent attempts to re-establish the connection, using the values stored in its connection configuration parameters for that connection.
- If the Sybase Replication Agent is unable to re-establish a connection after the number of retries specified in the pds_retry_count or rs_retry_count parameter, then the Sybase Replication Agent instance goes to *Admin* state and the Log Transfer Manager component attempts to return the Sybase Replication Agent instance to *Replicating* state, based on the settings of the ra_retry_count and ra_retry_timeout parameters.

ra_retry_timeout

The number of seconds the Sybase Replication Agent waits between attempts to restart replication after a failure.

Default: 10

Value: An integer greater than 0.

Comment: The value of the ra_retry_timeout parameter is the number of seconds that the Log Transfer Manager component will wait between its attempts to get the Sybase Replication Agent instance back into *Replicating* state after a failure causes the instance to go to *Admin* state.

rasd_backup_dir

Note This configuration is available for Oracle only.

The directory path for Replication Agent System Database (RASD) backup files.

Note The rasd_backup_dir parameter is used *only* by the Replication Agent for Oracle. Replication Agent for DB2 UDB, and Microsoft SQL Server do *not* use the rasd_backup_dir parameter.

Default

The path to the RASD *backup* directory created automatically when the Sybase Replication Agent instance was created. For example:

- On Microsoft Windows platforms:

  %SYBASE%\RAX-15_0*inst_name*\repository\backup

  where:

  - *%SYBASE%* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.

- On UNIX platforms:

  $SYBASE/RAX-15_0/*inst_name*/repository/backup

  where:

  - *$SYBASE* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.

Value

A valid path on the Replication Agent host machine.

Comments

- When you create a Sybase Replication Agent instance, an RASD *backup* directory is created automatically as part of the instance directory structure. The default value of the rasd_backup_dir parameter points to that directory.
- If you specify any valid path as the value of the rasd_backup_dir parameter, Sybase Replication Agent places its RASD backup files in that directory during RASD backup operations, and it looks in that directory for the RASD backup files during restore operations.

rasd_database

Note This configuration is available for Oracle only.

The directory path for the Replication Agent System Database (RASD) database file.

Note The rasd_database parameter is used *only* by the Replication Agent for Oracle. Replication Agent for DB2 UDB and Microsoft SQL Server do *not* use the rasd_database parameter.

Default

The path to the RASD database file created automatically when the Sybase Replication Agent instance was created. For example:

- On Microsoft Windows platforms:

  %SYBASE%\RAX-15_0*inst_name*\repository*inst_name*.db

  where:

  - *%SYBASE%* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.
- On UNIX platforms:

  $SYBASE/RAX-15_0/*inst_name*/repository/*inst_name*.db

  where:

  - *$SYBASE* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.

Value

A valid path and RASD database file name on the Replication Agent host machine.

Comments

- When you create a Sybase Replication Agent instance, the *repository* directory and the RASD database file are created automatically. The default value of the rasd_database parameter points to the RASD database file in that directory.
- If you specify any valid path and RASD database file name as the value of the rasd_database parameter, the Sybase Replication Agent instance looks in that directory for its RASD database file, with the file name you specified.

rasd_mirror_tran_log

Note The rasd_mirror_tran_log parameter is used *only* by the Replication Agent for Oracle. Replication Agent for DB2 UDB and Microsoft SQL Server do *not* use the rasd_mirror_tran_log parameter.

Enables or disables Replication Agent System Database (RASD) transaction log mirroring.

Default: false

Values: true – enables mirroring the RASD transaction log to another file.

false – disables mirroring of the RASD transaction log.

Comment: Setting the value of the rasd_mirror_tran_log parameter to true provides additional recovery options in the event of a device failure on the Replication Agent host machine.

rasd_trace_log_dir

Note The rasd_trace_log_dir parameter is used *only* by the Replication Agent for Oracle. Replication Agent for DB2 UDB and Microsoft SQL Server do *not* use the rasd_trace_log_dir parameter.

The directory path for the Replication Agent System Database (RASD) trace log file.

Default: The path to the *repository* directory created automatically when the Sybase Replication Agent instance was created. For example:

- On Microsoft Windows platforms:

  %SYBASE%\RAX-15_0*inst_name*\repository

  where:

  - *%SYBASE%* is the path to the Sybase Replication Agent installation directory.
  - *inst_name* is the name of the Replication Agent instance.

- On UNIX platforms:

  $SYBASE/RAX-15_0/*inst_name*/repository

where:

- *$SYBASE* is the path to the Sybase Replication Agent installation directory.
- *inst_name* is the name of the Replication Agent instance.

Value

A valid path on the Replication Agent host machine.

Comments

- When you create a Sybase Replication Agent instance, the *repository* directory is created automatically as part of the instance directory structure. The default value of the rasd_trace_log_dir parameter points to that directory.
- If you specify any valid path as the value of the rasd_trace_log_dir parameter, the Sybase Replication Agent instance writes its RASD trace log file in that directory.

rasd_tran_log

Note The rasd_tran_log parameter is used *only* by the Replication Agent for Oracle. Replication Agent for DB2 UDB and Microsoft SQL Server do *not* use the rasd_tran_log parameter.

The directory path for the Replication Agent System Database (RASD) transaction log file.

Default

The path to the RASD transaction log file created automatically when the Sybase Replication Agent instance was created. For example:

- On Microsoft Windows platforms:

```
%SYBASE%\RAX-15_0\inst_name\repository\inst_name.log
```

where:

- *%SYBASE%* is the path to the Sybase Replication Agent installation directory.
- *inst_name* is the name of the Replication Agent instance.

- On UNIX platforms:

```
$SYBASE/RAX-15_0/inst_name/repository/inst_name.log
```

where:

- *$SYBASE* is the path to the Sybase Replication Agent installation directory.
- *inst_name* is the name of the Replication Agent instance.

Value

A valid path on the Replication Agent host machine.

Comments

- When you create a Sybase Replication Agent instance, the *repository* directory and RASD transaction log file are created automatically. The default value of the rasd_tran_log parameter points to that transaction log file.
- If you specify any valid path and RASD transaction log file name as the value of the rasd_tran_log parameter, the Sybase Replication Agent instance looks in that directory for its RASD transaction log file, with the name you specified.

rasd_tran_log_mirror

Note The rasd_tran_log_mirror parameter is used *only* by the Replication Agent for Oracle. Replication Agent for DB2 UDB and Microsoft SQL Server do *not* use the rasd_tran_log_mirror parameter.

The directory path for the Replication Agent System Database (RASD) transaction log file mirror.

Default

The path to the RASD transaction log file mirror in the *tran_log_mirror* directory created automatically when the Sybase Replication Agent instance was created. For example:

- On Microsoft Windows platforms:

```
%SYBASE%\RAX-15_0\inst_name\repository\tran_log_mirror\inst_name.log
```

where:

- *%SYBASE%* is the path to the Sybase Replication Agent installation directory.
- *inst_name* is the name of the Replication Agent instance.

- On UNIX platforms:

```
$SYBASE/RAX-15_0/inst_name/repository/tran_log_mirror/inst_name.log
```

where:

- *$SYBASE* is the path to the Sybase Replication Agent installation directory.
- *inst_name* is the name of the Replication Agent instance.

Value

A valid path on the Replication Agent host machine.

Comment

If you specify any valid path and transaction log file name as the value of the rasd_tran_log_mirror parameter, the Replication Agent instance looks in that directory for its RASD transaction log file mirror, with the name you specified.

rs_charset

The character set used in communication with the primary Replication Server.

The Sybase Replication Agent default character set must be set to match the primary database's character set. rs_charset must be set to match Replication Server's character set. If they differ, Sybase Replication Agent will do character set conversion before sending data to Replication Server. See the Sybase Replication Agent *Administration Guide* for more information on setting rs_charset.

Note Setting this property to anything other than the character set of the primary Replication Server causes Replication Agent to do character set conversion which will cause problems during replication.

Default

Defaults to empty string ("")

Value

Any valid Sybase character set supported by the Java VM on the Replication Agent host machine.

Comments

- Sybase recommends that you use the exact same name as what is found in the RepServer *configuration* (*.cfg*) file. For example, *iso_1*.
- Sybase recommends that you configure the primary data server and primary Replication Server to use the same character set.

Note If rs_charset is not set at the time you try to resume replication, Sybase Replication Agent returns an error.

When the Sybase Replication Agent instance is created, the rs_charset parameter is set to its default value "" (empty string).

- If you specify a valid character set for the value of the rs_charset parameter, the Sybase Replication Agent instance sends replicated transaction data from the primary database to the primary Replication Server in that character set.
- If you do *not* specify a valid character set name for the value of the rs_charset parameter (including the default rs_charset value ""), the Sybase Replication Agent instance will not allow you to resume replication.
- If the values of the rs_charset and the system default character set are valid, but *not* the same value, Sybase Replication Agent converts the replicated transaction data from the system defined database character set to the Replication Server character set before sending it to the primary Replication Server.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

rs_host_name

The name of the primary Replication Server host machine.

Default: <not_configured>

Value: A valid host name.

Comments:

- The value of the rs_host_name parameter is the name of the host machine for the primary Replication Server.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_packet_size

The network packet size on the connection to the primary Replication Server.

Default: 2048

Value: An integer from 512 to 8192.

Comments:

- The value of the rs_packet_size parameter is the maximum size (in bytes) of the network packets handled by the TCP/IP network protocol.

- The Sybase Replication Agent rs_packet_size parameter is equivalent to the Replication Server rs_packet_size parameter.
- When the network packet size is smaller, more packets must be processed to transmit a given amount of data to the Replication Server. When the network packet size is larger, more system resources are consumed to process the packets.
- The optimum value of the rs_packet_size parameter is based on the nature of the typical data replicated. If the typical operation is very large, a larger packet size is more efficient.
- A larger value of the rs_packet_size parameter is more efficient when the value of the lti_batch_mode parameter is true.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_password

The password that Replication Agent uses for Replication Server access.

Default: "" (empty string)

Value: A valid password.

Comments:

- The value of the rs_password parameter is the password for the user login name that Replication Agent uses to log in to the primary Replication Server.
- The value of the rs_password parameter is encrypted in the Sybase Replication Agent instance configuration file.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

rs_port_number

The client port number of the primary Replication Server.

Default: 1111

Value: A valid port number on the Replication Server host machine.

Comments

- The value of the rs_port_number parameter is the client port number of the primary Replication Server.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_retry_count

The number of times the Replication Agent will retry establishing a connection to the primary Replication Server.

Default

5

Value

An integer greater than 0.

Comments

- The value of the rs_retry_count parameter is the number of times that Sybase Replication Agent will try to establish a network connection to the Replication Server after a connection failure.
- Sybase recommends a setting of 5 for this parameter.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_retry_timeout

The number of seconds the Sybase Replication Agent waits between attempts to connect to the primary Replication Server.

Default

10

Value

An integer greater than 0.

Comments

- The value of the rs_retry_timeout parameter is the number of seconds that the Sybase Replication Agent will wait between its retry attempts to establish a network connection to the primary Replication Server after a connection failure.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_source_db

The name of the database identified in the Replication Server primary database connection.

Default: <not_configured>

Value: A valid database name.

Comments:

- The value of the rs_source_db parameter is the name of the primary database by which the primary Replication Server recognizes the primary database transaction log.
- The value of the rs_source_db parameter must match the name of the database specified in the Replication Server create connection command for the primary database.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_source_ds

The name of the data server identified in the Replication Server primary database connection.

Default: <not_configured>

Value: A valid server name.

Comments:

- The value of the rs_source_ds parameter is the name of the primary data server by which the primary Replication Server recognizes the primary database transaction log.
- The value of the rs_source_ds parameter must match the name of the data server specified in the Replication Server create connection command for the primary database.
- The value of the rs_source_ds parameter should *not* be the same as the name of the Sybase Replication Agent instance.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rs_username

The user login name that the Replication Agent uses for Replication Server access.

Default: <not_configured>

Value: A valid user name.

Comments:

- The value of the rs_username parameter is the user login name that Sybase Replication Agent uses to log in to the primary Replication Server.
- The value of the rs_password parameter is the password for the login name specified by the rs_username parameter.
- The user login name that Sybase Replication Agent uses to log in to the Replication Server must have connect source permission in the Replication Server.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

rssd_charset

The character set used in communication with the RSSD of the primary Replication Server.

Default: "" (empty string)

Value: Any valid Sybase character set supported by the Java VM on the Replication Agent host machine.

Comments:

- The value of the rssd_charset parameter must match (or be compatible with) the RSSD character set. The RSSD character set is usually the same as the Replication Server default character set identified by the Replication Server rs_charset configuration parameter.
- If you specify a valid character set for the value of the rssd_charset parameter, the Sybase Replication Agent instance communicates with the RSSD using that character set.
- If you do *not* specify a valid character set name for the value of the rssd_charset parameter (including the default rssd_charset value ""), the Sybase Replication Agent communicates with the RSSD using the RSSD charset.

- The rssd_charset parameter need not be set if the Sybase Replication Agent use_rssd parameter is set to false.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rssd_database_name

The database name of the RSSD of the primary Replication Server.

Default: <not_configured>

Value: A valid database name.

Comments:

- The value of the rssd_database_name parameter is the database name of the RSSD of the primary Replication Server.
- The rssd_database_name parameter need not be set if the Sybase Replication Agent use_rssd parameter is set to false.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rssd_host_name

The name of the machine on which the RSSD of the primary Replication Server resides.

Default: <not_configured>

Value: A valid host name.

Comments:

- The value of the rssd_host_name parameter is the name of the host machine on which the RSSD of the primary Replication Server resides.
- The rssd_host_name parameter need not be set if the Sybase Replication Agent use_rssd parameter is set to false.
- The value of the rssd_host_name parameter is the name of the host machine on which the RSSD resides.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rssd_password

The password that the Replication Agent uses for access to the RSSD of the primary Replication Server.

Default: "" (empty string)

Value: A valid password.

Comments:

- The value of the rssd_password parameter is the password for the user login name that the Sybase Replication Agent uses to access the RSSD of the primary Replication Server.
- The value of the rssd_password parameter is encrypted in the Sybase Replication Agent instance configuration file.
- The rssd_password parameter need not be set if the Sybase Replication Agent use_rssd parameter is set to false.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

rssd_port_number

The client port number of the Replication Server System Database (RSSD) of the primary Replication Server.

Default: 1111

Value: A valid port number on the RSSD host machine.

Comments:

- The value of the rssd_port_number parameter is the client port number of the RSSD data server.
- The rssd_port_number parameter need not be set if the Sybase Replication Agent use_rssd parameter is set to false.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Sybase Replication Agent connection configuration parameters.

rssd_username

The user login name that the Replication Agent uses to access the RSSD of the primary Replication Server.

Default <not_configured>

Value A valid user login name in the RSSD data server.

Comments
- The value of the rssd_username parameter is the user login name that the Sybase Replication Agent uses to access the RSSD.
- The rssd_username parameter need not be set if the Sybase Replication Agent use_rssd parameter is set to false.
- See the Sybase Replication Agent *Administration Guide* for more information about setting up Replication Agent connection configuration parameters.

scan_sleep_increment

The number of seconds to add to each wait interval before scanning the transaction log, after a previous scan yields no transaction to be replicated.

Default 5

Value An integer from 0 to 3600.

Comments
- The value of the scan_sleep_increment parameter is the number of seconds added to each wait interval before the Log Reader component scans the mirror log device for a transaction to be replicated, after a previous scan yields no such transaction.
- The number of seconds specified by the scan_sleep_increment parameter is added to each wait interval, until the wait interval reaches the value specified by the scan_sleep_max parameter.
- For optimum Sybase Replication Agent performance, the value of the scan_sleep_increment parameter should be balanced with the average number of operations in the primary database over a period of time. In general, better performance results from reading more operations from the transaction log during each Log Reader scan.
- With a primary database that is less frequently updated, increasing the value of the scan_sleep_increment parameter may improve overall performance.
- If the database is continuously updated, the value of the scan_sleep_increment parameter may not be significant to Sybase Replication Agent performance.

scan_sleep_max

The maximum wait interval between Log Reader transaction log scans.

Default: 60

Value: An integer from 5 to 86400.

Comments:

- The value of the scan_sleep_max parameter is the maximum number of seconds that can elapse before the Log Reader component scans the transaction log for a transaction to be replicated, after a previous scan yields no such transaction.
- For reduced replication latency in an infrequently updated database, Sybase recommends lower number settings for the scan_sleep_max parameter.
- If the primary database is continuously updated, the value of the scan_sleep_max parameter is not significant to Replication Agent performance.

skip_ltl_errors

Determines whether the Sybase Replication Agent ignores Log Transfer Language (LTL) error messages.

Warning! Using the skip_ltl_errors parameter incorrectly may cause data inconsistencies between the primary and replicate databases.

Default: false

Values: true – enables skipping LTL errors to continue replication.

false – disables skipping LTL errors.

Comments:

- If the skip_ltl_errors configuration parameter is set to true, the Sybase Replication Agent instance logs any LTL error messages returned by the Replication Server, along with the offending LTL commands, and then it continues processing transaction log records.
- If the skip_ltl_errors configuration parameter is set to false, the Sybase Replication Agent instance stops all of its replication processing and goes to *Admin* state if it receives an LTL error message and the error is unrecoverable.

- The skip_ltl_errors parameter is intended for troubleshooting only, with assistance from Sybase Technical Support.

structured_tokens

Determines whether the Replication Agent uses LTL structured tokens.

Default: true

Values: true – enables LTL structured tokens.

false – disables LTL structured tokens.

Comments:

- If the structured_tokens configuration parameter is set to true, the Log Transfer Interface (LTI) component uses LTL structured tokens when it generates LTL commands.
- Using structured tokens in the LTL can significantly improve overall replication system performance.
- Using structured tokens in the LTL can improve Replication Server performance, especially when non-Sybase datatypes in the primary database must be translated by Replication Server.
- To replicate columns that have one or more spaces in the column name, you must set the value of the structured_tokens parameter to true.

truncation_interval

Specifies a time interval between automatic truncations of the Sybase Replication Agent transaction log.

Warning! If you configure automatic truncation, Sybase Replication Agent for UDB silently deletes the *primary database* log files it no longer needs. For more information, see the Sybase Replication Agent *Primary Database Guide*.

Default: 0

Value: An integer from 0 to 720.

Comments:

- The value of the truncation_interval parameter is the number of minutes between automatic transaction log truncations.

- Automatic transaction log truncation based on the value of the truncation_interval parameter takes place only when the value of the truncation_type parameter is interval.
- The maximum truncation interval is 720 minutes, or 12 hours.
- If the value of the truncation_interval parameter is 0 (zero) and the value of the truncation_type parameter is interval (the default values for both parameters), automatic truncation is disabled.
- To truncate the transaction log manually, use the pdb_truncate_xlog command.
- See the pdb_archive_path, pdb_archive_remove and truncation_type configuration properties.

truncation_type

Configures transaction log truncation behavior of the Sybase Replication Agent.

Warning! For UDB, the Replication Agent deletes the UDB primary database log files that it no longer needs. For more information, see the Sybase Replication Agent *Primary Database Guide.*

Default

interval (for Oracle: locator_update)

Values

command – Sybase Replication Agent truncates the transaction log only when the pdb_truncate_xlog command is invoked.

When the value of the truncation_type parameter is command, the only way you can truncate the transaction log is by invoking the pdb_truncate_xlog command. No automatic truncation takes place when the value of the truncation_type parameter is command.

locator_update – Sybase Replication Agent truncates the transaction log automatically whenever it receives a new LTM Locator value from the primary Replication Server.

When the value of the truncation_type parameter is locator_update, the transaction log will be truncated automatically when Sybase Replication Agent receives a new LTM Locator from the primary Replication Server.

Comments

Note For Oracle, Truncation of the old archive log files from the pdb archival path directory is performed only if the pdb_archive_remove property is *true*.

- Regardless of the value of the truncation_type parameter, the Sybase Replication Agent transaction log can be truncated manually at any time by invoking the pdb_truncate_xlog command.
- If the value of the truncation_interval parameter is 0 (zero) and the value of the truncation_type parameter is interval (the default values for both parameters), automatic truncation is disabled.
- Sybase Replication Agent receives a new LTM Locator based on the values of the lti_update_trunc_point property.
- See the pdb_archive_path, pdb_archive_remove, and truncation_interval configuration properties.

use_rssd

Determines whether the Replication Agent uses replication definitions.

Default

true

Values

true – enables using replication definitions.

false – disables using replication definitions.

Comments

- If the value of the use_rssd parameter is true, the Sybase Replication Agent instance connects to the Replication Server System Database (RSSD) to retrieve replication definitions for the primary database automatically whenever it goes from *Admin* state to *Replicating* state (for example, when the resume command is invoked).
  - Each time it retrieves replication definitions, Sybase Replication Agent stores the information in a cache. Sybase Replication Agent uses replication definitions stored in its cache when it generates Log Transfer Language (LTL) commands.
  - If the Log Transfer Interface (LTI) component encounters an operation on a database object for which it does not have a cached replication definition, Sybase Replication Agent reconnects to the RSSD to update its replication definition cache.
  - If a replication definition still cannot be found for the operation, the Replication Agent instance suspends all of its replication operations and goes to *Admin* state.

- Sybase Replication Agent can use information in table and function replication definitions (that is, replication definitions for individual primary database objects) stored in the RSSD to generate more efficient LTL, and thus improve throughput in the LTI component and the Replication Server.

  Accessing replication definitions in the RSSD enables the LTI component to improve performance by:

  - Omitting column names in LTL – When columns are sent in the order specified in the replication definition, column images can be sent without column names (headings), which reduces LTL overhead.
  - Omitting unneeded columns in LTL – When columns are sent as specified in the replication definition, images for unchanged columns need not be sent, which reduces LTL overhead.
  - Sending data for each column in the datatype specified by the replication definition – This allows data to be handled more efficiently all the way through the replication system.
  - Sending database object names in the same character case as defined in the replication definition.

- If the value of the use_rssd parameter is false, none of the previously described performance improvements are possible. In that case, the Replication Agent sends all data as a char datatype in the LTL.
- If you use owner-qualified table names for either primary tables or replicate tables, you must:
  - Set the value of the use_rssd parameter to true
  - Specify an owner-qualified primary table name and/or replicate table name when you create the replication definition in the primary Replication Server

Glossary

This glossary describes Replication Server and Sybase Replication Agent terms used in this book.

Adaptive Server

The brand name for Sybase relational database management system (RDBMS) software products.

- *Adaptive Server Enterprise* manages multiple, large relational databases for high-volume online transaction processing (OLTP) systems and client applications.
- *Adaptive Server IQ* manages multiple, large relational databases with special indexing algorithms to support high-speed, high-volume business intelligence, decision support, and reporting client applications.
- *Adaptive Server Anywhere* manages relational databases with a small RDBMS footprint, which is ideal for embedded applications and mobile device applications.

See also **database** 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. It provides an alternative to using SQL insert and select commands to transfer data. See also **BCP utility** and **materialization**.

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**, **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. Contrast with **rollback**. See also **DBMS** and **transaction**.

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 data to remote locations, and then keeping the replicated data synchronized with the primary data. Data replication is distinct from data distribution. Replicated data is stored copies of data at one or more remote sites throughout a system, and it is not necessarily distributed data. Contrast with **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 be any data repository with the interface and functionality a data client requires. See also **client**, **client application**, and **data client**.

database A collection of data with a specific structure (or schema) for accepting, storing, and providing data for users. See also **data server** and **relational database**.

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 defined in Replication Server. See also **Replication Server** and **route**.

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 DBMS is a computer-based system for defining, creating, manipulating, controlling, managing, and using databases. The DBMS can include the user interface for using the database, or it can be a stand-alone data server system. Compare with **RDBMS**. See also **database**.

function A Replication Server object that represents a data server operation, such as insert, delete, or begin transaction. Replication Server distributes operations to replicate databases as functions. See also **function string**.

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 replication between (homogeneous) non-Sybase data servers, and 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, Sybase Replication Agent, and Replication Server. See also **Open Client** and **Open Server**.

Java An object-oriented, platform-independent, "write once, run anywhere" programming language developed by Sun Microsystems. The Sybase Replication Agent is a Java application.

Java VM The Java Virtual Machine. The Java VM (or JVM) is the part of the Java Runtime Environment (JRE) that interprets 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 **client**, **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. To run a Java application, such as the Sybase Replication Agent, a JRE must be installed on the machine. See also **Java** and **Java VM**.

LAN An abbreviation for *local area network*, a computer network located on the user's 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. See also **transaction replication**.

LOB An abbreviation for *large object*. A LOB is a type of data element (or datatype) associated with a column that contains extremely large quantities of data.

Log Reader An internal component of the 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 the Replication Agent that interacts with Replication Server to forward transactions for distribution to a replicate database. See also **Log Reader** and **Log Transfer Manager**.

Log Transfer Manager An internal component of the 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 **non-atomic materialization**.

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*. ODBC is 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 version of a set of data that is the source used for replication. Primary data is stored and managed by the primary database. See also **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 transactions and data in a replication system. Sometimes called the *active database*. Contrast with **replicate database**. See also **primary data** and **replicated transaction**.

primary key — The column or columns whose data uniquely identify each row in a table.

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** and **transaction log**.

quiesce An action that causes 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 data-changing operations have been propagated to their destinations. Some Replication Server commands require that you quiesce the replication system.

In a database, a state in which all data-changing operations are suspended so that transactions cannot change any data.

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*. RCL is the command language used to manage Replication Server. See also **Replication Server**.

RDBMS An abbreviation for *relational database management system*. An RDBMS is 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 The data managed by a replicate database, which is the destination (or target) of a replication system. Contrast with **primary data**. See also **replicate database** and **replication system**.

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 transactions and/or data in a replication system. Sometimes called the *standby database*. Contrast with **primary database**. See also **replicate data**, **replicated transaction**, and **replication system**.

replicated 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**.

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**.

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**, **replicate 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 The Sybase software product that provides the infrastructure for a robust transaction replication system. See also **Replication Agent**.

RSSD An abbreviation for *Replication Server System Database*. The RSSD manages replication system information for a Replication Server. See also **Replication Server**.

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 **data replication** and **transaction replication**.

rollback An instruction to a database to reverse the data changes requested in a unit of work (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**.

SQL An abbreviation for *Structured Query Language*. SQL is a non-procedural 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**.

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 database, 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 to the table. See also **database** and **relational 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 **replicated transaction** and **SQL**.

transaction log — Generally, the log of transactions that affect the data managed by a database or a data server. Replication Agent reads the transaction log to identify and acquire the transactions to be replicated from the primary database. See also **primary database**, **Replication Agent**, and **transaction**.

transaction replication — A data replication method that copies data-changing operations from a primary database to a replicate database. See also **data replication**, **primary database**, and **replicate database**.

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. See also **primary database**, **replicate database**, and **transaction**.

WAN — An abbreviation for *wide area network*, a system of local-area networks (LANs) connected together with data communication lines. Contrast with **LAN**.

Index

D

E

F

G

H

M

N

O

P

Q

R

S

T

U

V

W