



Troubleshooting Guide

Replication Server®

12.6

DOCUMENT ID: DC35920-01-1260-02

LAST REVISED: March 2006

Copyright © 1992-2006 by Sybase, Inc. All rights reserved.

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

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

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase, the Sybase logo, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Warehouse, Afaia, Answers Anywhere, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, AvantGo Mobile Delivery, AvantGo Mobile Inspection, AvantGo Mobile Marketing Channel, AvantGo Mobile Pharma, AvantGo Mobile Sales, AvantGo Pylon, AvantGo Pylon Application Server, AvantGo Pylon Conduit, AvantGo Pylon PIM Server, AvantGo Pylon Pro, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convoy/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, Developers Workbench, DirectConnect, DirectConnect Anywhere, Distribution Director, e-ADK, E-Anywhere, e-Biz Impact, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTIP, eFulfillment Accelerator, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, M2M Anywhere, Mach Desktop, Mail Anywhere Studio, Mainframe Connect, Maintenance Express, Manage Anywhere Studio, M-Business Anywhere, M-Business Channel, M-Business Network, M-Business Suite, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, mFolio, Mirror Activator, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, Pharma Anywhere, PocketBuilder, Pocket PowerBuilder, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, QAnywhere, Rapport, RemoteWare, RepConnector, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Report-Execute, Report Workbench, Resource Manager, RFID Anywhere, RW-DisplayLib, RW-Library, Sales Anywhere, SDF, Search Anywhere, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, smart.partners, smart.parts, smart.script, SOA Anywhere, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase IQ, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SybFlex, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TradeForce, Transact-SQL, Translation Toolkit, UltraLite, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server, XcelleNet, and XP Server are trademarks of Sybase, Inc. 10/05

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

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

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

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

Contents

About This Book	vii
CHAPTER 1 Identifying the Problem	1
Overview	1
Tasks or events that can cause replication system problems ...	2
Reduced Replication System performance	3
Types of problems	5
Manual recovery problems	9
Analyzing error logs	10
rs_init	11
Replication Server error log	11
LTM error Log	12
Diagnostic tools	12
Troubleshooting configuration problems	13
Troubleshooting subscription failures	14
Materialization	15
Dematerialization	16
Troubleshooting replication failures	17
Identify the data that is failing to replicate	18
Verify that Replication Server threads are up	19
Replication Server is down	21
Verify that it is not a queue problem	21
Verify that all RepAgents and LTMs are up	21
Whether only text and image columns are not replicated	22
Check system tables	22
Activate traces and dump queues	23
Verify that no route problems exist	23
Troubleshooting manual recovery problems	24
CHAPTER 2 Analyzing Error Messages	25
Overview	25
What is not covered	26
Replication Server errors	26

	Error message format.....	27
	Analyzing Replication Server error messages	29
	RepAgent error messages	31
	RepAgent error message format	32
	Analyzing RepAgent error messages.....	33
CHAPTER 3	Common Errors	35
	Finding the error description	35
	Error description format.....	36
	rs_init error message descriptions	36
	Replication Server error message descriptions	37
	Correcting Adaptive Server error.....	38
	7035	44
	11061	45
	Replication Server informational and warning messages	53
	A cached row for system table was swapped out	53
	Detecting loss for database.....	54
	DSI detected rs_update_lastcommit not marked as replicated	56
	Stable storage use is above 75 percent.....	56
	RepAgent error message descriptions.....	57
	9202: nested stored procedures not allowed (LTM 1033).....	57
	9210: network problems (LTM 2027)	58
	9215: log truncated (LTM 1079 with ASE 624)	59
	9220: connect only one RepAgent to a Replication Server (LTM 1078)	60
	LTM-only error message descriptions.....	61
	2028: invalid enable marker in log.....	61
CHAPTER 4	Route Problems	63
	Overview	63
	Using the rs_helproute stored procedure	64
	Problems with creating routes.....	65
	Common problems	66
	Messages in the error log at the source Replication Server....	66
	Output from rs_helproute at the source Replication Server	67
	Output from rs_helproute at the destination server	69
	Problems with altering routes.....	70
	Problems with dropping routes.....	70
	Dropping routes.....	70
	Dropping routes using the nowait option	73
CHAPTER 5	Subscription Problems	75

	Overview	75
	Using check subscription	76
	Materialization problems	79
	Common problems	80
	Atomic materialization	82
	Nonatomic materialization	87
	Bulk materialization	92
	Dematerialization problems.....	95
	with purge dematerialization.....	96
	Bulk dematerialization	99
CHAPTER 6	Replication Server Interface Problems	103
	Overview	103
	Connection errors.....	104
	Incorrect RSI user login name or password	104
	Incorrect user permissions at replicate Replication Server ...	105
	Invalid RSI locator	106
CHAPTER 7	RepAgent and Log Transfer Manager Problems.....	107
	Overview	107
	Problems when starting the RepAgent or LTM	108
	Invalid login	109
	Invalid permissions	109
	Problems restarting the LTM after it disconnects	110
	Errors from the Replication Server.....	110
	Normalization errors	111
	Errors from the Adaptive Server.....	114
	Dropping primary objects	114
	Invalid truncation page	115
	Running out of Adaptive Server alarm resources.....	116
CHAPTER 8	Data Server Interface Problems.....	117
	Overview	117
	Listing databases and using admin who and admin who, dsi	118
	Listing databases controlled by a Replication Server.....	118
	Checking the DSI status for a database connection	119
	Troubleshooting procedure	120
	DSI down or suspended.....	121
	Connection failure to the database.....	122
	Data server errors	123
	Replication Server errors.....	125
	DSI active.....	126

	Incorrect duplicate transaction resolution.....	126
	Examining the exceptions log.....	127
CHAPTER 9	Adaptive Server Log Problems	129
	Overview	129
	Truncating an Adaptive Server log.....	130
	Verifying the state of the secondary truncation point	130
	Turning off the secondary truncation point in a database	131
	Setting the secondary truncation point.....	132
	Resetting the locator	132
	Log truncation problems.....	133
	Orphaned transactions	134
	Inbound queue requires more disk space	135
	Adaptive Server log corruption	136
APPENDIX A	Dumping Stable Queues	137
	Introduction	137
	Using trace	138
	Using sysadmin dump_queue.....	139
	Confirming suspected problems.....	141
	sysadmin dump_queue examples.....	142
	Interpreting the output of a dump queue	143
	Using sysadmin sqt_dump_queue	150
	Using sysadmin dump_file	152
	sysadmin purge commands	153
APPENDIX B	Language, Sort Order, and Character Set Issues.....	155
	Overview	155
	Message language.....	156
	Stored procedure messages	156
	Other situations that produce mixed languages	156
	Sort order	157
	Sort order problems.....	157
	Character sets.....	158
	Index	159

About This Book

Replication Server® maintains replicated data at multiple sites on a network. Organizations with geographically distant sites can use Replication Server to create distributed database applications with better performance and data availability than a centralized database system can provide.

This book, *Replication Server Troubleshooting Guide*, describes Replication Server troubleshooting tasks, such as dumping stable queues, debugging failures with the Data Server Interface (DSI) and Replication Server Interface (RSI), and diagnosing and correcting problems with subscriptions.

Note Use the information in this manual for system-level troubleshooting. For replication agent–specific troubleshooting information, see the replication agent documentation for your data server.

Audience

The *Replication Server Troubleshooting Guide* is for Replication Server System Administrators and ap

How to use this book

This book contains:

- Chapter 1, “Identifying the Problem” describes how to identify configuration, subscription, replication, and manual recovery problems.
- Chapter 2, “Analyzing Error Messages” describes how to analyze Replication Server, RepAgent, and Log Transfer Manager (LTM) error messages.
- Chapter 3, “Common Errors” describes the most common errors that can occur during configuration and operation of a replication system.
- Chapter 4, “Route Problems” explains how to troubleshoot routing problems.
- Chapter 5, “Subscription Problems” helps you resolve common problems with subscriptions.

-
- Chapter 6, “Replication Server Interface Problems” helps you resolve problems with the RSI.
 - Chapter 7, “RepAgent and Log Transfer Manager Problems” helps you troubleshoot problems with RepAgent and the LTM.
 - Chapter 8, “Data Server Interface Problems” provides troubleshooting information for the DSI.
 - Chapter 9, “Adaptive Server Log Problems” describes how and when to truncate the Adaptive Server transaction log and other log problems.
 - Appendix A, “Dumping Stable Queues” explains how to dump Replication Server stable queues.
 - Appendix B, “Language, Sort Order, and Character Set Issues” discusses problems associated with running a replication system in an international environment.

Related documents

The Sybase® Replication Server® documentation set consists of the following:

- The release bulletin for your platform – contains last-minute information that was too late to be included in the books.

A more recent version of the 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.
- *Installation Guide* for your platform – describes installation and upgrade procedures for all Replication Server and related products.
- *What’s New in Replication Server?* – describes the new features in Replication Server version 12.6 and the system changes added to support those features.
- *Administration Guide* – contains an introduction to replication systems. This manual includes information and guidelines for creating and managing a replication system, setting up security, recovering from system failures, and improving performance.
- *Configuration Guide* for your platform – describes configuration procedures for all Replication Server and related products, and explains how to use the rs_init configuration utility.
- *Design Guide* – contains information about designing a replication system and integrating heterogeneous data servers into a replication system.

- *Getting Started with Replication Server* – provides step-by-step instructions for installing and setting up a simple replication system.
- *Heterogeneous Replication Guide* – describes how to use Replication Server to replicate data between databases supplied by different vendors.
- *Reference Manual* – contains the syntax and detailed descriptions of Replication Server commands in the Replication Command Language (RCL); Replication Server system functions; Sybase Adaptive Server® commands, system procedures, and stored procedures used with Replication Server; Replication Server executable programs; and Replication Server system tables.
- *System Tables Diagram* – illustrates system tables and their entity relationships in a poster format. Available only in print version.
- *Troubleshooting Guide* – contains information to aid in diagnosing and correcting problems in the replication system.
- Replication Server plug-in help, which contains information about using Sybase Central™ to manage Replication Server.

Other sources of information

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

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

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

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

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

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

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

❖ Finding the latest information on product certifications

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

❖ Finding the latest information on component certifications

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

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

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

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

Sybase EBFs and software maintenance

❖ Finding the latest information on EBFs and software maintenance

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

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

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

Conventions

This section describes style and syntax conventions, RCL command formatting conventions, and graphic icons used in this book.

Style conventions Syntax statements (displaying the syntax and options for a command) are printed as follows:

```
alter user user
set password new_passwd
[verify password old_passwd]
```

See “Syntax conventions” on page xii for more information.

Examples that show the use of Replication Server commands are printed as follows:

```
alter user louise
set password somNific
verify password EnnuI
```

Command names, command option names, program names, program flags, keywords, configuration parameters, functions, and stored procedures are printed as follows:

Use `alter user` to change the password for a login name.

Variables, parameters to functions and stored procedures, and user-supplied words are in italics in syntax and in paragraph text, as follows:

The set password *new_passwd* clause specifies a new password.

Names of database objects, such as databases, tables, columns, and datatypes, are in italics in paragraph text, as follows:

The *base_price* column in the *Items* table is a money datatype.

Names of replication objects, such as function-string classes, error classes, replication definitions, and subscriptions, are in italics, as follows:

rs_default_function_class is a default function-string class.

Syntax conventions Syntax formatting conventions are summarized in the following table. Examples combining these elements follow.

Key	Definition
<i>variable</i>	Variables (words standing for values that you fill in) are in italics.
{ }	Curly braces mean you must choose at least one of the enclosed options. Do not include braces in the command.
[]	Brackets mean you may choose or omit enclosed options. Do not include brackets in the command.
	Vertical bars mean you may choose no more than one option (enclosed in braces or brackets).
,	Commas mean you may choose as many options as you need (enclosed in braces or brackets). Separate your choices with commas, to be typed as part of the command. Commas may also be required in other syntax contexts.
()	Parentheses are to be typed as part of the command.
...	An ellipsis (three dots) means you may repeat the last unit as many times as you need. Do not include ellipses in the command.

Obligatory choices

- Curly braces and vertical bars – choose only one option.
`{red | yellow | blue}`
- Curly braces and commas – choose one or more options. If you choose more than one, separate your choices with commas.

`{cash, check, credit}`

Optional choices

- One item in square brackets – choose it or omit it.
`[anchovies]`
- Square brackets and vertical bars – choose none or only one.

```
[beans | rice | sweet_potatoes]
```

- Square brackets and commas – choose none, one, or more options. If you choose more than one, separate your choices with commas.

```
[extra_cheese, avocados, sour_cream]
```

Repeating elements

An ellipsis (...) means that you may repeat the last unit as many times as you need. For the alter function replication definition command, for example, you can list one or more parameters and their datatypes for either the add clause or the add searchable parameters clause:

```
alter function replication definition function_rep_def
{deliver as 'proc_name' |
add @parameter datatype [, @parameter
datatype]... |
add searchable parameters @parameter
[, @parameter]... |
send standby {all | replication definition}
parameters}
```

RCL command formatting

RCL commands are similar to Transact-SQL® commands. The following sections present the formatting rules.

Command format and command batches

- You can break a line anywhere except in the middle of a keyword or identifier. You can continue a character string on the next line by typing a backslash (\) at the end of the line.
- Extra space characters on a line are ignored, except after a backslash. Do not enter any spaces after a backslash.
- You can enter more than one command in a batch, unless otherwise noted.
- RCL commands are not transactional. Replication Server executes each command in a batch without regard for the completion status of other commands in the batch. Syntax errors in a command prevent Replication Server from parsing subsequent commands in a batch.

Case sensitivity

- Keywords in RCL commands are not case-sensitive. You can enter them with any combination of uppercase or lowercase letters.
- Identifiers and character data may or may not be case-sensitive, depending on the sort order that is in effect.
 - If you are using a case-sensitive sort order such as “binary,” you must enter identifiers and character data using the correct combination of uppercase and lowercase letters.
 - If you are using a sort order that is not case-sensitive, such as “nocase,” you can enter identifiers and character data using any combination of uppercase or lowercase letters.

Identifiers

- Identifiers are names you give to database or replication objects. They include object names, column names, variable names, and parameter names.
- Identifiers can be 1–30 bytes long (equivalent to 1–30 single-byte characters). The first character must be a letter or the @ or _ character.
- Replication Server function parameters are the only identifiers that can begin with the @ character. Function parameter names can include up to 30 characters *after* the @ character.
- After the first character, identifiers can include letters, digits, and the #, \$, or _ character. You cannot embed spaces in identifiers.

Parameters in function strings

- Parameters in function strings have the same rules as identifiers, except:
 - They are enclosed in question marks (?), allowing Replication Server to locate them in the function string. Use two consecutive question marks (??) to represent a literal question mark in a function string.

- The exclamation point (!) introduces a parameter modifier that indicates the source of the data to substitute for a parameter at runtime. See Chapter 14, “Customizing Database Operations,” in the *Replication Server Administration Guide* for a complete list of modifiers.

Datatype support

Replication Server supports all of the Sybase® Adaptive Server datatypes.

User-defined datatypes are not supported. The timestamp, double precision, nchar, and nvarchar datatypes are indirectly supported by mapping them to other datatypes. Columns using the timestamp datatype are mapped to varbinary(8).

Refer to “Topics” in Chapter 2, “Topics,” in the *Replication Server Reference Manual* for more information about the supported datatypes, including how to format them.

Graphic icons

Illustrations in this book use the following graphic icons to represent the components of a replication system.

Replication Server



This icon represents Replication Server, the Sybase server program that maintains replicated data on a local-area network (LAN) and processes data transactions received from other Replication Servers on the wide-area network (WAN).

Adaptive Server or other data servers



This icon represents Adaptive Server, the Sybase data server. Data servers manage databases containing primary or replicate data. Replication Server works with heterogeneous data servers, so unless specifically noted, the graphic can represent any data server in a replication system.

Replication Agent



This icon represents RepAgent, LTM, or another replication agent.

Client application



This icon represents a client application. A client application is a user process or application connected to a data server. It may be a front-end application program executed by a user or a program that executes as an extension of the system.

Replication Server Manager



This icon represents Replication Server Manager™ (hereafter referred to as RSM), which is a Sybase Central graphical administration tool that enables a Replication System Administrator to monitor, diagnose, troubleshoot, and administer a replication system.

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.

This document has 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.

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.

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.

Identifying the Problem

This chapter describes the types of problems that can occur in Replication Server, procedures you can use to identify your specific Replication Server problem, and the diagnostic tools you can use to solve the problem.

Name	Page
Overview	1
Manual recovery problems	9
Analyzing error logs	10
Diagnostic tools	12
Troubleshooting configuration problems	13
Troubleshooting subscription failures	14
Troubleshooting replication failures	17
Troubleshooting manual recovery problems	24

Overview

A correctly configured Replication Server system is designed to be fault-tolerant. However, in the event of a serious failure, it might be necessary to manually intervene and fix the problem. This guide will help you locate, identify, and fix the cause of the problem.

You can solve most Replication Server problems by following the procedures in this guide. The key to finding the cause of a replication system failure is to eliminate possible causes. Use the following troubleshooting methods to narrow down the possible causes:

- Identify recent tasks or events that might have adversely affected the replication system. You or another coworker could have performed a task that caused a replication failure or reduced the performance of the replication system. Many things can happen when a number of people are working on the same replication system.

An event such as a temporary network outage may also have caused a replication system problem.

- Analyze error logs.
- Determine the general problem category (configuration, subscription, replication, recovery).
- Use diagnostic tools, such as Replication Server programs or stored procedures, or isql, to analyze the replication system.

If you see an error message displayed in a Replication Server error log, you can identify the problem by reading the error log and referring to Chapter 2, “Analyzing Error Messages” to learn how to read error log files, Chapter 3, “Common Errors” to see if the error is a common one, or to one of the other chapters that is relevant to the symptoms of your problem.

If you do not see an error message, you must use diagnostic tools to further analyze the replication system.

Replication Server Manager (RSM) problems are not covered in this guide. See the online help for Sybase Central and *Replication Server Administration Guide* for information about fixing problems related to running RSM; see the configuration guide for your platform for information about fixing RSM configuration problems.

This guide may be able to help you to identify hardware, network, and operating system problems, but solving these problems is beyond the scope of this book. Any time a server or a network connection is down, you should also check for hardware, network, or operating system problems.

On Windows NT, you can usually see a hardware or operating system problem when stack traces randomly or frequently occur at the same time you get errors in the Replication Server error log.

Check the operating system error log for errors that indicate hardware or operating system problems. Such failures might only partially resolve the effects on the replication system. You may still need to resynchronize data between the primary and destination databases.

Tasks or events that can cause replication system problems

Before attempting to categorize your problem, see if you have performed any of the following tasks or if any of these events have occurred. These tasks or events could have caused a replication system problem.

- If you changed binaries, you might have made the Replication Server, LTM/RepAgent, or Adaptive Server incompatible with each other. Incompatible binaries can produce Open Server™ and Open Client™ error messages. Check the latest release bulletins for version compatibility between replication system components.
- Network went down and restarted. See “Verify that Replication Server threads are up” on page 19.
- Rebuilt queues. See “Troubleshooting manual recovery problems” on page 24.
- Recovered primary database. See “Troubleshooting manual recovery problems” on page 24.
- Added, altered, or deleted routes. The route operation might not be complete. See Chapter 4, “Route Problems”.
- Added, altered, or deleted subscriptions. The subscription operation might not have completed. See Chapter 5, “Subscription Problems”.
- Added, altered, or deleted Replication Server, LTM/RepAgent, primary or replicate database, data server, or table. You might have changed the replication system incorrectly. See “Troubleshooting configuration problems” on page 13.

Reduced Replication System performance

Reduced performance usually occurs when the operating system or machines are overloaded as a result of increased demands on the system. Reduced performance can result from adding applications or Replication Server components, executing very large transactions, or even upgrading the operating system.

Performance problems can be critical or non-critical. You can eliminate non-critical performance problems by optimizing your replication system. For non-critical performance problems, see the *Replication Server Administration Guide*.

If critical performance problems are left unresolved long enough, performance degradation can lead to fatal problems, such as full stable queues (see “Poor performance and stable queue size increasing” on page 4), in which replication stops. A replication, materialization, or dematerialization failure can be caused by a critical performance problem.

Critical performance problems can stem from:

- Introducing new components, such as Adaptive Servers, databases, Replication Servers, LTMs/RepAgents, or replication agents, may cause resource contention and overload any component.
- Changing the operating system—such as upgrading the operating system, applying patches, changing kernel parameters, or rebuilding the kernel—may adversely impact your replication system, memory allocation, and resources.
- Adding applications to the replication system may impact memory requirements and use resources.
- If you are replicating a very large database, the latency might be very high. You may also have a very large transaction or an open transaction.

Poor performance and stable queue size increasing

Sometimes poor Replication Server performance is caused by the Stable Queue Transaction (SQT) cache size being too small. If the SQT cache is not big enough to hold all open transactions, transactions that cannot fit completely into the cache are processed one command at a time.

The inbound stable queues grow (more than 1MB) because Replication Server cannot process the transactions fast enough.

To solve this problem:

- 1 Make sure the problem is not caused by orphaned transactions. For more information about orphaned transactions, see “Log truncation problems” on page 133.
- 2 Find the optimal size for the *sqt_max_cache_size* and *sqt_max_prs_size* parameters for your application. You may want to increase the Replication Server default from 131072 bytes to 256 KB (262144 bytes). *sqt_max_prs_size* determines the memory available for pre-parsing of messages, a larger size will give a performance improvement. It should be at least as large as *sqt_max_cache_size*.
- 3 Check the SQT cache size. If any of the following conditions are false, you may have an SQT cache that is too small. Check the *sqt_max_cache_size* and *sqt_max_prs_size* parameters:
 - In the admin who,sqt output, the “Removed” field should always be 0.
 - In the admin who,sqt output, if the “Full” field is 1, the “Closed” field should have a value other than 0.

- In the admin who, sqm output, the “Next Read” field should always be increasing or be greater or equal to the “Last Seg.Block” field.
- 4 As long as these conditions above are not fulfilled, keep increasing the values for *sqt_max_cache_size* and *sqt_max_prs_size* parameters in the RSSD.
 - 5 Try to set the *sqt_max_cache_size* and *sqt_max_prs_size* parameters so that on the average about 20 transactions are closed at any one time. More than 20 closed transactions indicates the cache is larger than required, fewer than 20 indicates the cache is not large enough.
 - a Suspend the connection to the data server.
 - b Log in to Replication Server and run these commands:

```
> configure replication server
    set sqt_max_cache_size to 'new_value'
> go
> configure replication server
    set sqt_max_prs_size to 'new_value'
> go
```
 - c To activate the new values, resume the connection. If you want to dump the SQT cache, use the `sysadmin sqt_dump_queue` command.

Types of problems

This section describes the types of problems that can occur with Replication Server. These types roughly correspond to the different stages of development of the replication system. A replication system consists of Replication Server components (Replication Servers, replication agents, data servers, routes, connections) connected together such that data is copied reliably from source tables to destination tables.

- Configuration problems – occur while you are setting up Replication Server such as adding Replication Servers, RepAgents, and data servers.
- Subscription problems – occur while you are materializing or dematerializing subscriptions.
- Replication problems – occur when updates to a primary database are being copied to a replicate database.
- Manual recovery problems – occur when you are recovering a failed partition, truncated primary database log, primary database failure, or Replication Server System Database (RSSD).

Configuration problems

Configuration errors occur when you are running `rs_init`, the Replication Server program used to add and configure Replication Servers, RepAgents, and data servers to the replication system.

Usually, symptoms and problems are identified by error messages in the `rs_init` log file. After using the procedure in Chapter 2, “Analyzing Error Messages”, you may be able to resolve the problem by looking for the solution in Chapter 3, “Common Errors”.

Some configuration problems cause subscription materialization failures, and their symptoms do not appear until you attempt to materialize a subscription.

For more information about Replication Server configuration problems, see “Troubleshooting configuration problems” on page 13.

Subscription problems

The replication process begins with subscription *materialization*, which is the process by which data is first copied to the destination database. When you no longer want a subscription replicated to a destination database, you dematerialize the subscription at the destination database. *Dematerialization* is the process by which data is deleted from the destination database.

Subscription failures occur when materialization or dematerialization fails.

Usually, the person who is conducting the materialization or dematerialization monitors the operation and reports any problems.

Subscription problem symptoms are easily identified and include:

- Materialization failure – no data in the subscription’s replicate table at the destination database, the status for subscriptions at the primary and replicate Replication Servers is not Valid, or materialization has been taking longer than is reasonable.
- Dematerialization failure – data still exists in the subscription’s replicate table at the destination database, the status for subscriptions at the primary and replicate Replication Servers is not Invalid, or dematerialization has been taking longer than is reasonable.

Some subscription problem symptoms are displayed as error messages in the Replication Server error log. You might also need to use the diagnostic tools (12) to identify subscription problem symptoms. For more information about subscription problems, see “Troubleshooting subscription failures” on page 14.

Using incorrect variables in function strings

If you are having problems with customized function strings, which can be used to replicate changes, you might have used incorrect variables. Here are the restrictions:

- Pre-11.0.2 release Replication Server function strings can use system-defined variables. Function-string class scope functions cannot use the defaults.
- Only function strings for `rs_insert` and `rs_update` can use new column values.
- Only function strings for `rs_delete` and `rs_update` can use old column values.
- Only function strings for `rs_select` and `rs_select_with_lock` can have input templates. Only they can use user-defined variables.
- Only function strings for user-defined functions can use parameter values of functions. The parameter value of a function consists of the parameters passed to a replicated stored procedure.

Replication problems

Replication consists of copying data operations, such as updating or deleting data, from a primary database to the destination database. Replication begins after a subscription has been successfully materialized. Replication problems occur when data is changed at a primary database but not at the destination database.

You can find that a replication problem exists by observing that data is not replicating to a destination database or by hearing from someone that their application is retrieving incorrect data from a destination database.

If you are monitoring the replication system, you might directly observe that data is not replicating to a destination database. Use `rs_subcmp` to determine which subscription is not being replicated.

You should consider that a replication problem may exist if someone reports that their client application retrieved incorrect data from a destination database. Compare the primary and destination tables; if they are the same, then data is being replicated correctly. There is probably a problem with the client application that is causing incorrect data to be displayed in the client application. If data is not the same at the primary and destination databases, replication is failing, and you must troubleshoot the replication system.

Some symptoms of a replication problem directly identify the cause; other symptoms require more investigation to find the underlying cause. These symptoms can be organized into the following types in order of the most common first.

- Data Server Interface (DSI) thread is down.
- Threads other than DSI are down.

You use `admin who_is_down` to display information about threads that are down.

You use `admin who_is_down` to display information about threads that are down.

- DIST thread
- RepAgent User thread
- RSI thread
- RSI User thread
- RS User thread
- SQM thread
- SQT thread
- Major replication system component is down.

Use `isql` to check to see if a server is down by logging in to each server.

- RepAgent
- LTM
- Replication Server
- Data server
- A “detecting loss” message, which means that data replication messages were lost after queues were rebuilt. This message is shown in the Replication Server error log or in the `rs_oqid` system table.
- Inbound or outbound stable queues are growing larger.

Use `admin who`, `sqm` and `sysadmin dump_queue` to display information about inbound and outbound stable queues.

- Number of duplicate transactions is increasing.

Use `admin who,sqt` and `sysadmin dump_queue` to display information about inbound and outbound stable queues.

- Transactions remain open for longer than is reasonable. These transactions might be orphans or a very long transaction. Orphans are transactions that do not have an ending commit or rollback statement.

Use `admin who,sqt` and `sysadmin dump_queue` to display information about inbound and outbound stable queues.

- Primary and destination Replication Servers do not have the same locator.
Use `isql` to login to the RSSD and view the `rs_locator` system table.
- Replication is successful for other subscriptions on different data servers with connections to the same destination Replication Server.

Use `rs_subcmp` to compare a subscription's tables in the primary and replicate databases to make sure the tables are the same.

- Replication is successful for other subscriptions in the same or different tables on the same data server while replication for a particular subscription is failing

Use `rs_subcmp` to compare a subscription's tables in the primary and replicate databases to make sure the tables are the same.

Some symptoms are displayed as error messages in the Replication Server, LTM, and Adaptive Server error logs. Use the diagnostic tools (12) to identify replication problem symptoms. For more information about replication problems, see “Troubleshooting replication failures” on page 17.

Manual recovery problems

While Replication Server is designed to tolerate most failures and to recover from them automatically, some failures require you to intervene by manually executing recovery tasks. Sometimes after you have completed a recovery task, you run into replication problems or errors show up in the Replication Server error log. Although the *Replication Server Administration Guide* covers most recovery situations and you can refer to that manual to see if you missed a step in a recovery task, this guide covers the most common problems that you might run into after you think that you have successfully completed a recovery task.

For information about some common manual recovery problems, see “Troubleshooting manual recovery problems” on page 24.

Analyzing error logs

Always maintain open windows for all error logs, including:

- Replication Server error logs (including a window for Replication Server standard error output (*stderr*))
- Log transfer manager (LTM) error logs
- Replication agent logs (such as Replication Agent™ for DB2)
- Any data server error logs (such as Adaptive Server)
- Any gateway error logs (such as DirectConnect™ or OmniConnect™)

Usually an error message will be displayed in one of these logs. Before checking for failures throughout your replication system, check for error messages in the error logs:

- 1 Scan the replication agent, data server, and gateway error logs to see if you can immediately find an error message that describes the cause of the error. If you find an error, see the documentation for the replication agent, data server or gateway to solve the problem.
- 2 Check for error messages in the *rs_init*, Replication Server, and LTM error logs following the instructions in Chapter 2, “How to Read Error Log Files.”

Note RepAgent error messages are displayed in the Adaptive Server error log.

- 3 Find the error message in Chapter 3, “Common Errors.”
- 4 Apply the corresponding repair procedure.
- 5 If you cannot find the error message in Chapter 1, “Identifying the Problem”, look in one of the other chapters that corresponds to the symptoms of your problem or in one of the following error message text files in the *\$\$SYBASE/doc* directory on UNIX platforms or the *\$\$SYBASE\doc* directory on Windows NT. These text files contain the text of the error message, a brief description of the error, its cause, and, sometimes, how to fix it.

For Windows NT:

- *errmsg.rs* for Replication Server error messages
- *errmsg.rsm* for RSM error messages

- *errmsg.ltm* for LTM error messages
- *errmsg.sub* for rs_subcmp error messages

For UNIX:

- *error_messages_rs* for Replication Server error messages
- *error_messages_rsmsrvr* for RSM error messages
- *error_messages_ltm* for LTM error messages
- *error_messages_subcmp* for rs_subcmp error messages

rs_init

rs_init error messages are recorded in the rs_init error log. The rs_init error log is located in:

- The *\$\$SYBASE/init/logs* directory for UNIX
- The *\$\$SYBASE\init\logs* directory for Windows NT

The rs_init error log file name is determined by the following format:

logmonthday.session#

- *month* is a 2-digit integer that represents the current month.
- *day* is a 2-digit integer that represents the current day.
- *session#* is a 3-digit integer that represents the number of the session for that day. Each time a new rs_init session is started on the same day, the session number is incremented by one.

An example of an rs_init error log file name is log0108.001.

- *month* is 01
- *day* is 08
- *session#* is 001

Replication Server error log

The default location for a Replication Server error log is in:

- The *\$\$SYBASE/install* directory for UNIX

- The `$$SYBASE\install` directory for Windows NT

The Replication Server error log file name has this format:

```
rs_name.log
```

`rs_name` is the Replication Server name. On Windows NT, this string is limited to 8 characters.

LTM error Log

By default, the LTM error log is located in the directory in which you started the LTM. By default, the LTM error log file is named `ltm.log`.

Diagnostic tools

Diagnostic tools retrieve the status and statistics of Replication Server components which, depending on the type of problem, you can use to analyze the replication system. Check the troubleshooting section of the problem category for detailed information. This section summarizes the available diagnostic tools.

Use:

- `isql` to log in to a Replication Server or data server to see if servers are up. You can also use `isql` to execute SQL commands to see if data is the same in the primary and replicate databases, or if data has been materialized or dematerialized.
- `admin who_is_down` to find out which Replication Server threads are down.
- `admin who, sqm` to display information, such as the number of duplicate transactions or the size of stable queues, about stable queues at a Replication Server.
- `admin who, sqt` to display information, such as the number of open transactions, about stable queues at a Replication Server.
- `admin statistics, md` to display information, such as the number of messages delivered, about messages delivered by a Replication Server.

- `sp_config_rep_agent` to display the current RepAgent configuration settings.
- `sp_help_rep_agent` to display static and dynamic information about a RepAgent thread.
- `sysadmin dump_queue` to dump stable queues and view them.
- `rs_helproute` to display the status of routes at a Replication Server.
- `rs_subcmp` to compare a subscription's tables in the primary and replicate databases to make sure the tables are the same.
- `check subscription` to display the status of subscriptions at a Replication Server.
- `rs_helppub` to display publications.
- `rs_helppubsub` to display publication subscriptions.
- `sp_setrepcol` to check the replication status of text or image columns.

Troubleshooting configuration problems

Verify that a replication system is configured correctly by materializing subscription data. Some configuration problems do not appear until you attempt to materialize subscription data. The most common configuration problems that cause materialization failures are:

- Failure to log in to the primary Adaptive Server. The user who creates the subscription in the replicate Replication Server must have the same login name and password in the primary Adaptive Server and the primary Replication Server.
- Missing permissions in the primary database. The user who creates the subscription must be a user in the primary database and must have select permission in the primary table.
- Missing permissions in the replicate database. The maintenance user must have select, insert, update, and delete permissions on the tables in the replicate database.
- A Replication Server or Adaptive Server has stopped running. Try to log in to each server using `isql`. Restart any servers that are not running.

Other common configuration problems include:

- Host name resolution error.
- The Adaptive Server entry does not exist in the interfaces file.

See “rs_init error message descriptions” on page 36.

For more information about troubleshooting Replication Server configuration problems, see the configuration guide for your platform.

Troubleshooting subscription failures

Subscription failures consist of materialization and dematerialization failures. This section describes procedures you can use to troubleshoot these types of problems.

Before troubleshooting a subscription problem, always verify that data has failed to materialize or dematerialize by:

- 1 Using isql to log in to the replicate database
- 2 Executing a select command that selects the materialized or dematerialized columns from the replicate table

Also make sure that all tasks required for subscription materialization have been completed. Subscription materialization can fail if you have not completed:

- Creating replication definitions or function replication definitions
- Marking tables or stored procedures for replication
- Creating connections to the destination databases
- Creating articles, if you are using them
- Creating and validating publications, if you are using them
- Marking text or image columns for replication, if you are replicating text and image columns
- Creating direct and indirect routes, if the destination database is connected to a Replication Server different from the primary database’s Replication Server
- Creating logical connections, if you are using warm standby applications

Materialization

If a subscription has not materialized, follow the general procedure in this section. See Chapter 5, “Subscription Problems”, for specific instructions.

- 1 If you are materializing a large amount of data, make sure the *num_threads* and *num_concurrent_subs* parameters are large enough.
- 2 Log in to the destination Replication Server and check subscription status using `check subscription`. Do the same for the primary Replication Server. `check subscription` returns information that diagnoses the problem. The potential problems that `check subscription` returns information about include:
 - Other subscriptions to the same replication definition and replicate database have not yet processed
 - No connection to the primary Replication Server because of an incorrect login
 - Primary Replication Server down or out of stable queues
 - Stable Queue Manager (SQM), Stable Queue Transaction interface (SQT), and Distributor (DIST) threads down
 - Primary data server down, incorrect login, out of stable queues, or rows selected with holdlock
 - RepAgent/LTM problem
 - Route problem
 - Destination Replication Server— incorrect login or out of stable queues
 - Destination Replication Server DSI problem—use `admin who`, `dsi` or `admin who`, `sqm` to determine what the specific problem is
 - Incorrect user privileges on destination database
- 3 Use `rs_helppub` and `rs_helppubsub` to find the publications and articles that a subscription is using.
- 4 If only some columns in the table are not being materialized:
 - Check replication status of text and image columns.
 - Verify that the replication definition is correctly defined.
 - Verify that the publications and articles are correctly defined.
- 5 Fix the problem.

- 6 When you think you have solved the problem, run the replication system. If subscriptions are still not being materialized correctly:
 - Analyze errors (see “Analyzing error logs” on page 10), or
 - Complete any of this section’s procedures that you skipped.

Dematerialization

If a subscription has not dematerialized, follow the general procedure in this section. See Chapter 5, “Subscription Problems”, for specific instructions.

- 1 Log in to the destination Replication Server and check subscription status using check subscription. Do the same for the primary Replication Server. check subscription returns information that diagnoses the problem. The potential problems include:
 - Other subscriptions to the same replication definition and replicate database have not yet processed
 - No connection to the primary Replication Server because of an incorrect login
 - Primary Replication Server down or out of stable queues
 - Stable Queue Manager (SQM), Stable Queue Transaction interface (SQT), and Distributor (DIST) threads down
 - Primary data server down, incorrect login, out of stable queues, or rows selected with holdlock
 - RepAgent/LTM problem
 - Route problem
 - Destination Replication Server DSI problem
 - Incorrect user privileges on destination database
- 2 If some columns are not being dematerialized:
 - Check replication status of text and image columns.
 - Verify that the replication definition is correctly defined.
- 3 Fix the problem.
- 4 When you think you have solved the problem, run the replication system. If subscriptions are still not being dematerialized correctly:

- Analyze errors (see “Analyzing error logs” on page 10), or
- Complete any of this section’s procedures that you skipped.

Troubleshooting replication failures

Replication failures may occur after subscriptions have successfully materialized. This section documents procedures you can use to isolate replication failures.

These procedures assume that the replication system has:

- Been installed and configured correctly,
- Successfully completed subscription materialization, and
- Has correctly replicated data previously.

If you are having installation and configuration problems, see the configuration guide for your platform. If you are having subscription problems, see “Troubleshooting subscription failures” on page 14.

Before you troubleshoot a replication failure, verify that data is not replicating:

- 1 Use `isql` to log in to the primary and replicate databases.
- 2 Execute `select` commands that select the replicate columns from the replicate table and the columns to be replicated from the primary table.
- 3 Compare the data from both tables to see if it is the same.

Note If the subscription has many columns that are replicated, use `rs_subcmp` to compare data in the primary and replicate databases. See “Identify the data that is failing to replicate” on page 18.

After you have verified that data is not replicating, use the procedures and diagnostic tools in the rest of this section to investigate and identify the exact cause of your problem. These procedures can be executed separately and in any order; however, they are listed in the order in which they are most likely to solve the problem.

- “Identify the data that is failing to replicate” on page 18
- “Verify that Replication Server threads are up” on page 19
- “Replication Server is down” on page 21

- “Verify that it is not a queue problem” on page 21
- “Verify that all RepAgents and LTMs are up” on page 21
- “Whether only text and image columns are not replicated” on page 22
- “Check system tables” on page 22
- “Activate traces and dump queues” on page 23
- “Verify that no route problems exist” on page 23

When you think you have solved the problem, run the replication system. If replication is still not executing correctly:

- Analyze errors (see “Analyzing error logs” on page 10), or
- Complete any of this section’s procedures that you skipped.

Identify the data that is failing to replicate

Use this procedure to determine the specific subscription and columns that are failing to replicate. This procedure also checks to see if the primary and destination data servers and primary or destination Replication Server are running.

- 1 To find out which data in a suspect subscription is failing to replicate, use `isql` to log in to the primary or destination Replication Server and run `rs_subcmp`.

`rs_subcmp` logs in to the primary and destination data servers and compares the subscription’s data in the primary and destination tables. `rs_subcmp` can compare tables at Adaptive Servers only. To compare tables at a non-Adaptive Server data server, you can use a program equivalent to `bcp out` on the non-Adaptive Server data server and `bcp out` on the Adaptive Server data servers, and compare the output using the UNIX `diff` command.
- 2 If `rs_subcmp` successfully displays inconsistent rows, note the columns and rows that are not being replicated.
- 3 If no data exists for subscribed columns, then the subscription has not materialized. Go to “Materialization” on page 15.
- 4 If you cannot log in to a Replication Server, then that Replication Server is down. Go to “Replication Server is down” on page 21.
- 5 If `rs_subcmp` fails, then one or both of the data servers are down.

- If the primary data server is down, the Adaptive Server log may be corrupt or full. The data server may also have an operating system or hardware error. See Chapter 9, “Adaptive Server Log Problems”, for more information about troubleshooting these problems.
 - If the destination data server is down, then the data server may have a DSI problem, or an operating system or hardware error. See Chapter 8, “Data Server Interface Problems”.
- 6 Use `rs_subcmp` to check if other subscriptions on the same data server are replicating.
- If no other subscriptions are replicating, it is likely that a problem exists with that data server and not with a particular subscription.
 - If all other subscriptions are replicating, then a problem may exist with that particular subscription.
- 7 Use `rs_subcmp` to check if other subscriptions on databases controlled by the same destination Replication Server are replicating. If replication is working for other databases controlled by the destination Replication Server, then the problem is a specific database, database connection, or RepAgent/LTM. Take the following actions:
- Look for orphaned transactions in the primary Replication Server inbound queue for the database. See Chapter 9, “Adaptive Server Log Problems”.
 - Troubleshoot RepAgents and LTMs. See Chapter 7, “RepAgent and Log Transfer Manager Problems”.
 - Troubleshoot database connections. See Chapter 8, “Data Server Interface Problems”.

Once you have identified the data that is failing to replicate, go to the next procedure.

Verify that Replication Server threads are up

Use `admin who_is_down` to determine if the primary and destination Replication Server threads are up. This procedure also verifies that the Replication Servers themselves are up.

- 1 Log in to the destination Replication Server and execute `admin who_is_down`. This command displays all the threads on this Replication Server that are down. If any threads are down, an error message should have been displayed in the Replication Server error log. Check the error log again. Also see the section that corresponds to the thread that is down:
 - DSI – see Chapter 8, “Data Server Interface Problems”. The problem could be duplicate keys or permissions failure.
 - DSI EXEC – see Chapter 8, “Data Server Interface Problems”.
 - RSI User – see Chapter 6, “Replication Server Interface Problems”.
 - RS User – see Chapter 5, “Subscription Problems”.
 - SQM – the SQM should not go down. Restart the Replication Server; you cannot resume the SQM thread.
 - SQT – see “13045” on page 48.
 - User – this should have no effect on replication.
- 2 Log in to the primary Replication Server and use `admin who_is_down` to display all the threads on this Replication Server that are down. See the section that corresponds to the thread that is down as follows:
 - DIST – See “7035” on page 44 and “13045” on page 48.
 - DSI – see Chapter 8, “Data Server Interface Problems”. The problem could be a duplicate keys or permissions failure.
 - DSI EXEC – see Chapter 8, “Data Server Interface Problems”.
 - Rep AgentUser – see Chapter 7, “RepAgent and Log Transfer Manager Problems”.
 - RSI – see Chapter 6, “Replication Server Interface Problems”.
 - RSI User – see Chapter 6, “Replication Server Interface Problems”.
 - RS User – see Chapter 5, “Subscription Problems”.
 - SQM – the SQM should not go down. Restart the Replication Server; you cannot resume the SQM thread.
 - SQT – see “13045” on page 48.
 - User – this should have no effect on replication.
- 3 The *keepalive* value is set too low and the TCP connection was terminated and never restarted, if the following conditions exist:

- The DSI is down,
 - The LTM/RepAgent is not connected to the Replication Server or Adaptive Server, and
 - The entire (or part of the) network went down and was restarted.
- 4 If the DSI is up, check for data loss. Although data loss error messages show up in the Replication Server error log, these errors only show up once and may have occurred several days earlier. See Chapter 8, “Data Server Interface Problems”.
 - 5 If you cannot log in to a Replication Server, it means that Replication Server is down. See the next procedure.

Replication Server is down

If the primary, destination, or both Replication Servers are down, analyze the error messages in each Replication Server’s error log.

If both Replication Servers are down, there might be different problems for the primary and replicate Replication Servers. If both Replication Servers have their RSSDs on the same Adaptive Server data server, then troubleshoot the Adaptive Server data server.

Verify that it is not a queue problem

- 1 Run `admin who,sqm` to see if the duplicates count is rising.
- 2 Check outbound queues for stuck and open transactions.

See Chapter 8, “Data Server Interface Problems”.

Verify that all RepAgents and LTMs are up

This procedure verifies that the RepAgent or LTM for the primary Adaptive Server data server is up.

- 1 To verify that the RepAgent is active, use `isql` to log in to the primary Adaptive Server data server and execute `sp_who`.
- 2 To verify that the LTM is active, use `isql` to log in to the LTM.

- 3 If the RepAgent or LTM is down:
 - The Adaptive Server log is corrupt.
 - The Adaptive Server log is full.
 - The RepAgent or LTM set an illegal truncation point.
 - The RepAgent or LTM used an incorrect primary Replication Server login.
 - text/image columns have inconsistent replication status.
 - Incorrect replication system configuration.

See Chapter 7, “RepAgent and Log Transfer Manager Problems”, for more information about solving RepAgent or LTM problems. See Chapter 9, “Adaptive Server Log Problems”, for information about log problems.

Whether only *text* and *image* columns are not replicated

If you run `rs_subcmp` and see that only text and image columns at the destination table are not being replicated, then these columns may have inconsistent replication status. See “Replication status inconsistency of text or image columns” on page 112 for more information about repairing this problem.

Check system tables

This procedure checks to see if there is any error information in the RSSD system tables.

- 1 See if there are any recoverable actions in the `rs_recovery` system table.
- 2 See if there are any “detecting losses” or “rejecting messages after loss detected” validation status messages in the `rs_oqid` system table. (You might have missed the message in the error log.)

Activate traces and dump queues

You specify trace flags when you run the diagnostic versions of the Replication Server (*repsrvr.diag*) and LTM (*ltm.diag*). By turning on traces and examining the output of `sysadmin dump_queue`, you can determine the component in which replication has stopped.

If replication stopped in one of the following components, see the corresponding chapter or section:

- Primary data server – Chapter 8, “Data Server Interface Problems”, and Chapter 9, “Adaptive Server Log Problems”.
- LTM/RepAgent – Chapter 7, “RepAgent and Log Transfer Manager Problems”.
- Outbound queue – Chapter 4, “Route Problems”, and Chapter 8, “Data Server Interface Problems”.
- DSI thread – Chapter 8, “Data Server Interface Problems”.
- Custom function-string class or customized function strings – Chapter 8, “Data Server Interface Problems”, or “Using trace” on page 138.

Verify that no route problems exist

This procedure checks for route problems. See Chapter 4, “Route Problems”, for more information about using `rs_helproure` and solving route problems.

- 1 Make sure that no other databases are replicating through the same route as the non-replicating subscription. If other databases are replicating through that route, then problems within the primary Replication Server/database and replicate Replication Server/database are more likely. It is unusual for replication to have stopped for one database while replication continues through the same route for other databases. Run `admin statistics, md` on the replicate Replication Server to see if “Messages Delivered” increases over time.
- 2 If no other databases are replicating through the same route as the non-replicating subscription, run `rs_helproure` at the primary or destination Replication Server in the RSSD to determine if there is a route problem.

Troubleshooting manual recovery problems

If you recently recovered a primary database and replication is failing, one of these problems may have occurred:

- A new generation ID for the primary database has not been set. See Chapter 19, “Replication System Recovery,” in the *Replication Server Administration Guide*.
- If you loaded the primary and replicate databases from the same backup, make sure that the `rs_lastcommit` table for the replicate database has the correct entries. If the `origin_qid` in the `rs_lastcommit` table is incorrect (old or changed), the DSI may ignore certain transactions associated with that `origin_qid`. You should bcp out the `rs_lastcommit` table, load the data, and then bcp in the `rs_lastcommit` table. If DSI still ignores certain transactions, update the `origin_qid` value to “0” to force DSI to accept these transactions.
- If you used dump and load to synchronize the primary and replicate databases because of a failure, be sure to increase the generation number in the primary database. If you do not set the generation number correctly, replication from this database may stop because Replication Server perceives that it has already processed the messages.

Also see Chapter 19, “Replication System Recovery,” in the *Replication Server Administration Guide*, which contains information for recovering primary databases.

Analyzing Error Messages

This chapter describes the format of Replication Server and RepAgent error messages, how to analyze these error messages, and provides example analyses of these error messages.

Name	Page
Overview	25
Replication Server errors	26
RepAgent error messages	31

Overview

Error messages provide important information for identifying the cause of replication system problems. When a Replication Server or RepAgent error occurs, an error message is recorded in an error log. Error messages are displayed in a standard format; once you learn this format, you will be able to read and understand all Replication Server and RepAgent error messages.

Replication Server error messages are recorded in text files called error logs or sent to the standard error output (*stderr*), which is usually a terminal window. In general, replication agent error messages are recorded in their own error logs, except for RepAgent error messages, which are recorded in the Adaptive Server data server error log. LTM error messages are recorded in the LTM's error log. Data server error messages are recorded in the data server's error log.

Multiple error messages can be generated by a single problem, and can be generated from any or all Replication Server components (including Adaptive Server, Replication Server, and RepAgent) that are adversely affected by the problem.

Replication Server, RepAgent, and LTM can print messages to the error log and to clients in several languages. See Appendix B, "Language, Sort Order, and Character Set Issues" for problems that can result from output in mixed languages.

You can monitor error logs using Replication Server Manager (RSM). This guide, however, does not show how to use RSM to monitor error logs or read the RSM error log. See the online help for Sybase Central and the *Replication Server Administration Guide* for information about how to perform these tasks.

What is not covered

This guide does not explain how to analyze the error logs from networks, gateways, non-Adaptive Server data servers, and non-Adaptive Server replication agents. See the documentation for these products for information about how to analyze their error logs.

This guide does not explain how to analyze the error logs from the LTM error log. For more information about the LTM error log, see the *Replication Server Administration Guide*.

Replication Server errors

Replication Server errors are recorded in the Replication Server error log. A Replication Server error log contains informational, warning, thread-terminated, fatal, and internal error messages. Messages are appended to the error log while Replication Server is operating; each time you restart Replication Server, error messages are appended to the end of the existing error log.

Informational messages report the current status of a component; for example, a process or command has completed or a component has been shut down. Warning, thread-terminated, fatal, and internal error messages are reported when a component abnormally terminates, a process or command cannot be completed, or a fatal internal error occurs in a component.

When Adaptive Server errors cause a Replication Server error, you find references to the Adaptive Server errors in the text of the Replication Server error message. In these cases, you must fix the Adaptive Server problem, which should eliminate the Replication Server error.

Replication Server error logs contain error messages generated during the execution of asynchronous commands, such as create subscription and create route. When you are executing asynchronous commands, pay special attention to the error logs for the Replication Servers affected by the asynchronous command.

Dedicate a terminal window to display the standard error output from a Replication Server. If the Replication Server error log is unavailable, error messages are sent to *stderr*. Error messages that are sent only to *stderr* are extremely serious, because the only action available to the Replication Server is to dump the stack and exit.

Error message format

The format of Replication Server error messages is:

```
s. date time severity_text error_num thread  
(thread_context) module (line_num) error_text
```

The variables in the error message format are defined as:

- *s* – indicates the severity or seriousness of the error.
- *date* – date.
- *time* – military time.
- *severity_text* – brief phrase that describes the severity.
- *error_num* – integer that uniquely identifies the error.
- *thread* – name of the Replication Server thread that received the error.
- *thread_context* – more specific information about the thread, such as the database and data server if it is a DSI thread.
- *module* – name of the Replication Server source file where the error was reported. This name is only for use by a Sybase Technical Support Engineer.
- *line_num* – line in the Replication Server source file where the error was reported. This number is only for use by a Sybase Technical Support Engineer.
- *error_text* – error message text. If a Replication Server error was caused by an Adaptive Server error, the Adaptive Server error is displayed as part of the Replication Server error message text.

Error messages with a severity of “I” (informational) contain only the date and time the error occurred, and the full text of the message.

Error messages are displayed in the language specified in the *RS_language* configuration parameter.

Severity

Each error message begins with a letter that indicates the severity of the error. Table 2-1 lists the severities.

Table 2-1: Replication Server error severities

Severity Code	Description
I	An informational message.
W	A warning about a condition that has not yet caused an error, but may require attention, for example, running out of a resource.
E	An error that does not prevent further processing, such as a site that is unavailable.
H	A Replication Server thread has died, for example, a lost network connection.
F	Fatal. A serious error caused Replication Server to exit. For example, if you start the Replication Server using an incorrect configuration.
N	Internal error, caused by anomalies in the Replication Server software. Report these errors to Sybase Technical Support.

The format of error messages in the LTM and Replication Server are identical. For more information about error severity, see the *Replication Server Administration Guide*.

Error numbers

Error numbers from Replication Server are constructed using the following formula:

$$\text{module_number} * 1000 + \text{error_condition}$$

A *module_number* is assigned to each source code module in the server and the *error_condition* is a numbered error condition within the source code module.

Analyzing Replication Server error messages

To analyze a Replication Server error log, you must understand how multiple error messages are related. Multiple error messages can be generated by the same problem and usually have identical or very close dates and times. Error messages that occur close together constitute an *error block*.

Replication Server error messages may contain references to Adaptive Server errors. By examining the causes of the Adaptive Server error, you will be able to solve the problem.

The Replication Server error log may also contain Open Client/Server error messages, since Replication Server is an Open Server program with Open Client components. Any Open Server errors in the error log constitute internal errors.

To read a Replication Server error log:

- 1 Look for the most recent Replication Server error by scanning from the end to the beginning of the error log. As Replication Server errors occur, Replication Server appends the error messages to the end of the error log. The most recent Replication Server error marks the end of the error block.

- 2 Find the first Replication Server error message in the error block.

Note If you see an Open Server/Client error such as “no more threads,” fix that error first. An Open Server error “no more threads” message means that you have run out of an Open Server resource (and Replication Server might have terminated).

- 3 Find and diagnose the last non–Replication Server error message that occurred after the first Replication Server error. If there are no non–Replication Server error messages after the first Replication Server error message, investigate the causes of the first Replication Server error message.

Note In diagnosing non–Replication Server errors, you may need to use other troubleshooting or error message guides—such as those for Adaptive Server or Open Client (particularly with network communication or other client problems).

- 4 In the error block, find any Replication Server error messages that tell you if a Replication Server component or thread has terminated. You must restart the component or thread that has terminated. Usually, subsequent Replication Server errors in the same error block were caused by the first Replication Server error.

Example

Examine the following error message block in the Replication Server log:

```
H. 97/08/23 09:52:02. THREAD FATAL ERROR #5049
DSI(WESTERNSD.westDB) - dsigmint.c(3675) The DSI thread
for database 'westernDS.westDB' is being shutdown. DSI
received data server error #102 which is mapped to
STOP_REPLICATION. See logged data server errors for more
information. The data server error was caused by RS
output command #1 mapped from input command #1 of the
failed transaction.
I. 97/08/23 09:52:02. The DSI thread for database
'WESTERNSD.westDB' is shutdown.
```

Table 2-2 breaks this error message into its components.

Table 2-2: Replication Server error message example description

Error Message Text	Explanation
H. 97/08/23 09:52:02. THREAD FATAL ERROR #5049 DSI (WESTERNDS.westDB) - dsiqmint.c(3675)	<p>This the first error. The H identifies the error as a thread termination error. This message shows that the DSI thread terminated.</p> <p>The error message includes the dataserver.database name where the error occurred, the internal Replication Server module (<i>dsiqmint.c</i>), and the line of code (3675) where the error occurred.</p>
The DSI received data server error # 102 which is mapped to STOP_REPLICATION. See the logged data server errors for more information.	<p>This Adaptive Server error message identifies the cause of the problem.</p> <p>Find a description for Adaptive Server Error 102 by:</p> <ul style="list-style-type: none"> • Finding the error in the Adaptive Server error log (the error would occur at approximately the same time as it occurred in the Replication Server error log), • Performing a select * from sysmessages where error = 102 in the master database, or • Looking up the number in the <i>Adaptive Server Error Messages or Troubleshooting and Error Messages Guide</i>. <p>The Adaptive Server Error Messages or Troubleshooting and Error Messages Guide describes Error 102 as invalid syntax near %s for a Transact-SQL statement. It also gives suggestions about when this error occurs (a keyword is spelled incorrectly, a keyword or parameter is missing, or the order of the keyword is incorrect). Look for additional information in the data server error log, which displays the character string for the %s field in the error message.</p>
I. 97/08/23 09:52:02. The DSI thread for database 'WESTERNDS.westDB' is shutdown.	<p>This last error message is informational (I) and caused by the problem specified in a previous block. After fixing the Adaptive Server problem, you must restart the DSI thread for the specified database.</p>
The data server error was caused by RS output command #1 mapped from input command #1 of the failed transaction.	<p>This text describes the command position in a grouped transaction sent by the Replication Server.</p>

RepAgent error messages

RepAgent records RepAgent errors such as “Could not spawn Rep Agent Thread” and all error messages returned by Replication Server (including CT-Library error messages) in the Adaptive Server error log.

RepAgent error message format

RepAgent error messages are recorded in the Adaptive Server error message format and are identified by the string “RepAgent(*dbid*),” which appears in the first line of the message. *dbid* is the database identification number of the RepAgent that logged the error.

Errors that you can retry are logged only once in the Adaptive Server error log.

The RepAgent error message format is:

```
date time RepAgent (dbid): error_number, severity,  
state, error_text
```

The RepAgent error format includes the following information:

- *date* – date that the error occurred.
- *time* – time that the error occurred.
- *dbid* – Adaptive Server identification number of the database that RepAgent is using. You can find this database ID by executing this command:

```
select x = db_id()
```

- *error_number* – RepAgent error message numbers range from 9200 to 9299.
- *severity* – severity can be:
 - EX_INFO: informational error message
 - EX_USER: user error
 - EX_RESOURCE: resource error in which an operating system resource or Replication Server resource is not available
 - EX_CMDFATAL: fatal error in which RepAgent cannot continue processing a transaction
- *state* – is for internal use only.
- *error_text* – description of the cause of the error.

Example

Here is an example of a RepAgent error message:

```
01/08/1998 14:47:58 PM RepAgent (10) 9209, EX_CMDFATAL,  
1: Missing datarow in TEXT/IMAGE insert log record.  
Transaction log may be corrupt. Please contact SYBASE
```

Technical Support. (current marker = (%d, %d)).

Table 2-3 breaks this error message into its components.

Table 2-3: RepAgent error message example description

Error Message Text	Explanation
01/08/1998	Date
14:47:58 PM	Time
(10)	Database ID
9209	Error number
EX_CMDFATAL	Severity
1	State
Missing data row in TEXT/IMAGE insert log record. Transaction log may be corrupt. Please contact SYBASE Technical Support. (current marker = (%d, %d)).	Error message text

For more information about the Adaptive Server error message format, see the *Adaptive Server Enterprise System Administration Guide*.

Analyzing RepAgent error messages

Analyzing RepAgent errors is much easier than analyzing Replication Server errors.

Example

The following error message indicates that the RepAgent thread attempted to connect to a database that already has a RepAgent thread running.

```
01/12/1998 16:37:58.34 PM RepAgent (10) 9220, EX_User,
1: The Rep Agent Thread context for database westernDB
has already been reserved by another session.
```

Only one RepAgent thread can connect to a database at any one time. Because the RepAgent abnormally terminated, the RepAgent thread process still exists although it is in sleep mode. The RepAgent thread wakes after the time specified in the *scan_retry* parameter elapses. Since the RepAgent has already terminated, the RepAgent thread process terminates when it wakes.

Common Errors

This chapter describes some of the most common Replication Server errors and solutions.

Name	Page
Finding the error description	35
rs_init error message descriptions	36
Replication Server error message descriptions	37
Replication Server informational and warning messages	53
RepAgent error message descriptions	57
LTM-only error message descriptions	61

Note This chapter does not describe all errors. If you do not find the error you are looking for in this chapter, refer to Chapter 1, “Identifying the Problem”, for more troubleshooting information or refer to the chapter that corresponds to your error description.

Finding the error description

To find the error description in this chapter, go to the appropriate section and look for the error by:

- The error number of a Replication Server or LTM error, or
- The text of the rs_init error, or Replication Server informational or warning message.

If there are multiple errors in an error log, use the procedure in Chapter 2, “Analyzing Error Messages”, to determine which error to look for in this chapter.

Error description format

Each error description includes these sections:

- Symptoms – includes the text of any error messages that are displayed in an error log and error conditions, such as poor performance, replication failure, connection failure, and abnormal component termination.
- Explanation – describes the error messages and probable causes of the errors.
- Solutions – describes procedures, workarounds, upgrades, and EBF information.

rs_init error message descriptions

Because rs_init error messages do not have error numbers, they are arranged alphabetically.

Error Could not find entry for Adaptive Server

Message text Error messages stating that the Adaptive Server entry does not exist in the *interfaces* file are displayed in the rs_init error log.

```
08/23/96 02:06:49 PM Running task to check the Adaptive Server.  
08/23/96 02:06:49 PM Could not find entry for Adaptive Server 'westdss' in  
the interfaces file.  
08/23/96 02:06:49 PM Task to check the Adaptive Server failed; terminating  
configuration.  
08/23/96 02:06:49 PM Configuration failed.  
08/23/96 02:06:49 PM Exiting.
```

Explanation When you execute rs_init with a resource file (rs_init -r) to configure an LTM, rs_init looks in the *interfaces* file for the Adaptive Server name that matches the one in the resource file.

The errors listed above occur because the Adaptive Server entry in the *interfaces* file does not match the resource file's Adaptive Server entry due to unprintable characters in the resource file.

Action Check for unprintable characters (such as control characters) in the Adaptive Server name entry in the resource file. The Adaptive Server name entry is specified as the *rs.rs_ds_name* parameter.

Error Unknown host machine name.

Message text	The following messages are displayed in the <code>rs_init</code> error log. <pre>08/09/98 11:54:43 AM Running task to check the RSSD Adaptive Server. 08/09/98 11:54:45 AM CONNECTIVITY ERROR: DB-Library error: 'Unknown host machine name.'. 08/09/98 11:54:45 AM Unable to connect to the RSSD SQL Server 'axptst'. Make sure that the Adaptive Server is running, and the 'sa' password is correct. 08/09/98 11:54:45 AM Task to check the RSSD Adaptive Server failed; terminating configuration. 08/09/98 11:54:45 AM Configuration failed. 08/09/98 11:54:45 AM Exiting.</pre>
Explanation	<code>rs_init</code> cannot find the RSSD machine's host name. Sometimes, directory services do not recognize the machine's host name and, instead, may recognize a machine only by its Internet Protocol (IP) address.
Action	Add an additional query line for the machine to the <i>interfaces</i> file and specify the host machine's IP address instead of its name. For example, for the <code>bss.bsslt.ch.corma.com</code> host name substitute its IP address, <code>2.41.100.35</code> , as follows: <pre>query tcp ether bss.bsslt.ch.corma.com 31440 query tcp ether 2.41.100.35 31440.</pre> <p>If you run <code>rs_init</code> again, the following messages may be displayed in the <code>rs_init</code> error log, but you can ignore them:</p> <pre>Could not obtain Hostname for Internet address '2.41.100.35'. Could not obtain Hostname for Internet address '2.41.100.41'. Could not obtain Hostname for Internet address '2.41.100.35'. Could not obtain Hostname for Internet address '2.41.100.41'. Running task to check the RSSD Adaptive Server.</pre>

Replication Server error message descriptions

These error messages, which are recorded in the Replication Server error log, are arranged numerically in ascending order.

1028	Data server errors
Explanation	Error 1028 is displayed in the Replication Server error log. An Adaptive Server error is reported in the text of error 1028. <pre>E. 95/08/09 11:38:01. ERROR #1028 DSI(westernDS.westDB) - dsigmint.c(3675) Message from server: Message: ..., State: ..., Severity: ...-- '...'</pre>

Error 1028 occurs when a data server returns an error to Replication Server. Error 1028 is usually followed by several other errors. The most common error situations are described in these sections:

- DSI shuts down because of SQL error in transaction
- Adaptive Server and DB2 table names do not match
- Adaptive Server last-chance threshold passed
- DSI shuts down because of SQL error in transaction
- Accessing a database while in recovery

In some cases, the DSI thread shuts down. Adaptive Server errors caused the connection to be suspended if:

- The Adaptive Server error is assigned the `retry_stop` or `stop_replication` error action.
- You executed the `suspendconnection` command with the `nowait` clause.

If the DSI thread shuts down, you cannot simply resume the connection. Resuming the connection without fixing the problem sends the same transaction to the Adaptive Server and causes the same error.

Solutions

In general, you must identify and fix the Adaptive Server problem that is causing the error. See “Correcting Adaptive Server error” on page 38. If you cannot correct the cause of the problem, as a last resort, you can resume the connection and skip the transaction. Skipping a transaction, however, causes inconsistencies between the primary and replicate databases, which you must manually fix in the replicate database. See “Skipping the transaction” on page 40. Also see the most common situations for error 1028 for solutions to their particular problems.

Correcting Adaptive Server error

Attempt to fix the Adaptive Server error by analyzing the transaction text in the exceptions log as follows:

- 1 Log in to the RSSD and display the transaction ID by executing the `rs_helpexception` stored procedure.

```
1> rs_helpexception
2> go
```

```
Summary of Logged Transactions on 'westernRS'
Total # of Logged Transactions = 1
Xact ID Org Site Org User Org Date Dest Site # Recs/Xact
```



```
-----
107 mil01hprds. eur eurian Jan 26 1996 westernDS. eur 3
For Detailed Information on a Logged Xact., type 'rs_helpexception {Xact
ID}'
(return status = 0)
```

- 2 Using the transaction ID displayed in the previous step, display the full transaction text by executing the `rs_helpexception` stored procedure with the `v` option. The `v` option includes the text of the transaction in the output.

```
1> rs_helpexception 107, v
2> go
```

You see this information:

```
Detailed Summary of Logged Transaction # 107 on 'westernRS' Origin Site
Origin User Org. Commit Date #Cmds in Xact
```

```
-----
westernDS.westDB eurian Jan 26 1996 9:27 3
Dest. Site Dest. User Date Logged
```

```
-----
westernDS.westDB eurian Jan 29 1996 10:39
This transaction was logged by the 'sysadmin log_first_tran' command.
Rejected Records
textval
```

```
-----
A0100distribute :origin_time='Jan 26 1996 9:27:24:416AM',:origin_user=
',:mode=1
begin transaction 'logexec' for 'eurian'/'*****'
begin transaction
A0100distribute :origin_user='',:mode=1 exec "TT"."so_req_rep_all_allco
n" @"p01"=80000709,@"p02"='MIL'
execute tt_act_rep_all_allcon @p01 = 80000709, @p02 = 'MIL'
A0100distribute :origin_time='Jan 26 1996 9:27:24:416AM',:origin_user=
',:mode=1 commit transaction
execute rs_update_lastcommit @origin = 107, @origin_gid = 0x
0000001004620d300019296000effffffff000000008910009bd7cd00010000000000
1, @ secondary_gid =
0x0000000000000000000000000000000000000000000000000000000000000000
0000000000, @origin_time = 'Jan 26 1996 9:27:24:416AM' commit transaction
(return status = 0)
```

This text corresponds to what is sent to the database (in this case, function strings have been applied).

- 3 Use this transaction information to manually apply the update to the replicate database.

- 4 After fixing the error in the database, resume the connection in the Replication Server.

```
> resume connection to westernDS.westDB
```
- 5 Delete the transaction from the exceptions log to keep the log small:

```
1> exec rs_delexception 107
2> go
```

Skipping the transaction

If you cannot fix the Adaptive Server error, skip the transaction and manually apply the skipped transaction to the replicate table to resynchronize the primary and replicate tables.

See Chapter 8, “Data Server Interface Problems”, for information about how to apply skipped transactions to a replicate database.

DSI shuts down because of SQL error in transaction

Symptoms

The following error messages are displayed in the Replication Server error log. The Data Server Interface (DSI) thread shuts down.

An Adaptive Server error is included in the text of the Replication Server error.

```
E. 95/08/09 11:38:01. ERROR #1028 DSI(western.west1) -
dsiqmint.c(3645) Message from server: Message: 229,
State: 1, Severity: 14-- 'DELETE permission denied on
object real_run, database west1, owner dbo'.
```

```
H. 95/08/09 11:38:01. THREAD FATAL ERROR #5049
DSI(western.west1) - dsiqmint.c(3652) The DSI thread
for database 'western.west1' is being shutdown. DSI
received data server error #229 which is mapped to
STOP_REPLICATION. See logged data server errors for more
information. The data server error was caused by RS
output command #1 mapped from input command #2 of the
failed transaction.
```

```
I. 95/08/09 11:38:01. The DSI thread for database
'western.west1' is shutdown. The Message from server:
text is the message you would get from the Adaptive
Server had you sent the command batch to the server with
isql. The text is taken from the table sysmessages. In
isql you would get: Msg 229, Level 14, State 1: DELETE
permission denied on object real_run, database west1,
owner dbo
```

Explanation	<p>These errors occur when a Replication Server transaction that is sent to the replicate Adaptive Server has a SQL error. The replicate Adaptive Server detects the SQL error in the transaction and returns a message to the Replication Server.</p> <p>The Replication Server responds by shutting down the connection and suspending replication; this prevents data inconsistencies between the primary and replicate sites. It allows the user to correct the problem at the replicate Adaptive Server and to maintain data consistency.</p> <p>For example, when a SQL semantic error occurs in a referential integrity or duplicate keys situation, Adaptive Server sends its message back to Replication Server. In response, Replication Server shuts down the connection and records the Replication Server error in the Replication Server error log. The Adaptive Server error is referenced in the body of the Replication Server error.</p>
Solution	<p>Fix the problem in Adaptive Server and resume the connection as described in the Solutions section of 1028: data server errors. If the problem persists, verify that autocorrection for the subscription has been turned on in the <code>rs_reprobs.attributes</code> table of the RSSD. Also, check for triggers that enforce referential integrity rules.</p>

Adaptive Server and DB2 table names do not match

Symptoms	<p>These messages are displayed in the Replication Server error log.</p> <pre>E. 95/10/10 12:10:07. ERROR #1028 DSI(DSMA1.apptst) - dsiqmint.c(3668) Message from server: Message: 208, State: 1, Severity: 16 -- 'TEST not found. Specify owner.objectname or use sp_help to check whether the object exists (sp_help may produce lots of output). ` H. 95/10/10 12:10:08. THREAD FATAL ERROR #5049 DSI(DSMA1.apptst) - dsiqmint.c(3675) The DSI thread for database 'DSMA1.apptst' is being shutdown. DSI received data server error #208 which is mapped to STOP_REPLICATION. See logged data server errors for more information. The data server error was caused by RS output command #1 mapped from input command #2 of the failed transaction. E. 95/10/10 11:58:09. ERROR #32032 LTM USER(TCPIP.ZD60) - /nrm/nrm.c(1658) No table with name 'TEST' is defined at repserver with id TCPIP.ZD6</pre>
Explanation	<p>These errors occur when replicating a DB2 table with an uppercase name to an Adaptive Server. DB2 object names are in uppercase while Adaptive Server object names can be in mixed case.</p>
Solutions	<p>Use any one of the following solutions:</p> <ul style="list-style-type: none"> • Replicate into VIEWS with uppercase object names.

- Create a table name in uppercase, and use custom function strings with column names that match the case of the columns in the replicate table.
- Create the Adaptive Server object names in uppercase to match the DB2 object names.
- If the Adaptive Server table uses lowercase names, use function strings to force the Replication Server to generate lowercase table and column names when transactions are sent to the Adaptive Server.

Adaptive Server last-chance threshold passed

Symptoms The following errors, which include references to the Adaptive Server error 7415, are recorded in the Replication Server error log.

```
I. 96/08/08 10:45:07. Message from server: Message: 7415, State: 1,
Severity: 10 - - 'The transaction log in database northDB is almost full.
Your transaction is being suspended until space is made available in the
log.'
```

```
E. 96/08/08 10:45:07. ERROR #5046 DSI(axp st.northDB) - /dsioqid.c(1638)
When executing the rs_get_lastcommit function in database 'axp st.northDB',
received data server errors. See logged data server errors for more
information.
```

Explanation By default, the replicate Adaptive Server suspends all transaction processing when the destination database log segment size becomes greater than the Adaptive Server's last-chance threshold. The last-chance threshold is a parameter that keeps the log from expanding to greater than the maximum size of the log.

If the database is the RSSD, other serious consequences can occur. See “RSSD log device full” on page 47.

This problem is related to “DSI shuts down because replicate database log is full” on page 43, which occurs when the Adaptive Server log for the database is almost full.

Note This problem can also apply to Replication Server version 11.0.x.

Solutions Reduce the size of the log by dumping or truncating it as follows.

1 Manually dump the log using these commands in the RSSD:

```
> sp_helpdb northDB
> go
> dump tran northDB to ...
```

```
> go
> sp_helpdb northDB
> go
```

- 2 If this step fails, truncate the log by executing `dump tran` with the `truncate_only` or `no_log` option.

DSI shuts down because replicate database log is full

Symptoms The DSI thread shuts down and the following error message is displayed in the Replication Server error log.

```
E. 96/07/17 11:38:16. ERROR #1028 DSI EXEC(107(1) SYDNEY_DS.pubs2) -
dsiqmint.c( 2361) Message from server: Message: 1105, State 3, Severity 17
-- 'Can't allocate space for object 'syslogs' in database 'pubs2' because
the 'logsegment' segment is full. If you ran out of space in syslogs, dump
the transaction log. Otherwise, use ALTER DATABASE or sp_extendsegment to
increase the size of the segment.
```

```
H. 96/07/17 11:38:16. THREAD FATAL ERROR #5049 DSI EXEC(107(1)
SYDNEY_DS.pubs2) - dsiqmint.c(2368) The DSI thread for database
'SYDNEY_DS.pubs2' is being shutdown.
```

```
DSI received data server error #1105 which is mapped to STOP_REPLICATION.
See logged data server errors for more information. The data server error
was caused by output command #1 mapped from input command #1 of the failed
transaction.
```

Explanation The DSI thread to the replicate data server has shut down because the Adaptive Server log for the database is full. A full Adaptive Server log could be caused by:

- Problems at the replicate database, or
- Problems at the replicate Replication Server.

This problem is related to “Adaptive Server last-chance threshold passed” on page 42, which occurs when the Adaptive Server log for the database is almost full.

Solutions Complete the following tasks:

- 1 Add space to the Adaptive Server log, after verifying that there is space in the Replication Server stable device.
- 2 If there is not enough space in the stable device to increase the size of the log, add space to the stable device.
- 3 Make sure the DSI thread to the data server and the RSI thread to the Replication Server are up and running.

- Starts up, or
- Resumes a connection, or
- Replicates new transactions after a closed connection.

Accessing a database while in recovery

Symptoms If a Replication Server attempts to use an Adaptive Server database that is in the process of recovery, you see these messages in the Replication Server error log:

```
E. 96/01/19 18:36:45. ERROR #1028 DSI(western.west1) - /dsiexec.c(306)
Message from server: Message: 921, State: 1, Severity: 14 -- 'Database
'west1' has not been recovered yet - please wait and try again.'.
I. 96/01/19 18:36:46. Message from server: Message: 5701, State: 1,
Severity: 10 - - 'Changed database context to 'master'.'.
E. 96/01/19 18:36:46. ERROR #5051 DSI(western.west1) - /dsiexec.c(314)
Received errors from database 'western.west1'. See logged ct-lib and data
server messages for more information.
```

Explanation A Replication Server cannot connect to an Adaptive Server database that is recovering. The database is not accessible until it has recovered. Replication Server automatically attempts to connect to databases (to which connections have been created) when Replication Server:

- Starts up, or
- Resumes a connection, or
- Replicates new transactions after a closed connection.

Solution The connection resumes automatically after the database has recovered.

7035

Replication Server out of memory

Symptoms After a connection is started, these error messages are displayed in the Replication Server error log:

```
I. 95/12/19 16:42:09. The DSI thread for database 'westernDS.westDB' is
started.
E. 95/12/19 16:42:17. ERROR #7035 DIST(westernDS.westDB) - m/memseg.c(771)
.....Additional allocation would exceed the memory_limit of '20000'
specified in the configuration.
Increase the value of the memory_limit parameter in the rs_config table and
```

restart the Repserver.

A second attempt at resuming the connection displays only the following informational message:

```
I. 95/12/19 16:50:21. Attempt to start a DSI thread for
database 'westernDS.westDB' that has already been
started.
```

Running `adminwho_is_down` at the Replication Server indicates that threads are down.

Both DIST and DIST EXEC threads are down:

```
Spid Name State Info
-----
DIST      Down   westernDS.westDB
DIST EXECDown 105:1 westernDS.westDB
```

Explanation Replication Server ran out of segment memory. The maximum amount of operating system memory that a Replication Server can use is equal to the value specified by the `memory_limit` parameter. Replication Server directly uses operating system memory.

Solutions To fix this problem:

- 1 Increase the value of the `memory_limit` parameter.
- 2 Restart the Replication Server.

11061

RSSD problems.

Explanation Error 11061 is displayed in the Replication Server error log when an Adaptive Server error occurs. The most common error situations are:

- RSSD Deadlocks
- RSSD Log Device Full

RSSD deadlocks

Symptoms The RSSD deadlocks and these error messages are displayed in the Replication Server error log:

```
E. 96/08/12 12:00:25. ERROR #11061 USER(rho_dbo) - s/stscol.c(1717) Check
the log for error messages from RSSD.
```

E. 96/08/12 12:00:25. ERROR #1028 USER(rho_dbo) - s/stscol.c(1717) Message from server: Message: 1205, State: 2, Severity: 13 -- 'Your server command (process id #14) was deadlocked with another process and has been chosen as deadlock victim. Re-run your command.'

Explanation RSSD deadlocks usually occur when commands for the RSSD are issued faster than the server can process them. For example, RSSD deadlocks may occur even on a fast machine and network when you run scripts that create, alter, or delete many subscriptions or replication objects.

Note In a production environment, deadlock situations on the replicate database are automatically handled by the Replication Server.

RSSD deadlocks are known to occur when you are:

- Creating routes in parallel within a star configuration. A *star configuration* has one primary Replication Server with only direct routes to other destination Replication Servers and each destination Replication Server has only one direct route back to the primary Replication Server.
- Creating, activating, or validating subscriptions in one or more Replication Servers.
- Dropping replication definitions in parallel in different Replication Servers.

Solutions If routes are deadlocked, drop the routes and re-create them sequentially, allowing one minute between each creation.

If an RSSD deadlock occurs during the activation or validation of subscriptions:

- 1 Use `rs_helpsub` in the RSSD or check subscription in the Replication Server to check for subscriptions with status “active/activating” instead of “active/unknown.”
- 2 Drop the “active/activating” subscriptions using the `without purge` option and re-create subscriptions.

If deadlocks occur while you are dropping subscriptions, drop them again.

To prevent a large number of deadlocks, do not load several scripts into Replication Server simultaneously. In extreme situations, avoid loading scripts simultaneously in different Replication Servers; instead run scripts sequentially.

RSSD log device full

Symptoms	<p>8.2 RSSD Log Device Full.</p> <p>These error messages are displayed in the Replication Server error log:</p> <pre> E. 96/08/13 10:35:15. ERROR #11061 USER(western_dbo) - s/stscol.c(1717) Check the log for error messages from RSSD. I. 96/08/13 10:35:15. Message from server: Message: 7412, State: 1, Severity: 10 - - `Space available in the log segment has fallen critically low in database `rssd`. All future modifications to this database will be suspended until the log is successfully dumped and space becomes available.'. I. 96/08/13 10:35:15. Message from server: Message: 7415, State: 1, Severity: 10 - - `The transaction log in database rssd is almost full. Your transaction is being suspended until space is made available in the log.`. </pre>
Explanation	<p>During subscription validation, the RSSD ran out of log space even though the truncate on checkpoint option had been set for the RSSD. Replication Server halted and Adaptive Server suspended modifying all transactions. After you resolved the log space issue, Replication Server continued to process the subscription but did not validate it correctly.</p>
Solutions	<p>To solve this problem, reduce the size of the log by dumping or truncating it as follows.</p> <ol style="list-style-type: none"> 1 Manually dump the log using these commands in the RSSD: <pre> > sp_helpdb RSSD > go > dump tran RSSD to ... > go > sp_helpdb RSSD > go </pre> 2 If this step fails, truncate the log by executing dump tran with the truncate_only or no_log option. <p>To prevent this error, monitor the RSSD log. If the RSSD log becomes more than 80% full:</p> <ol style="list-style-type: none"> 1 Suspend the operating system process that creates replicated objects (subscriptions). 2 Wait one minute to allow Replication Server to finish its transactions.

- 3 Dump the transaction log.
- 4 Resume the process.

13045

Replication suspended because RSSD restarted

4.6 RSSD SQL Server Bounced

Symptoms

These error messages are displayed in the Replication Server error log:

```
E. 96/09/30 14:59:01. ERROR #13045 LTM USER(longdss.amercyp) -
seful/cm.c(2597)
    FAILED to connect to server 'westss' as user 'westrs_rssd_prim'. See
ct-lib and/or server error messages for more information.
I. 96/09/30 14:59:01. Trying to connect to server 'westss' as user
'westrs_rssd_prim' .....
```

After the Adaptive Server with the RSSD has been started again, the following error messages are displayed in the Replication Server error log:

```
E. 96/09/30 14:59:01. ERROR #1027 DIST(westss.amerttp) - seful/cm.c(2593)
    Open Client Client-Library error: Error: 84083972, Severity: 5 --
'ct_connect(): unable to get layer message string: unable to get origin
message string: Net-Lib protocol driver call to connect two endpoints
failed', Operating System error 0 -- 'Socket connect failed - errno 22'.
E. 96/09/30 14:59:01. ERROR #13045 DIST(westss.amerttp) - seful/cm.c(2597)
    FAILED to connect to server 'westss' as user 'westrs_rssd_prim'. See
ct-lib and/or server error messages for more information.
I. 96/09/30 14:59:01. Trying to connect to server 'westss' as user
'westrs_rssd_prim' .....
```

```
E. 96/09/30 14:59:08. ERROR #1027 DSI(westss.amerttp) - /dsiutil.c(278)
    Open Client Client-Library error: Error: 84083974, Severity: 5 --
'ct_results(): unable to get layer message string: unable to get origin
message string: Net-Library operation terminated due to disconnect'.
E. 96/09/30 14:59:08. ERROR #5097 DSI(westss.amerttp) - /dsiutil.c(281)
The ct-lib function 'ct_results' returns FAIL for database 'westss.amerttp'.
The errors are retryable. The DSI thread will restart automatically. See ct-
lib messages for more information.
...
E. 96/09/30 14:59:12. ERROR #13043 LTM USER(longdss.amercyp) -
seful/cm.c(2796)
    Failed to execute the 'USE westrs_rssd' command on server 'westss'. See ct-
lib and sqlserver error messages for more information.
E. 96/09/30 14:59:12. ERROR #1028 LTM USER(longdss.amercyp) -
seful/cm.c(2796)
    Message from server: Message: 921, State: 1, Severity: 14 -- 'Database
```

```
'westrs_rssd' has not been recovered yet - please wait and try again.'.
I. 96/09/30 14:59:12. Message from server: Message: 5701, State: 1,
Severity: 10 -- 'Changed database context to 'master''.
I. 96/09/30 14:59:15. LTM for longdss.amercyp connected in passthru mode.
E. 96/09/30 14:59:16. ERROR #13043 USER(westrs_ltm) - seful/cm.c(2796)
    Failed to execute the 'USE westrs_rssd' command on
    server 'westss'. See ct-lib and sqlserver error messages
    for more information.
E. 96/09/30 14:59:16. ERROR #1028 USER(westrs_ltm) - seful/cm.c(2796)
    Message from server: Message: 921, State: 1, Severity: 14 -- 'Database
    'westrs_rssd' has not been recovered yet - please wait and try again.'.
I. 96/09/30 14:59:16. Message from server: Message: 5701, State: 1,
Severity: 10 -- 'Changed database context to 'master''.
...
E. 96/09/30 14:59:23. ERROR #13043 DIST(westss.amerttp) - seful/cm.c(2796)
    Failed to execute the 'USE westrs_rssd' command on server 'westss'. See ct-
    lib and sqlserver error messages for more information.
E. 96/09/30 14:59:23. ERROR #1028 DIST(westss.amerttp) - seful/cm.c(2796)
    Message from server: Message: 921, State: 1, Severity: 14 -- 'Database
    'westrs_rssd' has not been recovered yet - please wait and try again.'.
...
```

Explanation

The Adaptive Server that controls the RSSD was shut down and restarted while the Replication Server was running. The DIST and SQT threads to the databases controlled by the Replication Server were terminated. Replication to those databases was terminated and will not resume even after the RSSD becomes available again.

Running the `admin_who_is_down` command at the Replication Server shows that both DIST and SQT threads are down as follows:

```
Spid Name State Info
-----
DIST      Down   westernDS.westDB
SQT       Down   105:1 westernDS.westDB
```

Solution

To solve the problem:

- 1 At the Replication Server, execute `resume distributor` for each database to resume SQT and DIST threads.
- 2 Run `admin_who_is_down` at each database to verify that the SQT and DIST threads are up.

15020

Replication definition could not be found.

4.6 RSSD SQL Server Bounced.

Symptoms	The following error messages are displayed in the Replication Server error log: <pre>Error #15020: ``%s' doesn't exist."</pre>
Explanation	The replication definition for which you want to create a subscription could not be found.
Solution	<ul style="list-style-type: none">• Verify that a replication definition has been created for the primary version of the table that is to be replicated.• If the primary Replication Server is not the same as the replicate Replication Server, verify that:<ul style="list-style-type: none">• A route has been created from the primary Replication Server to the replicate Replication Server, or• The replication definition has been replicated.• Verify that subscription DDL commands were entered at the replicate Replication Server.

15040

No connection to replicate database.

4.6 RSSD SQL Server Bounced

Symptoms	The following error message is displayed in the Replication Server error log: <pre>Error #15040: "This database '%s.%s' is not controlled by this site."</pre>
Explanation	A connection from the replicate Replication Server to the database has not been created.
Solution	Use the create connection command at the replicate Replication Server to make a connection to the replicate database. To create a connection to an Adaptive Server replicate database, use <code>rs_init</code> .

15052

Primary and replicate database cannot be the same

4.6 RSSD SQL Server Bounced

Symptoms	The following error message is displayed in the Replication Server error log: <pre>Error #15052: "The replicate database '%s.%s' for</pre>
----------	---

```
subscription '%s' is the same as the primary database
for the replication definition '%s'. This is not
allowed."
```

Explanation You cannot create a subscription in which the replicate database is the same as the primary database.

Solution

- 1 Create the replicate table on a different database.
- 2 Execute the create subscription command again.

28028

Connect source not executed

4.6 RSSD SQL Server Bounced

Symptoms The following error message is displayed in the Replication Server error log:

```
Error #28028: "Connect source as user '%s' failed at PRS
'%s' for subscription '%s', for replication definition
'%RS_RSID' with replicate at '%RS_SITEID'."
```

Explanation The replicate Replication Server failed to execute the connect source command at the primary Replication Server.

Solution The subscription recovery daemon should recover the subscription.

29024

No matching function strings found.

Symptoms The error message is similar to the following:

```
Jan 26 11:27:59.300 1993: Error: 29204, Line: 864, File:
'generic/dsi/fstrmap.c' - Cannot find a matching
function string for function 'stocks.rs_insert' and
function class 'rs_sqlserver_function_class'.
```

Explanation This error occurs when Replication Server cannot find a function string to match a function. The DSI thread is suspended.

Solutions To correct the problem:

- 1 Query the RSSD at the primary site of the function string to determine if the function string exists there. Use this query for strings with replication definition scope:

```
select name from rs_funcstrings, rs_functions,
rs_classes, rs_objects
```

```
where rs_funcstrings.classid = rs_classes.classid
and rs_funcstrings.funcid = rs_functions.funcid
and classname = function_class_name
and classtype = 'F'
and rs_functions.objid = rs_objects.objid
and rs_objects.objname = replication_definition
and objtype = 'R' and funcname = function_name
```

If the function string does not exist at the primary site, continue with step 2.

If the function string does exist at the primary site, go to step 3.

- 2 If the function string does not exist at the primary site, create it there. If you omit the output clause from the string, Replication Server generates the default function string.
- 3 If the function string does exist at the primary site, it hasn't been replicated to the Replication Server at the replicate site. Verify that the RSSD at the primary site is being replicated to the replicate site. Refer to Chapter 1, "Identifying the Problem", for details. If changes at the primary site are being replicated to the replicate site, then you may need to wait for the function strings to arrive. Then restart the DSI.
- 4 For `rs_select` and `rs_select_with_lock`, there must be a function string whose input template matches the where clause of the select command. Extend the above query with a join to `rs_systext` to retrieve the input templates of function strings.

If the DSI receiving the error is communicating with a server using a function-string class other than the `rs_sqlserver_function_class`, `rs_default_function_class`, or `rs_db2_function_class`, you may still need to create the function string mentioned in the preceding paragraph.

37022

primary subscribe permission required.

4.6 RSSD SQL Server Bounced

Symptoms

The following error message is displayed in the Replication Server error log:

```
Error #37022: "PRIMARY SUBCRIBE permission is required
to execute the subscription command."
```

Explanation

The user does not have the permission required to create a subscription. You must have primary subscribe, create object, or sa permission at the primary Replication Server.

Solution Use the grant command to change the user's permissions.

37023

37023: create object permission required

Symptoms The following error message is displayed in the Replication Server error log:

```
Error #37023: "CREATE OBJECT permission is required to
execute command."
```

Explanation The user does not have the required permission.

Solution Use the grant command to change the user's permissions to primary subscribe.

Replication Server informational and warning messages

This section contains informational and warning messages displayed in the Replication Server error log. Because informational and warning messages do not have error numbers, they are organized alphabetically.

A cached row for system table was swapped out

3.8 STS Cache Size

Symptoms This Replication Server informational message is displayed many times in the Replication Server error log:

```
I. 96/01/16 15:39:53. A cached row for system table
'rs_columns' was swapped out of the cache in order to
accomodate another row.
```

Explanation When there is no space available in the cache for any more rows in a specific RSSD system table, that system table's oldest rows are swapped out of the cache.

The *sts_cachesize* parameter specifies the maximum number of rows that are cached for a single RSSD system table.

Solution Increase the value of the *sts_cachesize* parameter for the RSSD.

Detecting loss for database

Symptoms

The message below is displayed in the Replication Server error log after you executed a rebuild queues command (which is typically issued after a stable device failed).

You should check for data loss detection messages in the Replication Server error log only after you have issued the rebuild queues command and the data loss detection process has completed. Data loss detection should only be performed when there are primary transactions to be replicated. The RSM heartbeats feature replicates a primary transaction to the destination data server. You can activate the heartbeats feature for the connection to force data loss detection to be performed.

Note Sometimes these messages are not displayed at the end of the error log file, so you must search back through the log for them. If you restart the Replication Server, these messages will be displayed again.

```
I. 96/10/02 13:31:19. DSI: detecting loss for database 'westss.eastlp' from
origin 'westss.norttp'
date = 'Oct  2 1996 12:01:35:313PM',
qid=00000001004fc7ac0000fb0e0016ffffff000000008a0a00c630b20001000000000
03.
  I. 96/10/02 13:31:19. DSI: no transactions for database 'westss.norttp'
from origin 'westss.eastlp' will be applied or deleted from the queue,
because loss was detected.
  I. 96/10/02 13:33:08. LTM for westss.eastlp connected in passthru mode.
  I. 96/10/02 14:31:28. Disconnecting LTM for westss.eastlp. LTM will shut
down
  I. 96/10/02 14:31:32. Resetting LTM Starting Log Position for westss.eastlp
  I. 96/10/02 14:31:34. Shutting down distributor for 106
  I. 96/10/02 14:31:34. The distributor for 'westss.eastlp' is Shutting Down
  I. 96/10/02 14:31:36. Disconnecting LTM for westss.eastlp to Rebuild
  I. 96/10/02 14:31:36. Shutting down the DSI thread for 'westss.eastlp'.
  I. 96/10/02 14:31:36. The DSI thread for database 'westss.eastlp' is
shutdown.
  I. 96/10/02 14:31:36. DSI: enabled loss detection for 'westss.eastlp'.
  I. 96/10/02 14:31:40. Rebuild Queues: Deleting Queue 106:1
  I. 96/10/02 14:31:40. SQM Stopping: 106:1 westss.eastlp
  I. 96/10/02 14:31:40. Rebuild Queues: Done Rebuilding Queue 106:1.
Restarting
  I. 96/10/02 14:31:40. SQM starting: 106:1 westss.eastlp
  I. 96/10/02 14:31:40. Rebuild Queues: Deleting Queue 106:0
  I. 96/10/02 14:31:41. SQM Stopping: 106:0 westss.eastlp
  I. 96/10/02 14:31:41. Rebuild Queues: Done Rebuilding Queue 106:0.
```


Restarting

```

I. 96/10/02 14:31:41. SQM starting: 106:0 westss.eastlp
I. 96/10/02 14:31:47. Starting the distributor for 'westss.eastlp'.
I. 96/10/02 14:31:49. Starting the DSI thread for 'westss.eastlp'.
I. 96/10/02 14:31:50. The DSI thread for database 'westss.eastlp' is
started.
I. 96/10/02 14:31:51. DSI: detecting loss for database 'westss.eastlp' from
origin 'westss.norttp'
date = 'Oct 2 1996 12:01:35:313PM',
qid=00000001004fc7ac0000fb0e0016ffffffff000000008a0a00c630b200010000000000
03.
I. 96/10/02 14:36:22. LTM for westss.eastlp connected in passthru mode.
I. 96/10/02 14:36:22. Restarting LTM for westss.eastlp at the beginning of
the log
I. 96/10/02 14:36:33. LTM for westss.eastlp connected in passthru mode.
I. 96/10/02 14:53:32. DSI: no transactions for database 'westss.norttp'
from origin 'westss.eastlp'
will be applied or deleted from the queue, because loss was detected.
I. 96/10/02 14:53:32. DSI: detecting loss for database 'westss.norttp' from
origin 'westss.eastlp' date = 'Oct 2 1996 12:01:28:313PM',
qid=0000000100764b0e0000fb0c000affffffffff000000008a0a00c6287e00010000000000
01.
I. 96/10/02 14:53:35. DSI: loss detected for database 'westss.norttp' from
origin 'westss.eastlp'.
I. 96/10/02 15:00:02. LTM for westss.eastlp connected in passthru mode.
I. 96/10/02 15:00:17. DSI: loss detected for database 'westss.eastlp' from
origin 'westss.norttp'.

```

Explanation	Replication Server detected loss for a primary Replication Server or database. User transactions from the primary Replication Server or database are no longer accepted. If the replication system has more than one Replication Server, then the detected data loss could be either a DSI loss or an RSI loss. If the replication system has only one Replication Server, the detected data loss is a DSI loss. A DSI loss means that data was not replicated from the Replication Server to the replicate database. An RSI loss means that data was not replicated from one Replication Server to another Replication Server.
Solution	For information on correcting data losses, see Chapter 19, “Replication System Recovery,” of the <i>Replication Server Administration Guide</i> . To prevent data loss, use DSI or RSI save intervals to create a backup strategy with backlogs large enough to keep all messages.

DSI detected rs_update_lastcommit not marked as replicated

Symptoms A warm standby database disconnects and these error messages are displayed in the Replication Server error log:

```
I. 96/01/16 15:39:53. DSI for %RS_SITEID detected that
stored procedure rs_update_lastcommit is not marked as
replicated. Please execute sp_setreplicate and resume
connection
```

When active DSI comes up in presence of the Warm Standby it checks whether rs_update_lastcommit stored procedure is marked as replicated. This DSI detected a problem with replication status of this stored procedure. Please, correct the problem and resume connection

Explanation This error occurs during a warm standby installation. You have used isql to execute the create connection command and did not mark rs_update_lastcommit as replicated.

- Solutions**
- 1 Set replication status on rs_update_lastcommit by marking it with the sp_setrepproc system procedure.
 - 2 Resume connection.

Instead of using the create connection command, you can use rs_init to create a connection to a warm standby database. This message does not occur when you use rs_init to set up the warm standby.

When you use rs_init to set up your active and standby database, the rs_update_lastcommit stored procedure is automatically marked for replication.

Stable storage use is above 75 percent

Growing Stable Queue

Symptoms The total size of all the queues on a Replication Server has passed the set threshold. This warning message is displayed in the Replication Server error log:

```
W. 96/04/10 18:41:12. WARNING #6089 SQM(129:0
TTTdss.eanp) - qm/sqmsp.c(1317) WARNING: Stable Storage
Use is Above 75 percent
```

Explanation A Replication Server has two default thresholds (75% and 90%) for stable devices. A warning message is issued each time these thresholds are exceeded.

Solutions	<p>You should follow these procedures when the 75% and 90% thresholds have been exceeded:</p> <ol style="list-style-type: none"> 1 Verify that the replicate database is running. 2 If all other components are functional and the inbound queues are growing, shut down the corresponding Adaptive Server and check for an open transaction. 3 If the outbound queues are growing, check the DSI save interval in conjunction with low block usage.
-----------	---

RepAgent error message descriptions

This section contains RepAgent error messages that are contained in the Adaptive Server error log. If the RepAgent error also occurs in the LTM, the LTM error number is included in the heading. If you do not find an LTM error number in this section, see “LTM-only error message descriptions” on page 61.

9202: nested stored procedures not allowed (LTM 1033)

6.5 Nested Replicated Stored Procedure

Symptoms	<p>These error messages are displayed in the Adaptive Server error log.</p> <pre> 96/03/04 14:01:53.34 RepAgent (10) Error: 9202, Severity: EX_CMDFATAL, State: 1 Nested replicated stored procedure detected. Transaction log may be corrupt. Please contact SYBASE Technical Support. W. 96/03/04 14:01:53. WARNING #1033 logscan thread(NYDS.nydb1) - /ltmscan.c(4689) Nested replicated stored procedure is not allowed. The procedure name = 'south_nested', pid = '456', rid = '2'. </pre>
Explanation	<p>A nested stored procedure is called from within another stored procedure. The stored procedure that calls the nested stored procedure is called the outer stored procedure.</p> <p>If stored procedures with nested stored procedures are marked for replication with <code>sp_setrepproc</code>, then:</p>

- The RepAgent shuts down.
- The RepAgent forwards only the outer stored procedure call to the Replication Server.
- An error message is displayed in the Adaptive Server error log.

Solution

Do not use nested replicated stored procedures.

To solve this problem:

- 1 Skip the nested stored procedure transaction:
 - a Find the page of the secondary truncation point, by executing:

```
dbcc gettrunc
```
 - b Set a new secondary truncation point on the page after the nested stored procedure transaction by executing:

```
dbcc settrunc ('ltm', 'pageid', pageid)
```

where *pageid* is the id for the page *after* the current page you retrieved in Step a.
 - c Reset the locator by executing:

```
rs_zeroltm
```
- 2 Reapply only the nested stored procedure transaction.
- 3 Restart RepAgent.

9210: network problems (LTM 2027)

Symptoms

These error messages are displayed in the Adaptive Server error log:

```
95/09/07 09:41:26.34 RepAgent (10) Error: 9210,  
Severity: EX_CMDFATAL, State: 1  
Fatal error from CT-Lib.
```

```
Open Client Client-Library error: Error: 84083974,  
Severity: 5 -- ct_results(): unable to get layer message  
string: unable to get origin message string: Net-Library  
operation terminated due to disconnect.
```

Explanation	<p>You see these error messages when there are network problems; for example, when a RepAgent could not initialize a CT-Lib connection or a Replication Server returned an unexpected error. Network problems can also cause RepAgent error 9212 (which is caused by a CT-Lib error).</p> <p>After you fix the network problems, execute the <code>admin who_is_down</code> command on the Replication Server. You see that the REP AGENT USER threads are down.</p> <p>In RSM, “suspect status” (yellow) is displayed for the Replication Server, while “normal” status (green) is displayed for the RepAgent.</p>
Solution	Fix the network problems and restart RepAgent.

9215: log truncated (LTM 1079 with ASE 624)

Symptoms	<p>This error message is displayed in the Adaptive Server error log:</p> <pre>97/10/17 11:57:19.34 RepAgent (10) Error: 9215, Severity: EX_USER, State: 1 The Secondary Truncation Point is not valid. Please make sure that the Secondary Truncation Point is valid, and then restart the Rep Agent Thread for this database.</pre>
Explanation	<p>This error occurs when the log is truncated past the secondary truncation point and the RepAgent tries to reconnect to the database. RepAgent requests the log page number from the Replication Server and then accesses that page in the database. Because the log was truncated, the RepAgent will not start.</p> <p>Attempts to retrieve rows from the page via RID fail because the requested RID is a higher value than the last RID on the page. For example, truncating the log frees only those log pages that the LTM has sent to the Replication Server with confirmation.</p>
Solutions	<p>Because you truncated the log, you must reset the information about the database log in the <code>rs_locator</code> table within the RSSD as follows.</p> <ol style="list-style-type: none"> Execute the following stored procedure in the RSSD: <pre>> rs_zeroltm primss, primdb > go</pre> Activate the secondary truncation point again in the primary database. <pre>> dbcc settrunc(ltm, valid) > go</pre> Restart RepAgent.

Check for open transactions by comparing the output of dbcc gettrunc with the content of the rs_locator tables attribute locator. Truncating the transaction log may have caused orphaned transactions. Orphaned transactions occur when the transaction's begin statement is transferred, but the commit or rollback statement is accidentally deleted during log truncation.

9220: connect only one RepAgent to a Replication Server (LTM 1078)

Symptoms

When you start the RepAgent thread, the process terminates; you see this error message in the Adaptive Server error log:

```
95/09/22 08:51:43.34 RepAgent (10) Error: 9202,
Severity: EX_User, State: 1
The Rep Agent Thread context for database 'westDB' has
already been reserved by another session.
I. 95/09/22 08:51:43. Connection to SQL Server
'westernDS' succeeded.
H. 95/09/22 08:51:43. THREAD FATAL ERROR #1078 logscan
thread(westernDS.westDB) - /ltmscan.c(734) This client
lib or server error is fatal. The server the thread
connects to is 'westernDS'.
E. 95/09/22 08:51:43. ERROR #2028 logscan
thread(westernDS.westDB) - /ltmscan.c(734) Message from
server: Message: 9121, State: 2, Severity: 16 -- 'The
log transfer context for the current database is already
reserved by SQL Server process 8. The log transfer
context for the current database is not reserved.'
```

Explanation

Only one RepAgent thread can connect to a database at any one time. Because the RepAgent abnormally terminated, the RepAgent thread process still exists although it is in the sleep mode. The RepAgent thread wakes after the time specified in the *scan_retry* parameter elapses. Since the RepAgent has already terminated, the RepAgent thread process terminates when it wakes.

Solution

After shutting down a RepAgent thread, wait for two minutes or twice the value of the *scan_retry* parameter, before restarting it. The default for the *scan_retry* parameter is 60 seconds.

If the problem persists:

- 1 Log in to the Adaptive Server.
- 2 Terminate the process using the process ID number found in the error log. In this example, the process ID number is 8.

```

1> select spid, program_name, hostname from
sysprocesses
2> where spid = 8
3> go

spid program_name hostname
-----
8 repagent bs10h.bs1
(1 row affected)
1> kill 8
2> go

```

- 3 Run `sp_who` to check whether the process is terminated.
- 4 If the process still exists, recycle the Adaptive Server and the RepAgent thread.

LTM-only error message descriptions

This section contains error messages that occur only with the LTM. LTM errors are recorded in the LTM error log. These errors are organized in ascending numerical order.

2028: invalid enable marker in log

Symptoms	<p>These error messages are displayed in the LTM error log:</p> <pre> E. 96/09/11 16:40:09. ERROR #2028 logscan thread(ost_ra.pub) - /ltmscan. c(2616) Message from server: Message: 9141, State 1, Severity 16 -- 'The stored proc. 'rs_marker' associated with the EXECBEGIN log record in database 'pub', XA CT ID Rid pageid = 0x1cb; row num = 0x2, RID Rid pageid = 0x1cb; row num = 0x3, was dropped after the log record was written. The EXECBEGIN log record is not replicated.' </pre>
Explanation	<p>If you reinstall the active database while replication to a warm standby is in progress, an invalid enable marker may be left in the log. During the reinstallation, system-defined stored procedures (such as <code>rs_marker</code> and <code>rs_update_lastcommit</code>) are re-created and overwrite the stored procedures created during the initial installation. As a result, the <code>rs_marker</code> value changes and the standby database waits for a marker that never arrives.</p>

Solution

Drop the active and standby connections and reinstall the databases.

Route Problems

This chapter describes problems that may occur when you are creating, altering, or dropping routes and how to troubleshoot these problems.

Name	Page
Overview	63
Problems with creating routes	65
Problems with altering routes	70
Problems with dropping routes	70

If you cannot find the solution in this chapter

If an error message is displayed in an error log and you cannot find it in this chapter, refer to:

- “Analyzing error logs” on page 10, or
- Chapter 3, “Common Errors”

If an error message is not displayed, refer to Chapter 1, “Identifying the Problem”.

Overview

Use the procedures in this chapter to successfully solve problems that occur after executing a create route, alter route, or drop route command. Each procedure requires you to obtain and analyze information from the Replication Server error logs, execute the `rs_helproute` stored procedure, or both.

To learn how to use the `rs_helproute` stored procedure and interpret its results, see “Using the `rs_helproute` stored procedure” on page 64 and the *Replication Server Reference Manual*.

A route is one-way message stream from a source Replication Server to a destination Replication Server. Routes carry data modification commands (including those for RSSDs) and replicated functions or stored procedures between Replication Servers.

There are two types of routes—direct and indirect. A direct route sends messages directly from source to destination Replication Servers, with no intermediate Replication Servers. An indirect route sends messages through one or more intermediate Replication Servers.

Routes are created, altered, and dropped according to the following general procedure. The `rs_marker` stored procedure contains the full create route, alter route, or drop route command.

- The source Replication Server executes the `rs_marker` stored procedure against its RSSD.
- The RSSD's RepAgent picks up the stored procedure execution and forwards it to the source Replication Server.
- The source Replication Server appends the stored procedure execution to its inbound queue and forwards it to the destination Replication Server. The identity of the destination Replication Server is obtained from parameters in the stored procedure.
- The destination Replication Server processes the stored procedure execution, which materializes subscriptions to the primary Replication Server RSSD `rs_routes`, `rs_subscriptions`, `rs_rules`, and `rs_locator` system tables.

Using the `rs_helproute` stored procedure

Routes go through various states while being created, altered, or dropped. Execute the `rs_helproute` stored procedure on the RSSD at the source or the destination Replication Server to see the current state of the route:

```
rs_helproute [replication_server]
```

`replication_servers_helproute`

is an optional parameter indicating the name of a Replication Server. If you enter this parameter, information is given for routes to and from the specified Replication Server. Otherwise, information is provided for all routes to and from the current Replication Server. For each route,

returns:

- Route status

Status identifies the state of the route protocol. The status of a route is different at the source and destination Replication Servers. You analyze the route status at the source and destination Replication Servers to determine the problem.

- A list of system table subscriptions that are currently being processed

Incomplete materialization or dematerialization of system table subscriptions is one of the most common problems.

Because creating, altering, and dropping routes includes processing system table subscriptions, the list that `rs_helproute` returns helps you determine which subscriptions prevent you from proceeding to the next step in the process.

If you are creating a route, the list includes the system table subscriptions that are being created. If you are dropping a route, the list includes the system table subscriptions that are being dropped.

If no system table subscriptions are listed for a route, there are no problems with the system table subscriptions.

Problems with creating routes

This section explains how to troubleshoot problems that can occur when you create a route.

When a route is created:

- The source Replication Server connects to the destination Replication Server. The source Replication Server needs the correct user name and password to connect to the destination Replication Server.
- The source Replication Server sends a create route message to the destination Replication Server.
- The destination Replication Server receives this message and starts creating subscriptions to the source Replication Server system tables.
- When the system table subscriptions are created, the process is complete.

Common problems

Table 4-1 provides solutions to common problems that can occur while creating a route.

Table 4-1: Creating a route—problems and suggested actions

Problem	Suggested Action
The destination Replication Server has no entry in its interfaces file for the source Replication Server.	Include this name in the interfaces file.
You are trying to create a route that has just been dropped.	Wait until the destination Replication Server has cleaned up its route to the source Replication Server. Then the destination can accept the new route. You may need to resume this route from the source Replication Server—the source may have suspended this route after being refused connection by the destination.
You attempt to create an indirect route but the log at the intermediate site contains the message “route is needed to reach the destination.”	Drop the route from the source to the destination using the drop route command with the with nowait clause. Next, create a route from the intermediate site to the destination site. Then, use the create route command with the set next site clause to create an indirect route from the source to the destination.

Messages in the error log at the source Replication Server

Table 4-2 provides the solutions to the problems described by error messages in the source Replication Server’s error log.

Table 4-2: Messages at the source Replication Server

Message	What Does It Mean?	Suggested Actions
RSI <destination name>: Login incorrect RSI <destination name>: Trying to connect	The source Replication Server tries to connect to the destination Replication Server with the username/password supplied with the create route command, but the destination Replication Server does not recognize this username/password.	<ol style="list-style-type: none"> 1. Perform one of these tasks: <ul style="list-style-type: none"> • Use the create user or alter user command to create or alter this username/password at the destination Replication Server. • Use the alter route command with the set username clause to alter the username/password for this route. • Drop the route using the drop route command with the with nowait clause. Then re-create the route. 2. Use the resume route command to resume this route at the source Replication Server.
cm: Could not find interfaces entry for <destination name>	The destination Replication Server name is not found in the interfaces file.	<p>Perform one of these tasks:</p> <ul style="list-style-type: none"> • Add the Replication Server name in the interfaces file, or • If the user did not mean to create a route to this destination, use the drop route command with the with nowait option to drop the route.

Output from *rs_helproute* at the source Replication Server

Try using a table format. This section describes the output when *rs_helproute* is executed at the source Replication Server.

Message:

```
Route is being created. Source RS has not yet attempted
to send the protocol message to the destination RS.
```

Description:

The source Replication Server could not execute the `rs_marker` stored procedure against its RSSD.

Suggested Action:

Restart the Replication Server.

If restarting does not fix the route, check the RSSD and any related messages in the error logs. There may be a message in the error logs indicating some problem with the RSSD.

Message:

```
Route is being created. Either (1) RS is waiting for a
route protocol message from the RSSD Replication Agent
or (2) the RSSD Replication Agent inbound queue is not
being processed.
```

Description: This message is the result of one of the following problems:

- *Problem:* A message may be stuck between the RSSD log and the Replication Server inbound queue. The sender Replication Server executed the stored procedure, but the RSSD replication agent cannot forward it to the Replication Server.

Suggested Action: Check if the RSSD replication agent is experiencing problems sending messages to the Replication Server.

If not, restart the RSSD replication agent.

- *Problem:* A message may be stuck between the inbound and outbound queues in the Replication Server.

Suggested Action: Check if the Replication Server's inbound queues are full or the distributor for the RSSD is not running.

If neither of these problems exist, restart the RSSD replication agent.

Message:

```
Route is being created. Create route protocol message
is waiting to be delivered to the destination.
```

Description: The sending Replication Server is not processing its outbound queue.

Suggested Action: Resume the route, if:

- The RSI experienced problems connecting to the destination Replication Server.
- The route was suspended.

Check to see if:

- The destination Replication Server is up.
- The RSI for the destination Replication Server running at the source Replication Server.

Check to see if:

- The destination Replication Server is up.
- The RSI for the destination Replication Server running at the source Replication Server.

Message:

```
Route is being created. Destination has not yet finished
creating system table subscriptions.
```

Description: The destination Replication Server is still creating subscriptions to system tables.

Suggested Action: Make sure the destination Replication Server is running. Resume the route, if:

- The RSI experienced problems connecting to the destination Replication Server.
- The source could not send messages to the destination
- The system suspends the route due to an error or the user explicitly suspends it.

Check the list of system table subscriptions that still need to be created.

Output from `rs_helproute` at the destination server

This section describes the output when `rs_helproute` is executed at the destination Replication Server.

Message:

```
This site received the create route protocol message
from the source RS and is in the process of creating
system table subscriptions.
```

Description: The destination Replication Server is completing its system table subscriptions.

Suggested Action: Check the list of system table subscriptions that still need to be created. If the destination Replication Server is suspended in this state, it could mean there is a problem with the system table subscriptions.

Problems with altering routes

If you are experiencing problems using the alter route command:

- Quiesce the replicated data system prior to executing the alter route command. For more information, see the *Replication Server Administration Guide*.
- If you are changing a direct route to an indirect route, or if you are changing the intermediate site in an indirect route, you must create a route from the new intermediate site to the destination site.
- To determine the success of an alter route command, execute `rs_helproute` at the destination Replication Server after you alter the route. Make sure that the output reflects the route change. If the route change did not succeed, follow the troubleshooting procedures in “Output from `rs_helproute` at the source Replication Server” on page 67.

Problems with dropping routes

This section explains how to troubleshoot problems that can occur when you drop a route. By default, the source Replication Server waits until the destination Replication Server has cleaned up its part of the route. You can use the `nowait` option to drop a route without waiting for the destination Replication Server to clean up its part of the route.

Dropping routes

When a route is dropped:

- The source Replication Server drops all the system table subscriptions created by the destination.

- The source Replication Server sends a drop route protocol message to the destination Replication Server.
- After this message is delivered to the destination Replication Server (or to the next site, if the route is an indirect route), the source Replication Server drops the route by deleting the entry from its `rs_routes` system table and, for direct routes, deleting the outbound queue.
- The destination Replication Server receives the drop route message and cleans up its part of the route by deleting the locator for the source Replication Server from its `rs_locator` system table.

Output from `rs_helproute` at the source Replication Server

This section describes the output when `rs_helproute` is executed at the source Replication Server.

Message:

```
Route is being dropped. System table subscriptions are
being dropped.
```

Description: The destination Replication Server is trying to drop the system table subscriptions.

Suggested Action: Check the list of system table subscriptions that still need to be dropped.

Message:

```
Route is being dropped. Waiting for the system table
subscriptions to be dropped.
```

Description: The system is in this state until the subscriptions are dropped.

Suggested Action: The Replication Server should not be in this state for a long period of time. If it is, it may indicate a subscription problem. Refer to Chapter 5, “Subscription Problems”.

Message:

```
Route is being dropped. Source RS has not yet attempted
to send the drop route protocol message to the
destination RS.
```

Description: The sender Replication Server could not execute the `rs_marker` stored procedure against its RSSD.

Suggested Action: Restart the Replication Server. Check the RSSD and any related messages in the error logs. There may be a message in the error logs indicating some problem with the RSSD.

Message:

Route is being dropped. Drop route protocol message should be either in RSSD Replication Agent queue or in the inbound queue.

Description: This message is the result of one of the following problems:

- *Problem:* A message may be stuck between the RSSD log and the Replication Server inbound queue. The sender Replication Server executed the stored procedure, but the RSSD replication agent cannot forward it to the Replication Server.

Suggested Action: Check if the RSSD replication agent is experiencing problems sending messages to the Replication Server.

If not, restart the RSSD replication agent.

- *Problem:* A message may be stuck between the inbound and outbound queues in the Replication Server.

Suggested Action: Check if the Replication Server's inbound queues are full or the distributor for the RSSD is not running.

If neither of these problems exist, restart the RSSD replication agent.

Message:

Route is being dropped. Drop route protocol message is waiting to be delivered to the destination.

Description: The sender Replication Server is not processing its outbound queue.

Suggested Action: Is the destination Replication Server operating? Did the RSI (Replication Server Interface) experience problems trying to connect to the destination Replication Server? If so, it may be necessary to resume this route.

Output from `rs_helproute` at the destination server

Try using a table format. This section describes the output when `rs_helproute` is executed at the destination Replication Server.

Message:

Route is being dropped.

Description: The route is being dropped by the source Replication Server. The destination Replication Server is not very active while a route is being dropped. All of the state changes are seen at the source Replication Server only.

Suggested Action: Look at the output from `rs_helproute` at the source Replication Server.

Dropping routes using the *nowait* option

Warning! Use the `with nowait` clause only when the destination is temporarily irrecoverable and you must drop the route without waiting for the destination Replication Server to become available.

This command instructs the source Replication Server to drop its route to the destination without informing the destination Replication Server. You must then clean up the destination Replication Server when it becomes available again.

When you execute `drop route` with the `with nowait` option:

- The source Replication Server drops all the system table subscriptions created by the destination.
- The source Replication Server cleans up the route by deleting the route entry from its `rs_routes` system table and, for direct routes, dropping the outbound queue.

Output from `rs_helproute` at the source Replication Server

The output is the same as the `drop route` command, except that the route status reads:

```
Route is being dropped with nowait. ...
```

instead of:

```
Route is being dropped. ...
```

Cleaning up the destination Replication Server

To clean up the destination Replication Server:

- 1 Verify that all the messages from the source Replication Server are gone from the DSI queues.
- 2 From the `rs_subscriptions` system table, delete all the system table subscriptions entered by this site when the route was created. Delete rows from the `rs_rules` system table for these subscriptions.

- 3 Clean up user subscription entries from the `rs_subscriptions` and `rs_rules` system tables.

- 4 Execute this command against each system table:

```
delete from system_table
       where prsid = sourceRS_siteid
```

system_table is the name of the system table (`rs_subscriptions` or `rs_rules`).

sourceRS_siteid is the site ID of the source Replication Server.

For more information about the delete command, see the *Adaptive Server Reference Manual*.

- 5 Delete the entry for this route from the `rs_routes` system table.
- 6 Delete the entry for the source Replication Server from the `rs_locator` system table.

Subscription Problems

This chapter describes problems that occur when subscriptions are created, altered, or dropped.

Name	Page
Overview	75
Using check subscription	76
Materialization problems	79
Dematerialization problems	95

If you cannot find the solution in this chapter

If an error message is displayed in an error log and you cannot find it in this chapter, refer to:

- “Analyzing error logs” on page 10, or
- Chapter 3, “Common Errors”

If an error message is not displayed, refer to Chapter 1, “Identifying the Problem”

Overview

This chapter explains how to use the check subscription command to display the status of a subscription and how to solve some common problems when using atomic, nonatomic, and bulk materialization. Use `rs_helppub` and `rs_helppubsub` to find the publications and articles that a subscription is using. For more information about `rs_helppub` and `rs_helppubsub`, see Chapter 11, “Managing Subscriptions,” in the *Replication Server Administration Guide*. The general procedure for troubleshooting subscription problems is in “Troubleshooting subscription failures” on page 14.

The replication process begins with subscription *materialization*, which is the process by which data is first copied to the destination database. When you no longer want a subscription replicated to a destination database, you dematerialize the subscription at the destination database. *Dematerialization* is the process by which data is deleted from the destination database.

Note If you are using a replication agent, your subscription materialization process may differ from the process described in this chapter. See your replication agent documentation for the replication agent-specific subscription materialization process.

Subscription failures occur when materialization or dematerialization fails. Usually, the person who is conducting the materialization or dematerialization monitors the operation and reports any problems.

Subscription problem symptoms are easily identified and include:

- Materialization failure – no data in the subscription’s replicate table at the destination database, the status for subscriptions at the primary and replicate Replication Servers is not valid, or materialization is taking longer than is reasonable.
- Dematerialization failure – data still exists in the subscription’s replicate table at the destination database, subscriptions still exist at the primary or replicate Replication Servers, or dematerialization is taking longer than is reasonable.

If a subscription problem caused the DSI thread for the replicate database to abnormally terminate, you can restart it using the resume connection command. For more information, see “DSI down or suspended” on page 121.

Using *check subscription*

If all of the servers and Replication Server threads are running, but the subscription is not being created or dropped, use the check subscription command, which returns messages that describe the status of the materialization process.

Execute the check subscription at the primary and replicate Replication Servers. If the primary and replicate Replication Servers are the same Replication Server, you need to execute check subscription only once; check subscription displays separate status messages.

Use check subscription for publication to find out the materialization status of publication subscriptions. For atomic, nonatomic, or incremental publication subscriptions, check subscription for publication shows the status of only the current article subscription being processed; whereas for bulk or no materialization, the status of all article subscriptions is shown because all article subscriptions are processed together.

Table 5-1 and Table 5-2 list the messages that can be returned by executing check subscription at the primary and replicate Replication Servers.

Table 5-1: check subscription status at replicate Replication Server

Status	Message
Invalid	<subscription name> doesn't exist.
Removing	REMOVING subscription <subscription name> from system tables at the Replicate.
Dematerializing	Subscription <subscription name> is DEMATERIALIZING at the Replicate.
Valid	Subscription <subscription name> is VALID at the Replicate.
Validating	Subscription <subscription name> is VALIDATING at the Replicate.
Materialized	Subscription <subscription name> has been MATERIALIZED at the Replicate.
Active	Subscription <subscription name> is ACTIVE at the Replicate.
Activating	Subscription <subscription name> is ACTIVATING at the Replicate.
Qcomplete and Active	Subscription <subscription name> is ACTIVE at the Replicate and Materialization Queue has been completed.
Qcomplete	Materialization Queue for Subscription <subscription name> has been completed.
Active and Qcomplete	Subscription <subscription name> is ACTIVE at the Replicate, but Materialization Queue for it has not been completed.
Defined	Subscription <subscription name> has been defined at the Replicate.
Error	Subscription <subscription name> has experienced an unrecoverable error during Materialization or Dematerialization. Please consult the error log for more details.
Pending	Other subscriptions are being created or dropped for the same replication definition/database. Subscription <subscription name> will be processed when previous requests are completed.
Recovering	Subscription <subscription name> has experienced a recoverable error during Materialization or Dematerialization. It will be recovered by Subscription Daemon (dSub).

Table 5-2: check subscription status at primary Replication Server

Status	Message
Invalid	<subscription name> doesn't exist.
Dematerializing	Subscription <subscription name> is DEMATERIALIZING at the Primary.
Valid	Subscription <subscription name> is VALID at the PRIMARY.
Active	Subscription <subscription name> is ACTIVE at the PRIMARY.
Activating	Subscription <subscription name> is ACTIVATING at the PRIMARY.
Defined	Subscription <subscription name> has been defined at the PRIMARY.

Materialization problems

This section provides information for troubleshooting materialization problems. It provides information on common materialization problems as well as detailed information about each step in the materialization process. Because the materialization process is slightly different for atomic, nonatomic, and bulk materialization methods, they are covered in separate sections.

When publication subscriptions are materialized atomically, nonatomically, or incrementally, only one article subscription is processed at a time. When publication subscriptions are materialized using the bulk or no materialization methods, all article subscriptions are processed together.

If, when you are materializing a new subscription, there are other materializing or dematerializing subscriptions for the same replication definition and replicate database, the new subscription is assigned a Pending status (the recovering column in `rs_subscriptions` is set to 2). If you execute `check subscription`, the subscription's status is returned as Pending at the replicate Replication Server and Invalid at the primary Replication Server.

If the primary RSSD is unavailable or any other problems occur, the recovering column in `rs_subscriptions` is set to 1 to denote that the subscription requires recovery.

Subscriptions with a pending status are processed one at a time, in the order entered. The `request_date` column in the `rs_subscriptions` table indicates the time the subscription request was entered.

Common problems

Problems described in this section are due to:

- Incorrect or missing login account and permissions
- Schema inconsistency
- Missing interfaces file entries

Error messages that identify these problems are displayed in the replicate Replication Server error log.

Incorrect or missing login account and permissions

Many problems with materialization are due to incorrect or missing permissions in primary or replicate databases.

User requirements for creating subscriptions

Verify that the login accounts for the user creating a subscription meet the following requirements:

- The user's login name and password is the same at the replicate Replication Server, the primary Replication Server, and the primary data server.
- The user has been added to the primary database.
- The user has select permission on the primary table.
- The user has execute permission on the rs_marker stored procedure.

This requirement applies to a configuration that uses RepAgent or the LTM. If you are using a replication agent for a foreign data server, you may want to check your replication agent documentation for a similar requirement

The default installation scripts grant execute permission on rs_marker to "public," so this should not be an issue unless you did not use the installation scripts.

- The user has at least create object permission at the replicate Replication Server.
- The user has at least primary subscribe permission at the primary Replication Server.

If you cannot give the user the required accounts and permissions, drop the subscription without purge and have a different user create it.

Maintenance user permissions

Make sure that the maintenance user for the replicate database has update, delete, insert, and select permission on the replicate table. Lack of update permission causes errors in the replicate database when the DSI attempts to apply transactions in the materialization queue.

Depending on error action assignments, some transactions may be rejected and recorded in the exceptions log.

Schema inconsistency

The problems in this section result from schema inconsistencies at the primary database, Replication Server, and the replicate database.

Conflicting table or view in the primary database

A table or view with the same name and columns as the replication definition has been created at the primary database. If you have created custom function strings, make sure they match the replicate table at the replicate database. This can cause errors at the primary database as a result of executing select during subscription materialization.

Conflicting table or view in the replicate database

A table or view with the same name and columns as the replication definition has been created in the replicate database. If you create custom function strings, make sure they match the replicate table at the replicate database.

This would cause errors returned by the replicate database when the data server interface attempts to apply transactions in the materialization queue. Some transactions may be rejected and placed into the exceptions log, depending on the error action assignments.

Missing *rs_select* function string

If the primary database does not use one of the system-provided function-string classes (*rs_sqlserver_function_class*, *rs_default_function_class*, *rs_db2_function_class*) or a function-string class that is a child of *rs_default_function_class* or *rs_db2_function_class*, make sure there are function strings for the *rs_select* function that match the where clause of the create subscription or define subscription command.

Missing interfaces file entries

Login attempts may fail due to missing interface file entries. The interfaces file used by the replicate Replication Server should contain:

- An entry for the primary Replication Server
- An entry for the primary data server

Atomic materialization

The atomic materialization method is invoked using the default version of the create subscription command. If there are no other subscriptions for the same replication definition and replicate database, the subscription is defined at the primary Replication Server. If the primary Replication Server is unavailable or any other problems occur, the recovering column in the *rs_subscriptions* system table is set to 1 to denote that the subscription requires recovery.

After the definition stage is complete, the replicate Replication Server builds the materialization queue for the subscription. You can monitor this activity with the `admin who, sqm` command.

Replication Server executes the *rs_select_with_lock* function to select subscription rows from the primary site. After the materialization queue is built, the replicate Replication Server sends an activation request to the primary Replication Server. This request is passed through the primary database via the *rs_marker* system function. When the primary Replication Server receives the activation request, the subscription is marked valid. All updates following the request are sent to the replicate database if they match the subscription.

The primary Replication Server sends the activation request back to the replicate Replication Server, where it is inserted into the DSI queue for the replicate database. When the DSI queue processes the request, the subscription status is changed to active.

The DSI thread also switches over to the materialization queue from its regular outbound queue for the site. The admin who, dsi command indicates which queue the DSI thread is processing. The materialization queue is applied to the replicate database. After the materialization queue is processed, the subscription is marked Valid and the materialization is complete.

If the replicate Replication Server runs out of stable queue segments:

- 1 Add a new partition to the replicate Replication Server.
- 2 If no partition is available, drop the subscription without purge.
- 3 Use nonatomic materialization.

Note To estimate how much space a subscription needs, refer to the *Replication Server Design Guide*.

If the replicate database log cannot hold all of the data for the subscription in a single transaction:

- 1 Drop the subscription without purge.
- 2 Use the incremental version of atomic materialization.

Table 5-3 describes solutions for atomic materialization problems based on the status returned by check subscription. The Replicate Status column shows the subscription status returned by check subscription at the replicate Replication Server. The Primary Status column shows the subscription status returned by check subscription at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Table 5-3: Atomic materialization problems

Replicate Status	Primary Status	Subscription State	Suggested Actions
Defined/ Pending	Invalid	Waiting for other subscriptions for the same replication definition and replicate database to be processed.	Check for other subscriptions being created and dropped for the same replication definition and replicate database. If there are no other subscriptions, wait five minutes.
Defined/ Recovering	Invalid	Cannot connect to the primary Replication Server to define the subscription.	Check the replicate Replication Server error log for messages. Make sure the user creating the subscription has the same login name and password at the primary Replication Server and the replicate Replication Server. Make sure the user has at least primary subscribe permission at the primary Replication Server.

Replicate Status	Primary Status	Subscription State	Suggested Actions
Defined/ Recovering	Defined	Cannot build the materialization queue.	<p>Verify that the primary data server is running.</p> <p>Make sure the user creating the subscription has the same login name and password at the primary data server and replicate Replication Server, select permission on the primary table, and execute permission for rs_marker.</p> <p>Investigate whether the replicate Replication Server is out of stable queue segments. Use the admin disk_space command.</p> <p>Use admin who, dsi and admin who, sqm to monitor the queues.</p>
Defined	Defined	Building materialization queue.	<p>Building the queue may take some time. You may need to wait until this process is complete.</p> <p>Use admin who, sqm to monitor materialization.</p> <p>Investigate whether the primary Replication Server is out of stable queue segments. Use the admin disk_space command.</p> <p>Check if rows are being selected using a holdlock at the primary database.</p>

Replicate Status	Primary Status	Subscription State	Suggested Actions
Qcomplete	Defined	Waiting for the activation request to get to the primary Replication Server. Materialization queue has been built.	Verify that the RepAgent for the primary database is running. Investigate whether the primary Replication Server is out of stable queue segments.
Qcomplete/ Recovering	Defined/ Valid	The replicate Replication Server stopped after the materialization queue was completed.	Wait. The status should return to Qcomplete at the replicate Replication Server.
Qcomplete	Activating	The activation request is being processed at the primary Replication Server.	Wait. The status at the primary Replication Server should change to valid.
Qcomplete	Valid	The subscription is complete at the primary Replication Server. Waiting for activation request or applying the materialization queue at the replicate Replication Server.	Check the route between the primary and the replicate Replication Server. Check the replicate Replication Server DSI thread for the replicate database. Investigate whether the replicate Replication Server ran out of queue segments. Use admin who, dsi and admin who, sqm to monitor the queues. If admin who, dsi shows that the DSI is processing the materialization queue, wait until the queue has been processed. The length of the wait depends upon the size of the queue.
Valid	Valid	Complete.	None.

Nonatomic materialization

To create a subscription using the nonatomic materialization method, execute the `create subscription` command with the `without holdlock` option at the replicate Replication Server. The subscription is saved in the replicate Replication Server RSSD. If there are no other subscription requests for the same replication definition and replicate database, the subscription is defined at the primary Replication Server.

After the definition stage is complete, the replicate Replication Server sends an activation request to the primary Replication Server. The replicate Replication Server immediately starts to build the materialization queue for the subscription. After the materialization queue is built, the subscription status becomes `Qcomplete`. The replicate Replication Server sends a validation request to the primary Replication Server through the primary database. The `admin who` command can be used to monitor this queue.

When the activation request arrives at the primary Replication Server, the subscription status becomes `Active`. All updates following the request are sent to the subscription.

The primary Replication Server sends the activation request back to the replicate Replication Server. When the `DSI` at the replicate Replication Server receives the request, the subscription status becomes `Active` and the transactions in the materialization queue are applied to the replicate database. If the materialization queue has not been built yet, the status returned by the `check subscription` command is `Active` and not `Qcomplete`. If the materialization queue has been built, the status is `Qcomplete` and `Active`. The `DSI` thread switches over to the materialization queue from its regular outbound queue for the site. The `admin who, dsi` command shows which queue the `DSI` thread is processing.

After the contents of the materialization queue are applied to the replicate database, the subscription status becomes `Materialized`.

While the replicate Replication Server is applying inserts from the materialization queue, the validation request is moving from the primary database log, through the `RepAgent`, to the primary Replication Server.

Once the validation request arrives at the primary Replication Server, the subscription status becomes Valid at the primary Replication Server and the request is forwarded to the replicate Replication Server. The subscription status becomes Valid at the replicate Replication Server after the materialization queue is applied and the validation request reaches the beginning of the DSI queue.

Warning! The subscription data may be inconsistent at the replicate database from the time the DSI thread starts applying the materialization queue until the subscription is validated at the replicate Replication Server. This is the result of not using a holdlock while selecting the subscription data from the primary database. Once the subscription status becomes valid, however, the replicate data is consistent with the primary data.

Table 5-4 describes solutions for nonatomic materialization problems based on the status returned by check subscription. The Replicate Status column shows the subscription status returned by check subscription at the replicate Replication Server. The Primary Status column shows the subscription status returned by check subscription at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Table 5-4: Nonatomic materialization problems

Replicate Status	Primary Status	Subscription State	Suggested Actions
Defined/ Pending	Invalid	The replicate Replication Server is waiting for other subscriptions for the same replication definition and replicate database to be created or dropped.	Look for other subscriptions being created or dropped for the same replication definition and replicate database. If there are no other subscriptions, wait for five minutes.
Defined/ Recovering	Invalid	The replicate Replication Server cannot connect to the primary Replication Server to define the subscription.	Check the replicate Replication Server error log for messages. Make sure the user creating the subscription has the same login name and password at the primary Replication Server and the replicate Replication Server. The user should have at least primary subscribe permission on the primary Replication Server.

Replicate Status	Primary Status	Subscription State	Suggested Actions
Defined	Defined	The replicate Replication Server is building the materialization queue. The primary Replication Server is waiting for the activation request.	Verify that the materialization queue is being built. Verify that the primary Replication Server is up and that the SQM, SQT, and DIST threads for the primary database are running.
Defined/ Recovering	Defined	The replicate Replication Server cannot build the materialization queue or cannot send the activation request to the primary Replication Server.	Verify that the primary data server is running. Make sure the user creating the subscription has the same login name and password at the primary data server, select permission on the primary table, and execute permission for rs_marker. Investigate whether the replicate Replication Server is out of stable queue segments. Verify that the primary Replication Server is running and that the SQM, SQT, and DIST threads for the primary database are running. Investigate whether the primary Replication Server is out of segments.
Defined	Active	The replicate Replication Server is building the materialization queue. The primary Replication Server has received the activation request and sent it to the replicate Replication Server.	Verify that the materialization queue is being built. Verify that the connection between the primary and replicate Replication Server is up. Verify that the DSI thread for the replicate database is running.
Qcomplete	Defined	Waiting for the activation request to be processed by the distributor at the primary Replication Server.	Investigate whether the primary Replication Server is out of stable queue segments. Verify that the primary Replication Server is up and the SQM, SQT, and DIST threads for the primary database are running.

Replicate Status	Primary Status	Subscription State	Suggested Actions
Qcomplete/ Recovering	Defined	The replicate Replication Server was recycled. Waiting for the activation request to be processed by the distributor at the primary Replication Server.	The state at the replicate Replication Server should soon change to Qcomplete.
Qcomplete	Active	The primary Replication Server received and processed the activation request. Now, this Replication Server is waiting for the validation request. The replicate Replication Server is waiting for the activation request.	Check the route from the primary Replication Server to the replicate Replication Server. Verify that the DSI thread for the replicate database is running. Investigate whether the replicate Replication Server has run out of queue segments.
Qcomplete	Valid	The subscription is done at primary Replication Server. The replicate Replication Server is waiting for the activation request.	Check the route between the primary Replication Server and the replicate Replication Server. Check the DSI to the replicate database. Investigate whether the replicate Replication Server has run out of queue segments.
Active and not Qcomplete	Active	The replicate Replication Server is applying and building the materialization queue at same time. The primary Replication Server is waiting for the validation request. This request is not sent until the queue is complete at the replicate Replication Server.	Verify that the DSI thread is processing the materialization queue. Investigate whether the replicate Replication Server is out of stable queue segments. You can monitor the materialization queue activity using <code>admin who</code> , <code>sqm</code> .

Replicate Status	Primary Status	Subscription State	Suggested Actions
Qcomplete and Active	Active	<p>The primary Replication Server is waiting for the validation request.</p> <p>The replicate Replication Server is applying the materialization queue.</p>	<p>Verify that the DSI is processing the materialization queue.</p> <p>You can monitor the materialization queue activity using <code>admin who</code>, <code>sqm</code>.</p> <p>Verify that the primary data server, the RepAgent for the primary database and the RepAgent User thread in the primary Replication Server are running.</p> <p>Verify that the primary Replication Server is running and that the SQM, SQT, and DIST threads for the primary database are running.</p> <p>Investigate whether the primary Replication Server is out of segments.</p>
Qcomplete and Active	Valid	<p>The primary Replication Server received the validation request.</p> <p>The replicate Replication Server is applying the materialization queue.</p>	<p>Verify that the DSI thread is processing the materialization queue.</p> <p>You can monitor the materialization queue activity using <code>admin who</code>, <code>sqm</code>.</p> <p>Check the route between the primary Replication Server and the replicate Replication Server.</p>
Materialized	Active	<p>The primary Replication Server is waiting for the validation request.</p> <p>The replicate Replication Server has finished applying the materialization queue.</p>	<p>Investigate whether the primary Replication Server is out of segments.</p> <p>Verify that the primary data server, the RepAgent for the primary database, and the RepAgent User thread at the primary Replication Server are running.</p> <p>Verify that the SQM, SQT, and DIST threads for the primary database are running.</p>

Replicate Status	Primary Status	Subscription State	Suggested Actions
Materialized	Valid	The primary Replication Server received the validation request. The replicate Replication Server has finished applying the materialization queue.	Check the route between the primary Replication Server and the replicate Replication Server. Verify that the DSI for the replicate database is running. Investigate whether the replicate Replication Server has run out of stable queue segments.
Valid	Valid	Complete.	None.

Bulk materialization

To use bulk materialization, execute:

- define subscription to add the subscription to the RSSD for the primary and replicate Replication Servers,
- activate subscription to add the subscription to the Subscription Resolution Engine (SRE), and
- validate subscription to verify that the data at the replicate data is consistent with the primary data.

Defining the subscription

Table 5-5 describes solutions for bulk materialization problems that occur during the definition phase based on the status returned by check subscription. The Replicate Status column shows the subscription status returned by check subscription at the replicate Replication Server. The Primary Status column shows the subscription status returned by check subscription at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Table 5-5: Bulk materialization problems—definition phase

Replicate Status	Primary Status	Subscription State	Suggested Actions
Defined/ Pending	Invalid	Waiting for other subscriptions for the same replication definition and replicate database.	Look for other subscriptions for the same replication definition and database. If there are no other subscriptions, wait for five minutes.
Defined/ Recovering	Invalid	Cannot connect to the primary Replication Server to define the subscription.	Check the replicate Replication Server error log for messages. Make sure the user creating the subscription has the same login name and password at the primary Replication Server and the replicate Replication Server.
Defined	Defined	Definition cycle is complete.	Activate the subscription.

Activating the subscription

After the subscription is defined, use `activate subscription` to set the subscription status to `Activating` at the replicate Replication Server and to send the activation request to the primary Replication Server.

When the primary Replication Server receives the activation request, it changes the subscription status to `Active`, sends the activation request back to the replicate Replication Server, and begins sending updates for the subscription to the replicate Replication Server. When the replicate Replication Server receives the activation request, it changes the subscription status to `Active` at the replicate Replication Server.

If the `with suspension` option was specified for the `activate subscription` command, the DSI thread is suspended when the activation request is received at the replicate Replication Server. At this stage, you load the replicate database. For more information, see the *Replication Server Reference Manual*. After loading, execute the `resume connection` command to continue applying transactions from the DSI.

Table 5-6 describes solutions for bulk materialization problems that occur during the activation phase based on the status returned by check subscription. The Replicate Status column shows the subscription status returned by check subscription at the replicate Replication Server. The Primary Status column shows the subscription status returned by check subscription at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Table 5-6: Bulk materialization problems—activation phase

Replicate Status	Primary Status	Subscription State	Suggested Actions
Activating/ Recovering	Defined	Cannot send the activation request to the primary Replication Server.	Verify that the primary Replication Server is running and the SQM, SQT, and DIST threads for the primary database are running. Check the error log in the replicate Replication Server. Investigate whether the primary Replication Server has run out of segments.
Activating	Active	The primary Replication Server received the activation request and returned it to the replicate Replication Server. The replicate Replication Server is waiting for the activation request.	Verify the connection between the primary Replication Server and the replicate Replication Server. Verify that the DSI thread for the replicate database is running.
Active	Active	The activation stage is complete.	Validate the subscription. If you specified the with suspension option for the activate subscription command, you may now load the replicate database.

Validating the subscription

After the subscription is activated, complete bulk materialization by using validate subscription at the replicate Replication Server. The subscription status is changed to Validating. The replicate Replication Server forwards the validation request to the primary Replication Server.

When the primary Replication Server receives the validation request, it changes the status at the primary to Valid and returns the validation request to the replicate Replication Server. When the replicate Replication Server receives the validation request, the subscription status becomes Valid at both sites, and the bulk materialization is complete.

Table 5-7 describes solutions for bulk materialization problems that occur during the validation phase based on the status returned by check subscription. The Replicate Status column shows the subscription status returned by check subscription at the replicate Replication Server. The Primary Status column shows the subscription status returned by check subscription at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Table 5-7: Bulk materialization problems—validation phase

Replicate Status	Primary Status	Subscription State	Suggested Actions
Validating/ Recovering	Active	Cannot send the validation request to the primary Replication Server.	Verify that the primary Replication Server is running.
Validating	Valid	The primary Replication Server received the validation request. The replicate Replication Server is waiting for the validation request.	Verify the route between the primary and replicate Replication Server. Verify the DSI is running. Verify that the replicate Replication Server has sufficient queue segments.
Valid	Valid	Complete.	None.

Dematerialization problems

When you execute drop subscription at the replicate Replication Server, the subscription status becomes Dematerializing in the replicate Replication Server `rs_subscriptions` table. If there are any other materializing or dematerializing subscriptions for the same replication definition and replicate database, the recovering column in the `rs_subscriptions` table is set to 2 to indicate that the dematerialization request is pending. The check subscription command indicates that the subscription has a Dematerializing or Pending status at the replicate Replication Server.

Subscription creation and drop requests are processed one at a time in the order entered. The `request_date` column in `rs_subscriptions` stores the time a request was entered.

There are two methods for dematerializing subscriptions:

- `with purge` – deletes the subscription data from the replicate database.
- `without purge` or `bulk` – does not delete the subscription data from the replicate database.

Both methods go through two stages:

- `Dematerialization` – stops sending updates for the subscription to the replicate database and, optionally, deletes the subscription data from the replicate database.
- `Removal` – deletes the subscription from the system tables of the primary and the replicate Replication Server.

with purge dematerialization

The `dematerialize with purge` method is invoked by using the `with purge` option of the `drop subscription` command. The subscription status is set to `Dematerializing` at the replicate Replication Server and a drop request is forwarded to the primary Replication Server.

When the primary Replication Server gets the drop request, it stops sending updates for the subscription to the replicate Replication Server. The subscription status is changed to `Dematerializing` at the primary Replication Server and a drop request is returned to the replicate Replication Server.

When the replicate Replication Server gets the drop request, it creates a dematerialization queue and starts a dematerialization thread to fill the queue. The dematerialization thread selects subscription data from the replicate database and puts a row delete operation into the dematerialization queue for each row that does not belong to any other active subscription.

While the dematerialization thread is writing to the dematerialization queue, the DSI thread switches to the dematerialization queue and begins applying the deletes to the replicate database.

If the dematerialization thread encounters an error, the entire dematerialization process restarts and the same error may recur. To avoid this problem, suspend the connection to the replicate database. Then resume it using the `skip transaction` option. This resumes the connection with the second transaction in the queue. The first transaction is written to the exceptions log.

When dematerialization has completed, the subscription status at the replicate Replication Server is changed to Removing. The replicate Replication Server logs in to the primary Replication Server and asks it to delete the subscription from its system tables. When that request succeeds, the replicate Replication Server removes the subscription from its own system tables and the dematerialization process is complete.

If the replicate Replication Server runs out of stable queue segments while dematerializing using the `with purge` option:

- 1 Add a new partition to the replicate Replication Server. If no partition is available, consider using bulk dematerialization.
- 2 If you have difficulty deleting subscription data, drop the subscription using the `without purge` option.
- 3 If you have difficulties because of permissions, a user with the appropriate permission should issue the `drop subscription` command.

Table 5-8 describes solutions for dematerialization problems based on the status returned by `check subscription`. The Replicate Status column shows the subscription status returned by `check subscription` at the replicate Replication Server. The Primary Status column shows the subscription status returned by `check subscription` at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Regardless of the problem you are having, you should first check the error logs for all of the servers involved for specific information. Also check that the user who is dropping the subscription has the same login name and password at the replicate and primary Replication Servers, and that the maintenance user for the replicate database has `select`, `delete`, and `update` permissions on the replicate table.

Table 5-8: Dematerialization problems—with purge option

Replicate Status	Primary Status	Subscription State	Suggested Actions
Dematerializing/ Pending		Waiting for other subscriptions for the same replication definition and replicate database.	Look for other subscriptions being created or dropped for the same replication definition and database. If no other subscription operations are in process, wait for five minutes.

Replicate Status	Primary Status	Subscription State	Suggested Actions
Dematerializing/ Recovering		Cannot connect to the primary Replication Server to drop the subscription.	<p>Check replicate Replication Server error log for messages.</p> <p>If the user dropping the subscription does not have the same login and password at the primary and replicate Replication Servers, then:</p> <ul style="list-style-type: none"> • Provide the user with a Replication Server account that has the same login and password at the primary and replicate, or • Have a different user with the appropriate permissions drop the subscription. <p>Also, the user must have at least primary subscribe privileges at the primary Replication Server.</p>
Dematerializing		The primary Replication Server is waiting for the drop request.	<p>Investigate whether the primary Replication Server is out of queue segments.</p> <p>Verify that the primary Replication Server is running and that the SQM, SQT and DIST threads for the primary database are running.</p>
Dematerializing	Dematerializing	<p>The primary Replication Server processed the drop request and sent it to the replicate Replication Server.</p> <p>The replicate Replication Server is waiting for the drop request.</p>	<p>Check the route between the primary Replication Server and the replicate Replication Server.</p> <p>Check the DSI thread for the replicate database.</p> <p>Investigate whether the replicate Replication Server has run out of queue segments.</p>
Dematerializing/ Recovering	Dematerializing	<p>The primary Replication Server processed the drop request and sent it to the replicate Replication Server.</p> <p>The replicate Replication Server failed to create and process the dematerialization queue.</p>	<p>Investigate whether the replicate Replication Server has run out of queue segments.</p> <p>Verify that the maintenance user has select privileges on the replicate table.</p> <p>Check the DSI thread for the replicate database.</p>

Replicate Status	Primary Status	Subscription State	Suggested Actions
Removing/ Recovering	Dematerializing	The subscription has dematerialized at the replicate Replication Server, but the replicate Replication Server could not log in to primary Replication Server to remove the subscription from the system tables.	Verify that the primary Replication Server is running. If the user dropping the subscription does not have the same login and password at the primary and replicate Replication Servers, then: <ul style="list-style-type: none"> • Provide the user with a Replication Server account that has the same login and password at the primary and replicate, or • Have a different user with the appropriate permissions drop the subscription.
Removing	Dematerializing	The primary Replication Server is deleting the subscription. The replicate Replication Server is waiting for it to finish.	Wait.
Removing	Invalid	The subscription has been removed from primary Replication Server. The replicate Replication Server will remove the subscription next.	Wait.
Invalid	Invalid	The subscription has been dropped.	None.

Bulk dematerialization

Bulk dematerialization is invoked using the `without purge` option of the `drop subscription` command. The subscription status becomes `Dematerializing` at the replicate Replication Server and a drop request is forwarded to the primary Replication Server.

When the primary Replication Server receives the drop request, it stops sending updates for the subscription to the replicate Replication Server. The subscription status becomes `Dematerializing` at the primary Replication Server and a drop request is returned to the replicate Replication Server.

When the replicate Replication Server receives the drop request, the subscription status is changed to Removing at the replicate. The replicate Replication Server logs in to the primary Replication Server and requests that it delete the subscription from its system tables. When that request has succeeded, the replicate Replication Server removes the subscription from its own system tables and the dematerialization is complete.

Table 5-9 describes solutions for bulk materialization problems based on the status returned by check subscription. The Replicate Status column shows the subscription status returned by check subscription at the replicate Replication Server. The Primary Status column shows the subscription status returned by check subscription at the primary Replication Server. The Subscription State column provides a detailed description of the subscription status. The Suggested Actions column provides the solutions.

Table 5-9: Dematerialization problems—without purge option

Replicate Status	Primary Status	Subscription State	Suggested Actions
Dematerializing/ Pending		Waiting for other subscription requests for the same replication definition and replicate database.	Check for other subscription being created or dropped for the same replication definition and database. If there are no other subscriptions, wait for five minutes.
Dematerializing/ Recovering		Cannot connect to the primary Replication Server to drop the subscription.	Check the replicate Replication Server error log for messages. Make sure the user who created the subscription has the same login name and password at the primary Replication Server and the replicate Replication Server. The user should also have at least primary subscribe privileges.
Dematerializing		The primary Replication Server is waiting for the drop request.	Determine whether the primary Replication Server has run out of queue segments. Verify that the primary Replication Server is up and that the SQM, SQT, and DIST threads for the primary database are running.

Replicate Status	Primary Status	Subscription State	Suggested Actions
Dematerializing	Dematerializing	The primary Replication Server processed the drop request and sent it to the replicate Replication Server. The replicate Replication Server is waiting for the drop request.	Check the route between the primary Replication Server and the replicate Replication Server. Check the DSI thread for the replicate database. Determine whether the replicate Replication Server ran out of queue segments.
Dematerializing/ Recovering	Dematerializing	The primary Replication Server processed the drop request and returned it to the replicate Replication Server. The replicate Replication Server abnormally terminated.	Wait for the subscription daemon to reset the recovering flag.
Removing/ Recovering	Dematerializing	The replicate Replication Server could not log in to the primary Replication Server to delete the subscription from the system tables.	Verify that the primary Replication Server is running.
Removing	Dematerializing	The primary Replication Server is deleting the subscription from the system tables. The replicate Replication Server is waiting for the primary Replication Server to finish.	Wait.
Removing	Invalid	The subscription is removed from the primary Replication Server. The replicate Replication Server will remove the subscription next.	Wait.
Invalid	Invalid	The subscription has been dropped.	None.

Replication Server Interface Problems

This chapter describes how to fix Replication Server interface problems.

Name	Page
Overview	103
Connection errors	104

If you cannot find the solution in this chapter

If an error message is displayed in an error log and you cannot find it in this chapter, refer to:

- “Analyzing error logs” on page 10, or
- Chapter 3, “Common Errors”

If an error message is not displayed, refer to Chapter 1, “Identifying the Problem”

Overview

The RSI allows two Replication Servers to exchange transactions across a route. The RSI consists of:

- A stable queue and thread at the primary Replication Server. The thread reads the stable queue and writes to the network.
- A thread at the replicate Replication Server. This thread reads from the network and writes into outbound stable queues.

There are two basic types of RSI problems:

- Those occurring when Replication Server attempts to connect to another Replication Server, and
- Those that occur when the locator is invalid.

Although many errors may occur when a primary Replication Server loses its connection to the replicate Replication Server, none of these errors should cause a loss of data or cause the RSI to be suspended. The primary thread continuously attempts to connect to the replicate Replication Server until the connection is reestablished.

Losing the connection terminates the RSI thread at the replicate Replication Server. The thread at a replicate Replication Server exists only while a connection exists. Executing `admin who` in the replicate Replication Server shows that the primary Replication Server is no longer logged in as an RSI source. When the primary Replication Server reestablishes the connection to the replicate Replication Server, data is replicated across the route to the replicate Replication Server.

An RSI locator identifies the last message that the replicate Replication Server stable queue received from the primary Replication Server. If the RSI locators at the primary and replicate Replication Servers are not identical, then the primary Replication Server might delete messages from its stable queue at an incorrect location.

Connection errors

These errors can occur when two Replication Servers attempt to connect through the RSI:

- Incorrect RSI user login name or password
- Incorrect user permissions at replicate Replication Server

Incorrect RSI user login name or password

The primary Replication Server logs in to the replicate Replication Server using the RSI login name and password. The RSI user login name and password are created by `rs_init` and are also required when you create a route to the replicate Replication Server.

If there is a problem with the login name and password at the replicate Replication Server, the primary Replication Server logs one of these errors:

```
Jan 20 09:44:45.656 1993: Error: 1337, Line: 1978, File:  
'/generic/useful/cm.c'- When trying to connect to
```

```
server 'Rep_server' as user 'Rep_server_RSI', the DB-
Library routine 'dbopen' returned the error 'Error:
20014, Severity: 2 -- 'Login incorrect.''. The server
error message is 'Message: 1421, State: 0, Severity: 11
-- 'Invalid login attempted by user 'Rep_server_RSI'''.
```

To repair this problem:

- 1 Retrieve the RSI user login name and password from the RSSD of the primary Replication Server by using isql to execute this query on the RSSD:

```
select username, password
from rs_maintusers, rs_sites
where name = remote_RS_name and destid = id
```

remote_RS_name is the destination Replication Server's name.

id is the destination Replication Server's site ID.

If you have password encryption, you cannot access the password by executing a query.

- 2 Perform one of these actions:
 - Use the alter route command to change the password for the primary Replication Server to the RSI user password, which you retrieved in the previous step.
 - Use the alter user command to change the password for the login name at the replicate Replication Server. Then use the resume route command to resume the route.

Incorrect user permissions at replicate Replication Server

The login must also have connect source or sa permission on the replicate Replication Server. If the login lacks the appropriate permission, an error message is sent to the primary Replication Server's error log.

To repair this problem:

- 1 Grant connect source permission to the login at the replicate Replication Server.
- 2 Resume the route at the primary Replication Server

Invalid RSI locator

An RSI locator points to the last message that the replicate Replication Server stable queue received from the primary Replication Server. The RSI uses the RSI locator in the `rs_locator` system table to delete messages in the stable queues. The replicate Replication Server updates the RSI locator in the `rs_locator` system table with the last messages that the replicate Replication Server added to its stable queues. Periodically, the primary Replication Server requests the RSI locator from `rs_locator` at the replicate Replication Server. The primary Replication Server deletes its stable queue messages up to and including the message identified by the RSI locator—these are the messages that the replicate Replication Server has already received.

Note The RSI locator is different from the locator that is used with database log truncation. See “Resetting the locator” on page 132 for more information about the database log truncation locator.

If the RSI locators at the primary and replicate Replication Servers are not identical, then the primary Replication Server might delete messages from its stable queue at an incorrect location and invalidate the RSI locator at the replicate Replication Server.

RSI locators must be identical. If they are not identical:

- 1 Make sure that the replicate Replication Server is not running.
- 2 Set the replicate Replication Server RSI locator to “0” by executing this command against the replicate Replication Server’s RSSD:

```
update rs_locator set locator=0x0
  where sender = primary_replication_server
```

RepAgent and Log Transfer Manager Problems

This chapter describes how to troubleshoot RepAgent and Log Transfer Manager (LTM) problems.

Name	Page
Overview	107
Problems when starting the RepAgent or LTM	108
Problems restarting the LTM after it disconnects	110
Errors from the Replication Server	110
Errors from the Adaptive Server	114

If you cannot find the solution in this chapter

If an error message is displayed in an error log and you cannot find it in this chapter, refer to:

- “Analyzing error logs” on page 10, or
- Chapter 3, “Common Errors”

If an error message is not displayed, refer to Chapter 1, “Identifying the Problem”

Overview

The RepAgent and LTM retrieve the data for primary objects from the Adaptive Server log and convert the log record information into Log Transfer Language (LTL) commands. These commands are sent to the primary Replication Server for distribution and replication. The RepAgent and LTM also coordinate database log truncation with the Adaptive Server and primary Replication Server.

Errors can be caused by problems with the Adaptive Server, Replication Server, RepAgent, or LTM. The RepAgent records error messages in the Adaptive Server error log. The LTM records LTM, Replication Server, and Adaptive Server error messages in its error log. These error messages identify the server that caused the error so that you can diagnose and correct the problem.

This chapter describes these problems:

- Incorrect logins and passwords
- Incorrect permissions
- Phantom connections
- Inconsistent replication definition setup
- Conflicting text or image column status
- Invalid truncation page
- Not enough Adaptive Server alarm resources

Note For more information about how the RepAgent and LTM processes errors, see the *Replication Server Administration Guide*.

Problems when starting the RepAgent or LTM

These errors can occur when the RepAgent or LTM starts up:

- Invalid login to the Replication Server or Adaptive Server
- Invalid permissions for the Replication Server

Invalid login

The RepAgent requires only a login name and password for the Replication Server user. The user must have connect source permission in the Replication Server. An LTM uses a configuration file that contains the login names and passwords for logging in to the Adaptive Server and Replication Server. To scan the Adaptive Server log you must specify the System Administrator or Database Owner login name and password. The login name and password the LTM uses to log in to the Replication Server must be defined to the Replication Server with the create user command.

When the RepAgent fails to log in to the Replication Server, RepAgent retries the login and error 9216 is displayed in the Adaptive Server log. If RepAgent continues to fail logging in to the Replication Server, RepAgent error 9214 is displayed in the Adaptive Server error log.

When the LTM fails to log in to the Replication Server, Replication Server error 1421 is displayed in the LTM error log:

```
Message: 1421, State: 0, Severity: 11 -- 'Invalid login
attempted by user 'john'
Jan 25 18:04:11.613 1997: Connection to server MY_PRS
failed.
```

When the LTM fails to log in to the Adaptive Server, Replication Server error 4002 is displayed in the LTM error log:

```
Message: 4002, State: 1, Severity: 14 -- 'Login
failed.'
Jan 25 18:05:54.034 1997: Connection to server MY_PDS
failed.
```

Invalid permissions

The RepAgent and LTM user must have connect source permission for the Replication Server; otherwise, RepAgent error 9211 is displayed in the Adaptive Server error log:

```
98/03/04 14:01:53.34 RepAgent (10) Error: 9211,
Severity: EX_CMDFATAL, State: 1
Received the following error message from the
Replication Server: Error: 37024, Line: 0, File: '' -
Connect source permission is required to execute
command.''
Replication Server error 37024 is displayed in the text
of error 9211 (for RepAgent) or the LTM error log:
```

Problems restarting the LTM after it disconnects

When the LTM disconnects, the Adaptive Server or Replication Server may not have terminated the LTM correctly. The Adaptive Server or Replication Server may still have an invalid LTM connection, called a “phantom” connection. (A phantom connection is usually a generic network problem.)

When the LTM attempts to reconnect to an Adaptive Server with a phantom connection, Adaptive Server error 9121 is displayed in the LTM error log:

```
Message: 9121, State: 1, Severity: 16 -- 'The log
transfer context for the current database is already
reserved by Adaptive Server process 19. The log transfer
context for the current context is not reserved.'
```

Error message 9121 is a retryable error. The LTM attempts to connect to the Adaptive Server in one minute intervals. If the LTM fails to connect ten times, it shuts down. You must kill the phantom process in the Adaptive Server.

For more information about this problem and how to kill a process, see “9220: connect only one RepAgent to a Replication Server (LTM 1078)” on page 60.

When the LTM attempts to reconnect to a Replication Server with a phantom connection, Replication Server error 14027 is displayed in the LTM error log:

```
Message: 14027, State: 0, Severity: 16 -- 'LTM for
'MY_PDS.MY_PDB' is already connected.'
```

Error 14027 is a retryable error, and the LTM attempts to reconnect.

Errors from the Replication Server

The RepAgent or LTM treats most Replication Server normalization errors, which have numbers between 32000 and 32999, as recoverable. The RepAgent or LTM logs the error and continues processing.

All other Replication Server errors, except network failures, are treated as fatal by the RepAgent or LTM. For example, an incorrect Replication Server login name and password in the configuration file is a fatal error. The LTM disconnects from the Adaptive Server and Replication Server and terminates. After you have corrected the problem, restart the RepAgent or LTM.

Normalization errors

Normalization errors result from transient inconsistencies in the setup of replication objects. For example, if a table has been marked replicated with the `sp_setreplicate` system procedure, but the replication definition has not yet been created for the table, the RepAgent or LTM retrieves log records for an object that is not yet known to the Replication Server. Replication Server error 32032 is displayed in the Replication Server error log:

```
Message: 32032, 'No table with name 'foo' is defined for
database with id MY_PDS.MY_PDB1.
```

Correct this problem by creating the replication definition for the table.

A normalization error can also occur if the replicated stored procedure is declared in the Replication Server for function replication but is marked in the Adaptive Server for table replication. Replication Server error 32047 is displayed in the Replication Server error log:

```
Message from server: Message: 32047, State 0, Severity
12 -- 'Function 't2' is associated with a function
replication definition, but an invoking stored
procedure was marked for a table replication
definition.'
```

See the *Replication Server Administration Guide* for instructions on setting up replication definitions.

Treating normalization errors as recoverable

When you treat normalization errors as recoverable, errors may still occur even when a replication definition is created and a subscription is successfully materialized. For example, an update may be lost, causing the replicate data to be inconsistent with the primary data. This kind of error should occur only during the initial setup of a replication system. The system administrator should monitor the error logs closely and correct errors as they are identified. To reconcile subscriptions, refer to Chapter 5, “Subscription Problems”

Normalization errors may also occur with replicated stored procedures.

Replication status inconsistency of *text* or *image* columns

Another normalization error occurs when there are inconsistencies in the replication status of text or image columns between the Adaptive Server database and the replication definition. Replication Server error 32046 is displayed in the Adaptive Server error log (for RepAgent) or the LTM error log:

```
Message: 32046 -- 'The status of column '%s' in repdef
is inconsistent with that of the LTL command.'
```

The replication status for text and image columns in the Adaptive Server database is carried in the modification LTL commands that the RepAgent or LTM sends to the Replication Server. If the status of a text or image column is not the same in the Adaptive Server database and in the replication definition, then Replication Server detects the inconsistency when the modification is being replicated, and the RepAgent or LTM shuts down.

If a text or image column has a status of `do_not_replicate` at the Adaptive Server database and the replication definition includes that column for replication, processing continues and the Replication Server sends the modifications to the replicate database without the text or image data. The Replication Server records a warning message in the Adaptive Server error log (for the RepAgent) or LTM error log.

Use the following procedures to resolve replication status inconsistencies for text or image column status and to resume operations after the RepAgent or LTM has shut down.

When *replicate_if_changed* conflicts with *always_replicate*

When the RepAgent or LTM shuts down because a text or image column has a status of `replicate_if_changed` at the Adaptive Server database and `always_replicate` in the replication definition, you must change the replication status so that they match, for example, either:

- Replicate text or image columns only when their values change, or
- Always replicate text or image columns.

To replicate text or image columns only when their values change:

- 1 Execute the alter replication definition command at the primary Replication Server and change the status of the text or image columns to `replicate_if_changed`. Wait for the modified replication definition to arrive at the replicate sites.
- 2 Restart the RepAgent or LTM.

To always replicate text or image columns:

- 1 Stop updates at the primary table.
- 2 Execute the alter replication definition command at the primary Replication Server, and change the status of the text or image columns to `replicate_if_changed`. Wait for the modified replication definition to arrive at the replicate sites.
- 3 Restart the RepAgent or LTM to let transactions with a `replicate_if_changed` status finish processing.
- 4 Execute the `sp_setrepcol` system procedure at the Adaptive Server and change the status to `always_replicate`.
- 5 Execute alter replication definition at the primary Replication Server and change the status of the text or image columns to `always_replicate`. Wait for the modified replication definition to be replicated to the replicate sites.
- 6 Resume updates to the primary table.

When `do_not_replicate` columns are included in the replication definition

When the Replication Server reports that the status of a text or image column is `do_not_replicate` at the Adaptive Server database and the replication definition includes that column for replication, you must change the replication status to either:

- Replicate text or image columns, or
- Not replicate text or image columns.

If you want to replicate text or image columns:

- 1 Execute the `sp_setrepcol` system procedure at the Adaptive Server database and change the status of the text or image column to `always_replicate` or `replicate_if_changed`. It should match the status in the replication definition.
- 2 Wait for subsequent transactions that modify the text or image column to be processed by the Replication Server.
- 3 Consider correcting any inconsistencies with the `rs_subcmp` program.

If you do not want to replicate text or image columns:

- 1 Stop updates to the primary table.
- 2 Drop subscriptions to the replication definition.
- 3 Drop the replication definition.

- 4 Re-create the replication definition without the text or image columns, and re-create subscriptions.
- 5 Resume updates to the primary table.

Errors from the Adaptive Server

This section describes errors from Adaptive Server that are caused by:

- An invalid truncation page, or
- Running out of Adaptive Server alarm resources.

The RepAgent does not use the Adaptive Server log scan thread; it will not encounter any errors related to the Adaptive Server log scan thread.

For an LTM, errors from Adaptive Server are usually related to the Adaptive Server log scan thread reading the Adaptive Server database log and have error numbers between 9100 and 9199. These errors are not recorded in the Adaptive Server log because they are not considered severe to the Adaptive Server. They are sent only to the LTM client. The LTM logs the errors and performs an action based on the error. Analyze the LTM error log for errors or conditions that occur when the log is being read.

Dropping primary objects

Dropping primary tables or stored procedures from an Adaptive Server affects data replication and should be done with caution.

If updates are made to the primary table and the table is dropped from the Adaptive Server, Adaptive Server error 9104 is displayed in the Adaptive Server error log (for the RepAgent) or LTM error log:

```
Message: 9104 'Cannot identify the object on the INSERT
log record for database 'MY_PDB1', XACT ID Rid pageid =
0xa0f; row num = 0x5, RID Rid pageid = 0xa0f; row num =
0x6. Information associated with the INSERT log record
is not replicated.
```

The log record for a replicated data operation references another log page, called the oampage, that has information about the replicated object. Since the table was dropped, the oampage is deallocated and the data cannot be replicated. If the oampage is assigned to a newly created object, the Adaptive Server may associate log records for the dropped object with the new object.

In the following example, the oampage for foo may be allocated to foo_bar after foo is dropped:

```
update table foo
drop table foo
create table foo_bar
update foo_bar
```

In the preceding example, Adaptive Server sends the first update record as an update to foo_bar instead of foo. If the new table has a much larger row size, then an Adaptive Server segmentation fault may occur.

To prevent this error from occurring, wait until the RepAgent or LTM has processed all updates for a primary table before dropping it. See the *Replication Server Administration Guide* for instructions to drop tables.

Except for replication definitions, schema changes to primary objects may also produce incorrect or unexpected log scan results. Make these changes with caution.

Dropping or recreating replicated stored procedures leads to unpredictable results; schema alterations have the same effect. Make these type of changes *only* after the RepAgent or LTM has processed all the records for the stored procedure.

Errors may also occur at the Replication Server, check the Replication Server for error messages.

Invalid truncation page

The RepAgent or LTM executes the dbcc settrunc command on the Adaptive Server to set the secondary truncation point in the database log. The secondary truncation point identifies the page that contains the begin transaction statement of the last transaction that was sent to the Replication Server. If the RepAgent or LTM specifies an invalid page to the Adaptive Server, an error is recorded and the RepAgent or LTM aborts. You may see Adaptive Server error 4213:

```
Message: 4213, State: 1, Severity: 16 -- 'Page 2561 in
database 'MY_DB' cannot be used as a truncation page
because the page is not allocated.'
```

The `rs_locator` table in the Replication Server system database tells the RepAgent or LTM where to set the secondary truncation point in the Adaptive Server log. This error is expected if the truncation state is set to ignore and the log is truncated. The RepAgent or LTM then restarts without resetting `rs_locator`.

To solve this problem, set `rs_locator` to zero and restart the RepAgent or LTM. For more information about log truncation, refer to “Resetting the locator” on page 132.

Running out of Adaptive Server alarm resources

The log scan code uses Adaptive Server alarm resources to wake up the log scan thread when new log records are appended to the log. The current default number of alarms configured for the Adaptive Server is 40. This should be enough when there are only a few LTMs running against the Adaptive Server.

If the Adaptive Server runs out of alarms during the log scan, Adaptive Server error 9123 is displayed in the LTM error log:

```
Message: 9123, State: 1, Severity: 16 -- 'Unable to set
a wakeup alarm for the logtransfer process.The
logtransfer scan of database 'MY_PDB1' will not honor
the specified timeout of 10 seconds.'
```

The LTM treats this as a recoverable error, because the error should disappear when Adaptive Server alarms become available. However, if there are many LTMs running against databases on the Adaptive Server, the system administrator should increase the configured number of alarms by executing the following command:

```
buildmaster -ycnalarm=200 -ddbmaster
```

This example sets the number of alarms on the Adaptive Server master device `ddbmaster` to 200.

This chapter describes problems you may encounter in the DSI.

Name	Page
Overview	117
Listing databases and using admin who and admin who, dsi	118
Troubleshooting procedure	120
DSI down or suspended	121
DSI active	126

If you cannot find the solution in this chapter

If an error message is displayed in an error log and you cannot find it in this chapter, refer to:

- “Analyzing error logs” on page 10, or
- Chapter 3, “Common Errors”

If an error message is not displayed, refer to Chapter 1, “Identifying the Problem”

Overview

The DSI applies transactions from a Replication Server stable queue to a data server, which can be an Adaptive Server or another data server with an interface to the Replication Server. Before a transaction is applied, Replication Server uses function strings to convert the commands to the format specified for the data server.

If the data server is not receiving updates, the DSI may be suspended or down. A down or suspended DSI can be caused by:

- Incorrect login or permissions
- A data server error
- A Replication Server error

A suspended DSI always results in an error message being sent to the Replication Server error log. You can analyze this error message to solve the problem. After correcting the error, restart the DSI.

If the DSI is active, unique transactions might be incorrectly resolved as duplicates, or the transactions might be failing on the data server. You must manually resolve the data inconsistency between the primary and replicate databases or reapply the transactions.

To determine the exact cause of the problem, use a select command, `admin who`, and `admin who, dsi`.

Note These problems also apply to parallel DSI threads. For more information specific to parallel DSI threads, see “Using Parallel DSI Threads” in Chapter 16, “Performance Tuning”, in the *Replication Server Administration Guide*.

Listing databases and using *admin who* and *admin who, dsi*

To investigate the DSI, find out which databases are controlled by the Replication Server and check the DSI status for all database connections on a Replication Server. Use this information to determine the cause of the problem.

Listing databases controlled by a Replication Server

To list the databases controlled by a Replication Server, log in to the system Adaptive Server, change to the RSSD, and enter:

```
select dsname, dbname from rs_databases, rs_sites
where prsid = id and name = <replication_server>
```


Checking the DSI status for a database connection

The command `admin who` displays the states of all threads in the Replication Server, including DSI scheduler and executor threads. In the output of `admin who`, the DSI scheduler threads have the name “DSI” and DSI executor threads have the name “DSI EXEC.” If the DSI is suspended when Replication Server starts up, the output shows only one DSI executor thread, even if additional DSI executor threads are configured.

The command `admin who, dsi` shows the states of all running DSI scheduler threads and provides configuration values and other information about them. If a DSI scheduler thread exists for a database but does not appear in the output of `admin who, dsi`, use the `resume connection` command to restart the data server interface for the database. For a complete description of `admin who, dsi` output, see the *Replication Server Reference Manual*.

The states of DSI scheduler threads in the output of `admin who` and `admin who, dsi` are:

- **Active** – the thread is starting, restarting after an internal error, or logging an exception to the RSSD.
- **Awaiting Command** – the thread is waiting for a transaction to become available in the stable queue for the database. The queue contains no complete transaction at this moment.
- **Awaiting Wakeup** – the thread is sleeping before it restarts. It sleeps for two minutes after an error that can be retried. During the two-minute interval, the cause of the error may disappear or the System Administrator may correct the problem. If the problem is corrected, the thread can restart without error after the sleep interval.

A failed transaction is retried when it causes a data server error that you assigned the `retry_stop` or `retry_log` option using the `assign action` command.

- **Awaiting Message** – the thread has dispatched transactions to the DSI executor threads and is waiting for them to complete.
- **Suspended** – the DSI connection has been suspended by a user command, by an error, or by a `drop subscription` or `activate subscription` command with the `with suspension` clause. This state appears only in `admin who` output.
- **Down** – the thread has not been started. This state appears only in `admin who` output.

The states of DSI executor threads in the output of `admin who` are:

- Active – the thread is starting, executing a transaction at the data server, logging an exception to the RSSD, or restarting after an internal error.
- Awaiting Message – the thread is processing a transaction and is waiting for another thread to complete processing its transaction, or the thread is waiting to receive another statement from the Stable Queue Transaction interface (SQT).
- Awaiting Command – the thread is waiting to receive another transaction from the DSI scheduler thread.
- Down – the meaning of this state depends on the state of the DSI scheduler thread as follows:
 - If the DSI scheduler thread is Down, the DSI executor thread connection was suspended when the Replication Server was started, and the connection has not been resumed.
 - If the DSI scheduler thread is Active or Awaiting Wakeup, the DSI executor thread connection is recovering from a retryable error and is starting or restarting.
- Suspended – the connection has been suspended by a user command, an error, or a drop subscription or activate subscription command using the with suspension clause.

A failed transaction does not cause the DSI connection to be suspended and is retried if the failed transaction caused a data server error that you assigned the `retry_stop` or `retry_log` option using the `assign` action command. For more information about the `assign` action command, see Chapter 3, “Replication Server Commands,” of the *Replication Server Reference Manual*.

Troubleshooting procedure

If changes made to a primary database are not arriving at a replicate database that has active subscriptions, it is possible that the DSI for the replicate database is suspended.

Follow these steps to troubleshoot the DSI for the replicate database:

- 1 Log in to the Replication Server that controls the replicate database and execute the `admin who, dsi` command. This returns one entry for each database with replicated data. To learn how to use `admin who, dsi`, see “Listing databases and using `admin who` and `admin who, dsi`” on page 118.
- 2 If a database does not have an entry, check the Replication Server error log to see if the DSI for the database was not started or was suspended.

These are samples of error messages that are displayed in the Replication Server error log as a result of a DSI that did not start or was suspended:

```
The DSI thread for 'RDS.rdb2' is not started, because
the connection to the database has been suspended or
the connection has not been completed.
The DSI thread for database 'RDS.rdb2' is being
shutdown. DSI received data server error #1105 which
is mapped to STOP_REPLICATION. See logged data
server errors for more information. The data server
error was caused by RS output command #1 mapped from
input command #1 of the failed transaction.
The DSI thread for database 'RDS.rdb2' is shutdown.
```

If the DSI is down or suspended, see “DSI down or suspended” on page 121.

- 3 If there is an entry for the database, but the state is `Awaiting Command`, the DSI is waiting for a transaction from the primary. See Chapter 1, “Identifying the Problem” to determine why updates are not reaching this Replication Server.
- 4 If the state of the DSI thread is `Active`, then:
 - Unique transactions might be incorrectly resolved as duplicates, or
 - Transactions might be failing on the data server and written to the exceptions log.

Refer to “DSI active” on page 126 for specific instructions.

DSI down or suspended

This section describes solving these problems:

- Incorrect login or permissions

- A data server error
- A Replication Server error

If a subscription problem causes the DSI thread for the replicate database to terminate abnormally, you can restart the thread using the `resume connection` command. If possible, fix whatever condition caused the problem before resuming the connection. For example, if the maintenance user does not have update permission on the replicate table, first grant the user update permission and then resume the connection.

If you cannot fix the problem, resuming the connection will cause the DSI thread to reexecute the command that failed, and suspend the DSI again. To prevent this cycle, assign a different action to the error returned to the DSI. You must assign error actions at the Replication Server where the error class is created. For more information about error actions and classes, see the *Replication Server Administration Guide*.

Connection failure to the database

If the error is a connection failure to the database:

- Verify that the data server is defined in the *interfaces* file.
- Verify that the data server is running.
- Verify that the maintenance user name and password are correct.

To find the maintenance user name and password for a database, log in to the system Adaptive Server, use the RSSD, and execute the following query:

```
select username, password
       from rs_maintusers, rs_databases
       where destid = dbid
          and dsname = data_server
          and dbname = database
```

If the maintenance user name or password is not correct, you can either change them on the data server or use the `alter connection` command to change them in the Replication Server. If you have password encryption, you cannot access the password by executing a query. After correcting the user name or password, resume the connection.

Asynchronous transaction connection failures

When an asynchronous stored procedure reaches the primary Replication Server, the DSI uses the original login and password to connect to the primary data server. The maintenance login is *not* used.

If the login fails, use the `sysadmin log_first_tran` command to log the asynchronous transaction into the exceptions log. The `app_user` and `app_pwd` columns in the `rs_exceptshdr` system table contain the login and password the DSI uses to log in to the primary data server. See “Examining the exceptions log” on page 127 for more information.

Data server errors

Data server errors (for example, a permissions violation or duplicate key) are logged in the Replication Server error log.

If you cannot correct the data server error, restart the DSI by requesting that the first transaction (the one causing the error) be skipped. Do this by executing the resume connection command with the skip transaction option.

Warning! Skipping transactions may cause replication inconsistencies. See *Skipping transactions* for details.

If errors are caused by unique-key violations, you can use the `set autocorrection` command to turn automatic error correction on, before you resume the connection. These errors are likely to occur during nonatomic materialization. For more information about the `set autocorrection` command, see the *Replication Server Reference Manual*.

As a last resort, use the skip transaction option of the resume connection command. This option skips the materializing or dematerializing transaction altogether. If the problem occurred during subscription materialization, drop the subscription and create it again in a way that avoids the same problem.

For more information about subscription problems, see Chapter 5, “Subscription Problems”

Skipping transactions

A transaction intended for a database is skipped when:

- The resume connection command is executed with the skip transaction option.
- A data server error is encountered and the assigned action for the error is log, or retry_log.

Skipping a transaction that causes errors is a convenient way to force the DSI to resume applying transactions to the database. However, skipping a transaction has important implications for the database:

- Replicate data may become inconsistent with primary data.
- If an asynchronous transaction originates in a replicate database, and the transaction is skipped at the primary data server, the primary database is not updated.
- The inconsistency caused by skipping a transaction may result in even more errors if subsequent transactions depend upon the unapplied actions of the skipped transaction.
- Subscription materialization and dematerialization requests are special types of transactions. If you skip a subscription materialization request, it may invalidate the subscription and must be dropped.
- A subscription materialization or dematerialization request may cause Replication Server to create a separate stable queue, called the materialization or dematerialization queue. You cannot explicitly skip transactions in this queue using the resume connection command. They can, however, be skipped due to data server errors that are assigned the log or retry_log action.
- If a subscription materialization or dematerialization request is skipped, all transactions in the corresponding queue may be skipped. However, if a DSI is suspended in the middle of applying the transactions in a materialization or dematerialization queue, some of the transactions in the queue may already have been committed while others may not. Later, if you use the resume connection command to skip the request, the replicate database will have the effects of the previously committed transactions.

Because of these implications, you should skip transactions only after other means of correcting the error have been exhausted. After skipping a transaction, determine what you must do to bring the replicated data back to a state of consistency.

Customized handling of data server errors

You can customize how Replication Server handles data server errors. If you want the DSI to continue applying updates even when a data server error is encountered, you can change the error action from `stop_replication` to `log` or `retry_log`. To make this change, use the `assign action` command at the Replication Server where the error class was created. All databases using the error class are affected. If the primary database for a class is at a different site, then you may need to wait until the new error action is replicated to the local site.

To limit the change to one database, create a new error class with the `create error class` command. Specify the desired error actions for the class, then change the connection to the database to use the new error class.

If you want to examine the commands in the transaction that caused the connection to be suspended, use the `sysadmin log_first_tran` command. This command writes the first transaction into the exceptions log without skipping it.

Replication Server errors

If a Replication Server error is displayed in the Replication Server error log, refer to the chapter that corresponds to the following conditions.

- If the error message is about function strings or is “function string not found,” refer to “Using incorrect variables in function strings” on page 7 or “29024” on page 51.
- If the error message is about subscriptions, refer to Chapter 5, “Subscription Problems”
- If the error message is about missing objects, there is probably some inconsistency in the RSSD. Call Sybase Technical Support for assistance.

DSI active

When the DSI is active but the replicate data server is not receiving updates, unique transactions might be incorrectly resolved as duplicates, duplicate transactions might be reapplied, or the transactions might be failing on the data server. If the error's action was set to log or `retry_log` using the `assign` action command, the error and failed transaction are written to the exceptions log. You must manually resolve the data inconsistency between the primary and replicate databases or reapply the transactions.

Incorrect duplicate transaction resolution

If the `origin_qid` values stored in a data server or the `rs_exceptslast` are modified by mistake, non-duplicate transactions can be ignored or duplicate transactions could be reapplied. If you suspect that this is happening in your system, check the stored values and compare them with the transactions in the database's stable queue to determine the validity of the values. If the values are wrong, you must modify them directly.

The DSI records the last transaction committed or written into the exceptions log so that it can detect duplicates after a system restart. Each transaction is identified by a unique origin database ID and an origin queue ID that increases for each transaction. For more information about the exceptions log, see "Examining the exceptions log" on page 127.

The last transaction committed from each origin database is recorded on a data server through execution of the function strings defined for the data server's function-string class. For the default function-string class, `rs_sqlserver_function_class`, this is done in the function string of a `commit` command, that is, the `rs_commit` function. The function-string class of every data server must support the `rs_get_lastcommit` function which returns the `origin_qid` and `secondary_qid` for each origin database. The `secondary_qid` is the ID of the queue used for subscription materialization or dematerialization.

The `origin_qid` and `secondary_qid` for the last transaction written into the exception log from each origin is recorded in the `rs_exceptslast` table in the RSSD. However, transactions logged explicitly by the `sysadmin log_first_tran` command are not recorded in this table. These transactions are logged, but they are not skipped.

When a data server interface is started or restarted, it gets the `origin_qid` returned by the `rs_get_lastcommit` function and the one stored in the `rs_exceptslast` table. It assumes that any transaction in the queue with an `origin_qid` less than the larger of these two values is a duplicate and ignores it.

Refer to Appendix A, “Dumping Stable Queues” for more information on how to dump transactions in a queue.

Examining the exceptions log

If a transaction is skipped, either explicitly using the `resume connection` command with the `skip transaction` option, or as a result of the `log action` assigned to data server errors, the skipped transaction is written into the exceptions log. Orphan transactions and transactions logged by the `sysadmin log_first_tran` command are also in the exceptions log.

If a replicate database is not receiving updates, it is possible that the update transactions were skipped and written into the exceptions log. Examine the exceptions log to find out what transactions have been logged.

The exception log consists of three tables: `rs_exceptshdr`, `rs_exceptscmd`, and `rs_systext`. The `rs_exceptshdr` table has one entry for each transaction. The `rs_exceptscmd` table has one entry for each command (either source or output) of the transaction. The `rs_systext` table stores the text of the commands.

To view the header information for all logged transactions intended for a database, log in to the system Adaptive Server for the Replication Server controlling the database, and execute the following query against the RSSD:

```
select * from rs_exceptshdr
       where error_site = data_server
       and error_db = database
       order by log_time
```

The rows are returned in ascending order, by the time the transactions were logged. If you want the rows to appear in descending order, attach “`desc`” to the end of the query.

Join the three system tables to view all information about a transaction in the log. The following query gives the source commands and their corresponding output commands for each logged transaction:

```
select hdr.sys_trans_id, cmd_type, textval
       from rs_exceptshdr hdr,
            rs_exceptscmd cmd,
            rs_systext
```

```
where error_site = data_server
and error_db = database
and hdr.sys_trans_id = cmd.sys_trans_id
and cmd_id = parentid
order by log_time, src_cmd_line,
output_cmd_index, sequence
```

See “7035” on page 44 for more information about identifying the text of the skipped transactions in the exceptions log.

For more information about the system tables that make up the exceptions log, see the *Replication Server Reference Manual*.

Adaptive Server Log Problems

This chapter describes problems that occur in Adaptive Server transaction log.

Name	Page
Overview	129
Truncating an Adaptive Server log	130
Log truncation problems	133
Adaptive Server log corruption	136

If you cannot find the solution in this chapter

If an error message is displayed in an error log and you cannot find it in this chapter, see:

- “Analyzing error logs” on page 10
- Chapter 3, “Common Errors”

If an error message is not displayed, see Chapter 1, “Identifying the Problem”

Overview

The Adaptive Server primary database log is the source of the data that Replication Server distributes. A RepAgent or LTM retrieves transactions from the log and sends them to the Replication Server. Adaptive Server log problems include:

- Log files that have reached the maximum size and need to be truncated
- Missing or incomplete transactions at the primary Replication Server as a result of truncated primary database logs
- Log files corrupted by software and hardware sources

Truncating an Adaptive Server log

When the database log is full, Adaptive Server cannot continue to process updates. To solve this problem you must truncate the log. However, you should not truncate transactions from the log before the RepAgent or LTM has passed them successfully to the Replication Server.

Adaptive Server uses truncation points to ensure that only transactions processed by the RepAgent or LTM are truncated. A **secondary truncation point** marks the place in the primary database log up to which the RepAgent or LTM has processed transactions. The RepAgent or LTM periodically updates the secondary truncation point to reflect transactions successfully passed to the Replication Server. Adaptive Server will not truncate the log past the secondary truncation point. For more information about database log truncation, see Chapter 2, “Replication Server Technical Overview,” in the *Replication Server Administration Guide*.

There are, however, times when it is necessary to truncate the log beyond the secondary truncation point. For example, if the RepAgent or LTM cannot access the Replication Server and the log fills, you may want to truncate the log at the secondary truncation point rather than extending the log or preventing clients from updating the primary database. Truncating the log can cause inconsistencies between the primary and replicate databases. To resolve these inconsistencies, see “Log truncation problems” on page 133.

To truncate the database log at the secondary truncation point:

- 1 Verify that the secondary truncation point is set for the database. See “Verifying the state of the secondary truncation point” on page 130.
- 2 Turn off the secondary truncation point in the database. See “Turning off the secondary truncation point in a database” on page 131
- 3 Dump the database log. See the *Adaptive Server Reference Manual*.
- 4 Set the secondary truncation point. See “Setting the secondary truncation point” on page 132

Verifying the state of the secondary truncation point

To check whether the secondary truncation point is set for a database, use the `dbcc gettrunc` command. For example, to see if the secondary truncation point is set in the Parts database, log in to the primary database and execute:

```
> use Parts
```

```
> go
> dbcc gettrunc
> go
```

In the output that these commands generate, a “0” in the `ltn_trunc_state` column means the secondary truncation point is turned off for the database. A “1” in the `ltn_trunc_state` column indicates the secondary truncation point is turned on.

Turning off the secondary truncation point in a database

Truncating the log frees only log pages that the RepAgent or LTM has sent to the Replication Server with confirmation. To truncate portions of the log that have not been transferred, you must turn off the secondary truncation point for the database by executing the Transact-SQL command `dbcc settrunc(ltm, ignore)` in the database.

Note Only the Adaptive Server System Administrator and the Database Owner have permission to execute the `dbcc settrunc` command.

First, suspend the RepAgent or LTM forwarding the log. Otherwise, executing `dbcc settrunc` results in the following error:

```
The log transfer context for the current database is
already reserved by Adaptive Server process 7. The log
transfer context for the current database is not
reserved.
```

To turn off the secondary truncation point in the database named `Parts`, for example, log in to the Adaptive Server with the database and execute:

```
> use Parts
> go
> dbcc settrunc(ltm, ignore)
> go
```

After you have turned off the secondary truncation point, you can truncate the log using the Transact-SQL `dump transaction` command. See the *Adaptive Server Reference Manual* for more information.

Finally, you should dump the transaction log. This may result in the loss of updates to the replicate database. The primary database is not affected. You must reapply the lost updates to the replicate database. For more information about reapplying updates, see “Log truncation problems” on page 133.

Setting the secondary truncation point

Before you restart a RepAgent or LTM, you need to reestablish the secondary truncation point using `dbcc settrunc(ltm,valid)`. For example, to reestablish the secondary truncation point in the Parts database, log in to the Adaptive Server with the database and execute:

```
> use Parts
> go
> dbcc settrunc(ltm, valid)
> go
```

Use the `dbcc gettrunc` command to verify that the secondary truncation point has been enabled:

```
> use Parts
> go
> dbcc gettrunc
> go
```

The `ltm_trunc_state` column in the output from this command should contain a 1.

After resetting the locator for the database in the `rs_locator` table, resume the Replication Server to restart the RepAgent or LTM. The RepAgent or LTM begins scanning the log from the page returned in the `ltm_truncpage` column by the `dbcc gettrunc` command.

Resetting the locator

When a RepAgent or LTM starts, it uses a **locator** to determine where to begin scanning the database log. The locator for a database log is stored in the `rs_locator` table in the RSSD of the Replication Server that controls the database.

Note This locator is different from the RSI locator. See “Invalid RSI locator” on page 106 for more information about the RSI locator.

The locator points to a record on a page in the database log. If the truncation point for the database is turned off, you or Adaptive Server may truncate the page the locator points to, invalidating the locator. When the locator is invalid, the RepAgent writes error message 9215 (see “9215: log truncated (LTM 1079 with ASE 624)” on page 59) to the Adaptive Server error log. The LTM writes error message 1079 to the LTM error log:

```
E. 95/10/30 14:18:13. ERROR #1079 logscan thread(vishy_pds.pdb) -
/ltmscan.c(1459)
The action for this error is invalid. Internal LTM error. H. 95/10/30 14:18:13.
THREAD FATAL ERROR #1078 logscan thread(vishy_pds.pdb) - /ltmscan.c(1459)
This CT-Lib or Open Server error is fatal. The server the thread connects to
is 'vishy_pds'. I. 95/10/30 14:18:13. No Sybase library information is available
97/10/17 11:57:19.34 RepAgent (10) Error: 9215, Severity: EX_USER, State: 1
The Secondary Truncation Point is not valid. Please make sure that the Secondary
Truncation Point is valid, and then restart the Rep Agent Thread for this
database.
```

To recover, reset the locator to 0x0. This directs the RepAgent or LTM to begin scanning the log from the truncation point.

To reset the locator, update the `rs_locator` table in the RSSD for the Replication Server controlling the database using the Adaptive Server stored procedure `rs_zeroltm`:

```
rs_zeroltm data_server, database
```

where `data_server` and `database` are the respective Adaptive Server and database whose LTM locator value you want to reset.

For example, to reset the LTM locator for the MYDS Adaptive Server and the Parts database, execute;

```
rs_zeroltm MYDS, Parts
```

Log truncation problems

Truncating the log of the primary database may lead to these types of inconsistencies:

- Orphaned transactions can cause the stable queues to fill.
- Inbound queue requires more disk space.
- Parts of transactions, or entire transactions, may be missing, causing inconsistent data at replicate sites. See Chapter 19, “Replication System Recovery,” in the *Replication Server Administration Guide* for more information.
- Subscription commands may be lost, disrupting the materialization or dematerialization protocol and requiring manual cleanup.

- If a log is truncated while a route is being created, subscription commands against the RSSD may be lost. As a result, the route cannot be created. If this occurs, drop and re-create the route.

For more information about how to recover lost transactions, see Chapter 19, “Replication System Recovery,” in the *Replication Server Administration Guide*.

Orphaned transactions

An orphaned transaction is a transaction in an inbound stable queue that is missing a terminating commit or rollback command. Because Replication Server does not free up a queue segment until every transaction in the queue segment (or in any of the preceding queue segments) has been committed or rolled back, orphaned transactions can cause the inbound queue to fill up.

To check for orphaned transactions:

- 1 Log in to the Replication Server and execute `admin who, sqt` to display information about the stable queues at the Replication Server.
- 2 Use the output from `admin who, sqt` to identify the entry for the inbound queue of the database whose log was truncated.

The queue has a two-part name formed from the database ID and the queue-type identifier. For an inbound queue, the queue-type identifier is 1. For example, if the database ID is 101, the inbound queue name is 101:1.

- 3 If the output for the inbound queue shows an open transaction that does not change over long periods of time, an orphaned transaction is probably in the queue. However, because it is difficult to distinguish between an orphaned transaction and a very long transaction, you must further verify that an orphaned transaction exists.

To verify that there is an orphaned transaction:

- a Dump the stable queue and examine the information about the transaction. Refer to Appendix A, “Dumping Stable Queues” for instructions. You can dump just the begin record to find the user name and the time the command was executed, and then see if the user still has open transactions in the database.

- b Even if you find that the user does not have open transactions, the queue may still have an orphaned transaction. If the queue is large, the RepAgent or LTM may not be keeping up with it. To determine whether this is the case, dump the last block in the queue and look at the date for the commands. Use `admin who`, `sqm` to find the last block.
- 4 If the transaction is orphaned, use the `sysadmin purge_first_open` command to skip it.
- 5 Use the output from dumping the stable queue to manually apply the orphaned transaction to the replicate database.

Inbound queue requires more disk space

If the inbound queue requires more disk space, a message similar to the following is displayed in the Replication Server error log:

```
_SQM_ADD_SEGMENT: Going to wait for a segment to be
freed. Administrative action is needed.
```

To determine which stable queues hold the segments, log in to the RSSD and execute the following query:

```
select q_number, q_type, count(*)
from rs_segments
group by q_number, q_type
```

The output of this query shows the number of segments each stable queue is using. Each segment is 1MB of disk space. Look at the queues with the largest number of segments:

- A *q_type* of 0 means an outbound queue. If any of the outbound queues have more than two segments, verify that the connections they service are operating. See Chapter 8, “Data Server Interface Problems” and Chapter 6, “Replication Server Interface Problems” for more details.
- A *q_type* of 1 is an inbound queue. If one of these queues has a large number of segments, it may contain an orphaned transaction. See “Orphaned transactions” on page 134 for more information.

Adaptive Server log corruption

Adaptive Server log corruption can cause failures anywhere along the replication path. The most common symptoms are:

- RepAgent or LTM shuts down when it reads log data it cannot process.
The RepAgent or LTM shutting down is the most common symptom of a corrupt Adaptive Server log. In addition, the error log may reveal error messages that are inconsistent with the current operation; however, no specific log corruption error message is displayed.
- Adaptive Server shuts down. The LTM logscan thread, which reads the Adaptive Server log, terminates when it reads corrupted log data.
- A component of the Replication Server (such as the DSI) shuts down.
When the RepAgent or LTM passes incorrect data to the Replication Server, Replication Server components, such as the DSI, may shut down. For instance, if the DSI encounters a mismatched data type and fails to resolve the differences, it shuts down.
- Replicated data appears out of sync.

The procedures for determining log corruption and where the corruption occurs is complex. It requires knowledge of the log format as well as Adaptive Server and Replication Server internal components.

Note If you suspect Adaptive Server log corruption, contact Sybase Technical Support to help you diagnose and resolve the issue.

Dumping Stable Queues

This appendix describes how to dump stable queues.

Name	Page
Introduction	137
Using sysadmin dump_queue	139
Using sysadmin sqt_dump_queue	150
Using sysadmin dump_file	152
sysadmin purge commands	153

Introduction

Stable queues are composed of segments stored on disk partitions. Each segment has a megabyte of message space, divided into 64 blocks of 16KB. Segments store messages passed by the system. Each block in a segment corresponds to a physical data transfer between disk and memory. A segment is allocated to just one queue—different queues cannot share segments. Each block within a segment may store one or more messages. A map of the system segments is stored in the `rs_segments` table of the RSSD.

Replication Server deletes messages from queues on a segment basis. Internally, it deletes blocks from the segments. However, the freed space does not appear in the system tables until all of the blocks contained in the 1MB segment are deleted.

When a stable queue is dumped, a segment may contain deleted blocks along with the undeleted blocks. You can identify the undeleted blocks of a segment by checking the `First Seg.Block` column in the `admin who, sqm` output for the queue.

Sometimes queues are dumped while Replication Server is actively adding and deleting segments from the queues. At such times, the output of the `admin who, sqm` command may not accurately represent the contents of the stable queues. Segments may be deleted before you can enter the `sysadmin dump_queue` command. To prevent Replication Server from changing a stable queue while you examine its contents, you can suspend the database connection or route that uses the stable queue. You can also start the Replication Server in standalone mode to prevent changes to a stable queue. For more information, see the *Replication Server Administration Guide*.

This list briefly describes each command:

- `sysadmin dump_queue` – dumps the contents of the transaction cache of an inbound or outbound stable queue.
- `sysadmin sqt_dump_queue` – dumps the contents of the transaction cache of an inbound queue.
- `sysadmin dump_file` – specifies an alternate log file to be used by `sysadmin dump_queue` or `sysadmin dump_queue`.
- `sysadmin purge_all_open` – purges all open transactions from an inbound queue.
- `sysadmin purge_first_open` – purges only the first open transaction from an inbound queue.

Using trace

The outbound queue for the replicate database contains default commands. If you have created function strings in addition to the defaults, you may want to use a special DSI trace flag to print the actual commands that are being sent to replicate databases.

To set traces, you need to run the debug replication server, `repserver.diag`, which should be located in the *bin* directory.

To set traces, add the following lines into the `.cfg` file of the Replication Server you want to trace.

- 1 This line forwards the output to the standard output or to the *errorlog*:

```
trace=GEN,TRC_STDERR or trace=GEN,TRC_ERRLOG
```

- 2 This line gives the output of commands that are sent to the replicate site.

```
trace=DSI,DSI_CMD_DUMP
```

It logs the commands in the form:

```
Command sent to 'RDS.RDB' : <command>
```

Using *sysadmin dump_queue*

The `sysadmin dump_queue` command dumps the contents of a Replication Server stable queue to the Replication Server log, an alternate log file, the client, or to tables in the RSSD. You can examine the dump to see the messages in the stable queue. This may help you to identify and reapply transactions that have not been applied or applied incorrectly to the replicate database.

Here is the syntax for the `sysadmin dump_queue` command:

```
sysadmin dump_queue, q_number, q_type, seg, blk, cnt [,RSSD | client]
```

q_number, *q_type* – identifies the stable queue to dump. Examine the output of the `admin who` command or the `admin who, sqm` command to find these values.

seg – identifies the starting segment. If set to -1, the starting segment is the first active segment in the queue. If set to -2, the starting segment is the first segment in the queue and includes any inactive segments retained by setting a save interval.

blk – identifies the 16K block in the segment where the dump is to begin. Block numbering starts at 1 and ends at 64. The following two additional options are used with the *seg* setting:

- If *seg* is -1 and *blk* is -1, then start with the first undeleted block of the queue.
- If *seg* is -1 and *blk* is -2, then start with the first unread block of the queue. For example, if the RSI and DSI are down, then executing `sysadmin dump_queue` with these parameters dumps the unread segments in the queue.

cnt – specifies the number of blocks to dump. This number can span multiple segments. If set to -1, the end of the current segment is the last block dumped. If set to -2, the end of the queue is the last block dumped.

RSSD – optional flag that forces the output to system tables in the RSSD. If you use this option, insert a comma before it. Sending output to the RSSD is the slowest option.

- If the RSSD option is used, then:

- System tables are first cleared of the segments with the same *q_number*, *q_type*, *seg*, and *blk* as the blocks are being dumped, and
- Dump is written into two tables in the Replication Server system database: *rs_queuemsg* and *rs_queuemsgtxt*.
- The *rs_queuemsg* table has the following columns:
 - *q_number*
 - *q_type*
 - *q_seg*
 - *q_blk*
 - *q_row*
 - *len*
 - *origin_site_id*
 - *origin_q_id*
 - *origin_time* (used only at the beginning of a transaction)
 - *origin_user*
 - *tran_name*
 - *local_q_id*
 - *status*
 - *tran_len*
 - *txt_len*
 - *tran_id*
 - *lorigin_site_id*
 - *version*

Only the begin transaction command uses the *origin_time*, *origin_user*, and *tran_name* columns.

- The *rs_queuemsgtxt* table holds the text of commands dumped from the stable queue. If the text of a command exceeds 255 characters, it is stored in multiple rows numbered with the *q_seg* column.

client – optional flag that forces the output to the client that is issuing this command.

- If the client option is used, the dump is written into the client that is issuing this command, such as isql or Replication Server Manager.
- The client output has the following columns:
 - Q Number
 - Q Type
 - Segment
 - Block
 - Row
 - Message Len
 - Orgn Siteid
 - Orgn Time
 - Orgn Qid
 - Orgn User
 - Tran Name
 - Local Qid
 - Status
 - Tranid
 - Logical Orgn Siteid
 - Version
 - Command Len
 - Seq No
 - Command

If you do not specify the RSSD or client option, the output will go into the Replication Server log. If you use the `sysadmin dump_file` command to specify an alternative log file for dumping queues, the output will go into the alternate dump file.

Confirming suspected problems

This section describes how to use `sysadmin dump_queue` to confirm that a suspected problem exists.

When the Replication Server runs out of stable queue segments, an uncommitted transaction may remain in the inbound stable queue. To find out if this has happened, follow these steps:

- 1 Log in to the Replication Server.
- 2 Execute `admin who, sqt`.
- 3 Use the output from `admin who, sqt` to determine which inbound queue has an open transaction. It will appear as an open transaction with the state marked "st:O". The Info column of this row holds the queue number and type. The transaction state is followed by the number of commands and then a local queue ID for the begin transaction command in the form `segment:block:row`.

In the following example of `admin who, sqt` output, the transaction has a queue number of 103 and a queue type of 1, is in an open state and consists of three commands. The begin transaction command is located in segment 21, block 28, row 0:

```
103:1 st:O,cmds:3,qid:21:28:0
```

- 4 Execute the following command to dump block 28 of segment 21:

```
sysadmin dump_queue, 103, 1, 21, 28, 1, RSSD
```

- 5 Replication Server writes the text of the commands in block 28 of segment 1 into the `rs_queuemsg` and `rs_queuemsgtxt` system tables or into the Replication Server log.
- 6 If you use the RSSD option, you can view information about the command in the `rs_queuemsg` system table by executing this command on the RSSD:

```
select * from rs_queuemsg
where q_number = 103 and q_type = 1 and
q_seg = 21 and q_blk=28 and q_row = 0
```

Execute a similar query against `rs_queuemsgtxt` to obtain the text of the command.

***sysadmin dump_queue* examples**

The following examples demonstrate the use of the `sysadmin dump_queue` command.

- This command acts on queue 103:1. It dumps blocks 15 through 64 of segment 0 and blocks 1 through 15 of segment 1 into the Replication Server log:


```
sysadmin dump_queue, 103, 1, 0, 15, 65
```

- The following command dumps all of queue 103:1 into the Replication Server log:

```
sysadmin dump_queue, 103, 1, -1, 1, -2
```

- This command dumps all of queue 103:1 into the RSSD:

```
sysadmin dump_queue, 103, 1, -1, 1, -2, RSSD
```

- This command dumps all of queue 103:1 into the client:

```
sysadmin dump_queue, 103, 1, -1, 1, -2, client
```

- The following series of commands dump all of queue 103:1 into a file named *qdump*. The last `sysadmin dump_file` command closes *qdump*.

```
sysadmin dump_file, qdump
sysadmin dump_queue, 103, 1, -1, 1, -2
sysadmin dump_file
```

Interpreting the output of a dump queue

Stable queues consist of partitions. Partitions are created with the `add partition` command. Whenever a partition is added to the replication server, it is divided into 1 megabyte segments. Segments are divided into 64 blocks of 16 KB. Messages have fixed size headers, variable size messages, and are written into blocks.

The following examples shows the output of an `sysadmin dump_queue` command on various types of queues and how to read the output of each.

Example 1: outbound queue after *create subscription*

In this example, the `create subscription` command at the replicate Replication Server starts the materialization process by selecting from the primary database:

```
create subscription emp_queue_sub
for emp_queue
with replicate at LNDS.lndb
```

To view the entire queue:

```
sysadmin dump_queue, 103, 0, -1, 1, -2
```

Output

```
95/06/28 16:24:15. BLOCK BEGIN q_number=103 q_type=0
blk=0:20 cnt=1
I. 95/06/28 16:24:15. Begin Transaction Origin User=
Tran Name=
I. 95/06/28 16:24:15. ENTRY len=176 orig=102
oqid=0000000100000a5c0000028e000dfffffff00000000883c0
10dd700000100000000000 lqid=0:20:0 st=21 tr=
'0000000100000a' \ '0009' NYDSnydb comlen=61 activate
subscription 0x010000658000006f 0
```

Explanation

- **BLOCK BEGIN**
The beginning of a 16K block.
- **q_number=103**
The queue ID.
- **q_type=0**
The queue type (0 = outbound, 1 = inbound).
- **blk=0:20**
Block identifier (segment 0: block 20).
- **cnt=1**
The number of entries (commands) in this block.
- **Begin Transaction Origin User= Tran Name=**
Beginning of a transaction with originating user and the transaction name. *Tran Name* is empty when the original transaction at the primary database was explicitly started with a begin transaction without supplying a transaction name.
- **ENTRY**
The actual message.
- **len=176 orig=102**
The message length and the originating site ID.
- **oqid=0000000100000a5c0000028e000dfffffff00000000883c010dd700000100000000000**

The origin queue ID: the queue ID assigned to the command string at the place the command originated from. This ID is used for duplicate detection.

- Lqid=0:20:0

Local queue ID: Commands logged by an SQM are assigned a local queue ID. This queue ID is sent to the next queue as the sender queue ID and is used by the destination to acknowledge the command. *Lqid* consists of segment number, block number and row number. Row numbers start at 0.

- St=21

Status and command type. The field is a bitmask which has status and command type set. This bit's settings are 0x01 and 0x20, a commit and materialization command.

0x01 = status of a commit

0x02 = status of a rollback

0x04 = status of a begin

0x08 = status of a orphan statement

0x10 = status of a materialization statement

0x20 = status of a materialization statement

0x40 = status of a tran going to primary

0x80 = ignore duplicates for DDL commands

0x0100 = status of a materialization command

0x0200 = dump/load message

0x0400 = routing message

0x0800 = flag for DDL generation

0x1000 = DDL generation at DSI flag

0x2000 = special SQM message

0x4000 = SQT mini-abort

0x8000 = message has been zapped

- tr= '000000010000a' \ '0009' NYDSnydb comlen=61 activate subscription 0x010000658000006f 0

Transaction ID and command, *comlen* indicates the length of the command.

Example 2: outbound queue after update

In this example, an update has been performed at the primary database. The minimal column feature has not been set for the replication definition:

```
update emp_queue set emp_first =
"General" where emp_id = "123456789"
```

To view the outbound queue or the replicate:

```
sysadmin dump_queue, 103, 0, -1,1,-2
```

Output

```
BLOCK BEGIN q_number=103  q_type=0 blk=0:21 cnt=3
I. 95/06/28 16:42:37.      Begin Transaction Origin
User=sa Tran Name=upd
I. 95/06/28 16:42:37.      ENTRY len=184 orig=102
oqid=0000000100000a670000028f0007ffffffff00000000883c0
1131c10000100000000001 lqid=0:21:0 st=4 tr=
'0000000100000a' g '0002' NYDSnydb comlen=71 begin
transaction upd
I. 95/06/28 16:42:37.      ENTRY len=328 orig=102
oqid=0000000100000a670000028f0007ffffffff00000000883c0
1131c10000100000000002 lqid=0:21:1 st=0 tr=
'0000000100000a' g '0002' NYDSnydb comlen=216 update
emp_queue set emp_id='123456789',emp_last='Davis
',emp_first='General          ',emp_title='Process
Engineer                      ',emp_date='Jan  1 1995
12:00:00:000AM',branch_id=1111 where
emp_id='123456789'
I. 95/06/28 16:42:37.      ENTRY len=168 orig=102
oqid=0000000100000a670000028f0007ffffffff00000000883c0
1131c10000100000000003 lqid=0:21:2 st=1 tr=
'0000000100000a' g '0002' NYDSnydb comlen=55 commit
transaction
```

Explanation

- cnt=3
There are 3 rows in this block.
- Tran Name=upd

If you perform an update and did not explicitly start a transaction, Adaptive Server uses *upd* as the transaction name. Likewise, Adaptive Server uses *del* for a delete command and *ins* for an insert command.

- `st=4, st=0, st=1`

Status for begin, other, and commit.

Example 3: inbound queue

In this example, these commands were executed at a primary database, NYDS.nydb:

```
sp_setreplicate emp_queue, true
insert emp_queue values("123456789", "Davis", "Gen",
"Process Engineer","1/1/95", 1111)
insert emp_queue values("987654321", "Irvine", "Ben",
"Microwave Engineer","3/5/94", 2222)
grant all on emp_queue to public
```

To view the inbound queue of the primary Replication Server:

```
sysadmin dump_queue, 102,1,-1,1,-2
```

Output

```
I. 95/06/28 16:28:27. BLOCK BEGIN q_number=102 q_type=1
blk=0:35 cnt=2
I. 95/06/28 16:28:27.      Begin Transaction Origin
User=sa Tran Name=user_transaction I. 95/06/28
16:28:27.      ENTRY len=312 orig=102
oqid=0000000100000a5c0000028e00060000028e000600000883c0
10ca88900000000000000000 lqid=0:35:0 st=4 tr=
'0000000100000a' \ '0006' NYDSnydb comlen=200 begin
transaction user_transaction
I. 95/06/28 16:28:27.      ENTRY len=312 orig=102
oqid=0000000100000a5c0000028e0008ffffff000000000883c0
10ca88900010000000000000 lqid=0:35:1 st=1 tr=
'0000000100000a' \ '0006' NYDSnydb comlen=200 commit
transaction
I. 95/06/28 16:28:27. BLOCK BEGIN q_number=102 q_type=1
blk=0:36 cnt=3
I. 95/06/28 16:28:27.      Begin Transaction Origin
User=sa Tran Name=sub010000658000006f_ffffff
I. 95/06/28 16:28:28.      ENTRY len=312 orig=102
oqid=0000000100000a5c0000028e00090000028e000900000883c0
10dd70000000000000000000 lqid=0:36:0 st=4 tr=
```

```
'0000000100000a' \ '0009' NYDSnydb comlen=200 begin
transaction sub010000658000006f_ffffff
I. 95/06/28 16:28:28.          ENTRY len=156 orig=102
oqid=0000000100000a5c0000028e000b0000028e000900000883c0
10dd70000000000000000000 lqid=0:36:1 st=16 tr=
'0000000100000a' \ '0009' NYDSnydb comlen=44 activate
subscription 0x010000658000006f 0
I. 95/06/28 16:28:28.          ENTRY len=312 orig=102
oqid=0000000100000a5c0000028e000dffffff000000000883c0
10dd70000001000000000000 lqid=0:36:2 st=1 tr=
'0000000100000a' \ '0009' NYDSnydb comlen=200 commit
transaction
I. 95/06/28 16:28:28. BLOCK BEGIN q_number=102 q_type=1
blk=0:37 cnt=2
I. 95/06/28 16:28:28.          Begin Transaction Origin
User=sa Tran Name=grrev
I. 95/06/28 16:28:28.          ENTRY len=312 orig=102
oqid=0000000100000a5c0000028e000e0000028e000e00000883c0
10e290700000000000000000 lqid=0:37:0 st=4 tr=
'0000000100000a' \ '000e' NYDSnydb comlen=200 begin
transaction grrev
I. 95/06/28 16:28:28.          ENTRY len=312 orig=102
oqid=0000000100000a670000028f0001ffffff000000000883c0
10e2907000010000000000000 lqid=0:37:1 st=1 tr=
'0000000100000a' \ '000e' NYDSnydb comlen=200 commit
transaction
I. 95/06/28 16:28:28. END QUEUE DUMP FOR 102:1
```

Explanation

- Tran Name=user_transaction
A transaction was started with the name user_transaction.
- Q_type=1
An inbound queue.
- Tran Name=sub010000658000006f_ffffff
The Replication Server starts the activate subscription transaction with a generated name.
- Tran Name=grrev

The grant and revoke permission commands get a special transaction name from the Adaptive Server.

Note You will see all begin, commit, and rollback commands regardless of whether or not the transactions contain changes that need to be replicated. When the LTM reads the log, it does not know whether or not the transaction includes changes to objects that are marked for replication. Because future versions of Replication Server may change this scheme to optimize for network traffic, you should check your current version of Replication Server for this behavior.

Example 4: inbound queue

In this example, this update command was executed at the primary database:

```
update emp_queue set emp_first = "General" where emp_id
= "123456789"
```

Execute the sysadmin dump_queue command to view the inbound queue of the primary Replication Server:

```
sysadmin dump_queue, 102,1,-1,1,-2
```

Output

```
I. 95/06/28 16:46:42. BLOCK BEGIN q_number=102 q_type=1
blk=0:38 cnt=3
I. 95/06/28 16:46:42.      Begin Transaction Origin
User=sa Tran Name=upd
I. 95/06/28 16:46:42.      ENTRY len=312 orig=102
oqid=0000000100000a670000028f00020000028f00020000883c0
1131c1000000000000000000 lqid=0:38:0 st=4 tr=
'0000000100000a' g '0002' NYDSnydb comlen=200 begin
transaction upd
I. 95/06/28 16:46:42.      ENTRY len=776 orig=102
oqid=0000000100000a670000028f00050000028f00020000883c0
1131c1000000000000000000 lqid=0:38:1 st=0 tr=
'0000000100000a' g '0002' NYDSnydb comlen=661 update
emp_queue set emp_id='123456789',emp_last='Davis
',emp_first='General
',emp_title='Process
Engineer
',emp_date='Jan  1 1995
12:00:00:000AM',branch_id=1111 where
emp_id='123456789'
I. 95/06/28 16:46:42.      ENTRY len=312 orig=102
oqid=0000000100000a670000028f0007ffffff00000000883c0
```

```
1131c1000010000000000000000000 lqid=0:38:2 st=1 tr=  
'0000000100000a' g '0002' NYDSnydb comlen=200 commit  
transaction
```

Explanation

- Tran Name=upd
Adaptive Server uses transaction name upd for implicit update transactions. RepAgent reads this name from the log.

Using sysadmin sqt_dump_queue

Use this command to dump contents of the transaction cache for each inbound queue. `sysadmin sqt_dump_queue` dumps all the statements of transactions that have not yet been deleted.

Commands are dumped into the Replication Server log or an alternate log file created using the `sysadmin dump_file` command. In the current version of Replication Server, you cannot dump the transaction cache into the RSSD.

Before using this command, make sure the transaction cache for the database exists using `admin who, sqt`.

Output from `sysadmin sqt_dump_queue` shows the state of transactions in the transaction cache: Open, Closed, or Read. Open transactions do not yet have a commit. Closed transactions have a commit but have not yet been completely read. Read transactions have been completely read but have not yet been deleted.

Following is the syntax for the `sysadmin dump_queue` command:

```
sysadmin sqt_dump_queue, q_number, q_type, reader [,open]
```

q_number, *q_type* – identifies the stable queue to dump. These values can be found by using the `admin who` command, the `admin who, sqm` command, and the `admin who, sqt` command.

reader – identifies the reader you want to dump the stable queue. This parameter applies to features that require multiple readers, such as warm standby applications. You can get the reader number by running `admin sqm_readers` or `admin who, sqt`. If you are not using multiple readers, enter “0.”

open – allows you to dump only open transactions. If you use this option, insert a comma before it.

Examples

The following examples demonstrate the use of `sysadmin sqt_dump_queue`.

- This command acts on inbound queue 103:1. All undeleted transactions in the queue are dumped from the transaction cache:

```
sysadmin sqt_dump_queue, 103, 1, 0
```

- The following series of commands dumps open transactions into the file `qdump` from the transaction cache for queue 103:1. The last command closes `qdump`:

```
sysadmin dump_file, qdump
sysadmin sqt_dump_queue, 103, 1, 0, open
sysadmin dump_file
```

- In an `isql` session against the `NYDS.nydb` database, run a `begin tran` test to start an open transaction.

The output of the `admin who, sqt` command shows one queue with an open transaction:

```
SpidState      Info
-----      -
20  Awaiting Wakeup 102:1 NYDS.nydb

ClosedReadOpen Trunc
-----
0 0      1      1

RemovedFullSQM BlockedFirst Trans
-----
0 0      1      st:0,cmds:1,qid:0:42:0
Parsed
-----
0
```

- Running `sysadmin sqt_dump_queue,102,1, 0` returns the dump output:

```
I. 96/06/28 17:18:50. TRANSACTION CACHE for 102:1
I. 96/06/28 17:18:51.      TRAN_OPEN TID=
'0000000100000a' g '0013' NYDSnydb Origin=102. 1
commands. in memory
I. 96/06/28 17:18:51.      Tran Name 'test'
User:'sa'
I. 96/06/28 17:18:51.      LQID 0:42:0      begin
transaction test
I. 96/06/28 17:18:51. END TRANSACTION CACHE for
102:1
```

Using sysadmin dump_file

Use this command to specify an alternate log file to be used by the sysadmin dump_queue and sysadmin sqt_dump_queue commands for output.

Here is the syntax for the sysadmin dump_file command:

```
sysadmin dump_file [, file_name]
```

where *file_name* is the local name of the file where dumps will be directed.

If you omit *file_name*, the current dump file is closed. If you specify *file_name* and a dump file currently exists, then the current dump file is closed and a new file is opened. The new file uses the name specified with the optional *file_name* entry.

Examples

These examples demonstrate how to use sysadmin dump_file.

- This command dumps the contents of queue 100:1 into the Replication Server log:

```
sysadmin dump_queue, 100, 1, -1, 1, -2
```

- The following series of commands use sysadmin dump_file to specify an alternate log file.

This command specifies *file1* as the alternate log file:

```
sysadmin dump_file, file1
```

This command outputs the contents of queue 100:1 into *file1*:

```
sysadmin dump_queue, 100, 1, -1, 1, -2
```

This command closes *file1* and opens *file2*. *file2* receives the output from subsequent dumps:

```
sysadmin dump_file, file2
```

This command dumps transaction statements from queue 100:1 into *file2*:

```
sysadmin sqt_dump_queue, 100, 1, 0
```

This command closes *file2*:

```
sysadmin dump_file
```

This command dumps the output from queue 103:1 into the Replication Server log:

```
sysadmin dump_queue, 103, 1, 0, 11, 63
```

***sysadmin purge* commands**

Use the `sysadmin purge_all_open` command to purge all open transactions from a Replication Server inbound queue. Use `sysadmin purge_first_open` to purge only the first open transaction from a Replication Server inbound queue.

Replication Server needs sufficient storage to execute the `sysadmin purge` commands. If you do not have enough storage, the following error message may appear:

```
This RS is out of Disk Space. Use another session to add
disk space for this command to proceed.
```

If this occurs, start another `isql` session and add stable storage to the Replication Server. `sysadmin purge` commands cannot proceed until sufficient storage is available.

If you want to review the contents of the transactions being dropped, then execute the `sysadmin sqt_dump_queue` command prior to using the `sysadmin purge_all_open` or `sysadmin purge_first_open` command.

The syntax for the `sysadmin purge_all_open` command is:

```
sysadmin purge_all_open, q_number, q_type
```

where *q_number*, *q_type* identifies the stable queue to purge. These values can be found by using the `admin who` command, the `admin who, sqm` command, and the `admin who, sqt` command. Only inbound queues can be purged.

The syntax for the `sysadmin purge_first_open` command is:

```
sysadmin purge_first_open, q_number, q_type
```

where *q_number*, *q_type* identifies the stable queue to purge. These values can be found by using the `admin who` command, the `admin who, sqm` command, and the `admin who, sqt` command. Only inbound queues can be purged.

Examples

The next set of examples describes the use of `sysadmin purge_first_open` and `sysadmin purge_all_open`.

- The following command purges the first open transaction from the specified queue:

```
sysadmin purge_first_open, 103, 1
```

- This command purges all open transactions from the specified queue:

```
sysadmin purge_all_open, 103, 1
```


Language, Sort Order, and Character Set Issues

This chapter describes localization of messages into different languages, sort order, character sets, and the problems that can occur with each feature.

Name	Page
Overview	155
Message language	156
Sort order	157
Character sets	158

Overview

Replication Server provides support for international environments with:

- Localization of messages into English, French, German, and Japanese languages
- Support for all Sybase-supported character sets, with character set conversion between Replication Server sites
- Support for non-binary sort orders

Because Replication Server works in a distributed environment with many independent processes, you must set up its languages, sort orders, and character sets carefully to avoid problems. This appendix describes some of the problems that can arise when you do not set up these items according to the guidelines recommended in the *Replication Server Design Guide*.

Note RepAgent uses the same language, sort order, and character set settings as the Adaptive Server.

Message language

You can configure Replication Servers and LTMs to print messages to the error log and to clients in English, French, German, or Japanese. This allows server programs in your replication system, including Replication Servers, LTMs, and Adaptive Servers, to write messages to their error logs in the configured language. However, depending on the server, they may or may not send messages to a client in the client's language.

For example, Adaptive Server checks for the language setting of its client (Replication Server or LTM) and returns messages in that language. Replication Server and the LTM, however, do not check for a client's language; instead, they return messages to a client in their own language. This can result in error logs with messages in different languages if the servers are configured in different languages.

To avoid the confusion that can result from a mixed-language error log, Sybase recommends that you configure the Replication Server and LTM so that all languages are the same.

Stored procedure messages

Stored procedures use the language of their Adaptive Server connection. Therefore, if you have a mixed-language environment, Adaptive Server stored procedure messages may be in a different language from the Replication Server and LTM.

For example, if you log in to Adaptive Server using French and execute a stored procedure, the stored procedure generates messages in French, even if the language setting of Replication Server or LTM is English or Japanese. If you log in to Adaptive Server with a language that is not installed at the Adaptive Server, Adaptive Server returns stored procedure messages in English.

Other situations that produce mixed languages

- In order to quickly and efficiently release EBFs and maintenance releases, there may be times when Sybase does not translate new or changed messages into French, German, or Japanese. If this is the case, the Replication Server, LTM, or rs_subcmp messages use the English version of the message.

- Message localization files may be misplaced. Verify that the files are in the correct directories.

Sort order

The sort order, or collation sequence, used by a server determines how character data and identifiers are compared and ordered. A sort order is determined by the character set.

Replication Server supports all Sybase-supported sort orders, including non-binary sort orders. Non-binary sort orders are necessary for the correct ordering of character data and identifiers in European languages.

Use a non-binary sort order if you have *both* of the following:

- Differing character sets between the primary and replicate data servers, and
- Data with 8-bit characters,

and you intend to do *either* of the following:

- Include columns with this data in the where clause of a subscription, or
- Query the database with an order by clause involving the columns with this data.

Sort order problems

Most sort order problems in Replication Server are a result of different sort orders being used in different servers throughout your system. If you are using different sort orders, you may experience the following problems:

- You receive a DSI error when trying to insert or update a record into a table with a unique index on a character column. Under the primary data server's sort order, the value is unique; but under the replicate data server's sort order, it is not unique. For example, "Ivis" and "LVIS" are considered distinct under a binary sort order but are considered equal under the nocase sort order.
- A drop subscription with purge command does not completely remove the rows at the replicate database. A query to select records to purge misses some records under the replicate data server's sort order.

- Replicated data is not what the replicate site expects.

When this occurs, a subscription probably has one or more character clauses. Updates during the lifetime of a subscription are processed with the primary Replication Server's sort order, leading to unexpected semantics at the replicate data server if the sort orders differ.

These problems can be corrected by editing the Replication Server and LTM configuration files to use the same sort order and, if necessary, by running `sybinit` to make the primary data server and replicate data server use this sort order as well.

Character sets

The following problems can arise when servers have different or incompatible character sets. These problems can be corrected by configuring all servers in your replication system to use compatible character sets.

- 8-bit or multibyte data is corrupted when replication is taking place between a multibyte and a single-byte data server.

Sybase does not support any character-set conversion between single and multibyte systems; it simply passes the data and object names as is. Since 8-bit and multibyte characters are character-set-specific, their semantics change after replication.

- Character data in the replicate data server has some question marks (?) substituted for non-ASCII characters, even though the primary and replicate data servers use compatible character sets. Also, object names originating at a remote site have some question mark (?) characters.

This happens because different character sets, even compatible ones, have some unique characters. When trying to convert a character that does not exist in the replicate data server's character set, a question mark (?) is substituted for the unrecognized character.

If you are replicating between compatible character sets (for example, `iso-1` and `cp850`), make sure that your object names and character data do not include any of the characters that are not common to both character sets.

Index

Numerics

8-bit data 158

A

Adaptive Server

adding data servers 3

error logs 10

errors 114

errors causing Replication Server errors 26

errors, correcting 40

alter route command 70

analyzing error logs 10

analyzing Replication Server error messages 29

asynchronous commands, errors from 27, 123

asynchronous transaction failures 123

atomic materialization 82, 83

determining if subscription needs recovery 82

monitoring 82

automatic error correction 123

B

binaries, changing 3

binary sort orders 157

bulk dematerialization 96, 99, 100

bulk materialization 92

C

cache for RSSD too small 53

case in RCL commands xiv

character sets

conversion of 158

guidelines 158

problems with 158

client applications, adding 4

configuration problems 6

connect source command, not executed 51

connections

DSI failures 122

none 50

phantom 110

resuming 122

RSI errors 103

conventions

examples xii

icons xv

syntax statements xii

conversion, character set 158

create route command 65

create subscription errors 79

D

data loss

DSI 55

RSI 55

data servers

customizing error handling 125

error logs 10

error logs, non-Adaptive Server 26

errors 123

database log, locator 132

databases

finding maintenance user name and password 122

large 4

listing 118

no connection to 50

skipping transactions for 123

truncation point 130

dbcc gettrunc command 130

dematerialization

bulk 99

bulk, using without purge option 100

Index

- deleting subscription data and 96
- errors 96
- options for 95
- overview 95
- purge, list of states 97
- purging subscriptions and 96, 97
- without deleting subscription data 96, 100
- detecting loss message 22
- diagnostic tools, description of 12
- DIST thread down 49
- drop route with nowait command 73
- drop subscription command 95
 - with purge option 96, 97
 - without purge option 96
- dropping replicated objects 114
- DSI
 - connection failures 122
 - data loss 55
 - detecting duplicates 126
 - rs_update_lastcommit not replicated 56
 - states 119
 - verifying for replicate 120
- dumping
 - stable queues 137
 - transaction logs 42, 47

E

- enable markers, invalid 61
- errmsg.ltm 11
- errmsg.rs 10
- errmsg.rsm 10
- errmsg.sub 11
- error block, definition 29
- error correction, automatic 123
- error logs
 - Adaptive Server 10
 - analyzing 10
 - data server 10
 - gateway 10, 26
 - LTM 10, 26
 - mixed languages 156
 - monitoring with RSM 26
 - networks 26
 - overview of analyzing 25

- replication agents 10
- Replication Server 10, 11, 29
- rs_init 11
- error messages
 - analyzing Replication Server 29
 - error block 29
 - fatal 26
 - format for Replication Server 27
 - informational 26
 - internal 26
 - multiple, for same problem 29
 - Open Client/Server 29
 - procedure for finding 10
 - RepAgent, analyzing 33
 - RepAgent, format 32
 - RepAgent, overview 31
 - RepAgent, state 32
 - thread-terminated 26
 - warnings 26
- error numbers 28
- errors
 - 1033 57
 - 11061 45
 - 13045 48
 - 15020 50
 - 15040 50
 - 2028 61
 - 29024 51
 - 37022 52
 - 9202 57
 - asynchronous transactions 123
 - connection failures 123
 - creating routes 66
 - data server 123
 - data server, customizing handling 125
 - error numbers 28
 - from Adaptive Server 114
 - normalization 111
 - recoverable 111
 - RSI connection 103
- events causing problems 2
- examples
 - RepAgent error message 32
 - Replication Server error log 30
 - style conventions xi
- exceptions log 127

F

- fatal errors 26
- formatting, RCL commands xiii
- function strings
 - incorrect variables 7
 - no match found 51

G

- gateway error logs 10, 26
- generation number, increasing 24
- graphic icons xv

H

- hardware problems, identifying 2

I

- icons xv
- identifiers, defined xiv
- image columns
 - LTM shutdown and 112
 - replication problem 22
- inbound queue requires more disk space 135
- informational messages 26
- interfaces file
 - no destination Replication Server entry 67
 - no source Replication Server entry 66
- internal errors 26
- international environments, problems with 155, 158

K

- kernel, rebuilding 4
- key violation errors 123

L

- languages

- configuring 156
 - mixed 156
 - stored procedure 156
- locator
 - database log 132
 - inconsistencies between primary and replicate 106
 - resetting 132
- log files
 - Replication Server 11, 26
 - rs_init 11
- Log Transfer Manager. See LTM
- log, RSSD 47
- log, transaction
 - dumping 42, 47
 - full
 - truncating 42, 47
- login, incorrect when creating route 67
- LTM
 - adding 3
 - common errors 61
 - diagnostic version, running 23
 - error logs 10, 26
 - error message information file 11
 - errors from SQL Server 114
 - locator 132
 - normalization errors 110
 - problems starting 108
 - text or image columns 112
 - truncation point, setting 133
- ltm error log 12
- ltm.diag, using 23

M

- maintenance user name, finding for database 122
- matching function string, cannot find 51
- materialization
 - atomic 82, 83
 - bulk 92
 - description of 79
 - methods 79
 - nonatomic 87, 92
 - schema inconsistency 81
- memory_limit parameter, increasing 45
- message language for stored procedures 156

multibyte data 158

N

network error logs 26
network problems
 identifying 2
 LTM connection 110
non-ASCII characters replaced by ? 158
nonatomic materialization 87
non-binary sort orders 157
normalization errors 110

O

objects, replicated, dropping 114
Open Client/Server error messages 29
operating system problems
 identifying 2
 on Windows NT 2
 upgrading 4
origin_qid, incorrect 24
orphaned transactions 134

P

password, finding for database 122
performance problems, critical 3
permissions
 missing or incorrect 52, 80
 primary subscribe required for subscriptions 52
phantom connections 110
problems
 configuration 6
 hardware 2
 network 2
 operating system 2
 recovery, manual 9
 replication 7
 subscriptions 6
 types of 5

Q

question marks replacing non-ASCII characters 158

R

RCL commands, formatting xiii
recoverable normalization errors 111
recovery, manual
 generation number, increasing 24
 problems 9
 troubleshooting procedure for 24
rejecting messages after data loss 22
REP AGENT USER threads down 59
RepAgent
 adding 3
 error messages, analyzing 33
 error messages, format 32
 error messages, overview 31
 error messages, severity of 32
 error messages, state 32
 nested stored procedures 57
replicated objects, dropping 114
replication
 definitions not found 50
 determining where it stopped 23
 image column not replicating 22
 problems 7
 text column not replicating 22
replication agents
 adding 3
 error logs, non-Adaptive Server 26
 error logs, role in troubleshooting 10
Replication Server
 adding 3
 common errors 53
 common informational and warning messages 57
 diagnostic version, running 23
 error logs, overview of analyzing 25
 error logs, reading 29
 error logs, role in troubleshooting 10, 11
 error message format 27
 error message information file 10
 errors caused by Adaptive Server errors 26
 errors, types of 26
 incorrect login when creating route 67

- interfaces file, no entry 66, 67
 - intermediate for indirect route 66
 - Replication Server Manager. See RSM
 - Replication Server System Database. See RSSD
 - repserver.diag, using 23
 - resource contention 3
 - resume connection command
 - skip transaction option and 96
 - using 122
 - routes
 - altering 70
 - cleaning up after drop route with nowait 73
 - create route protocol message waiting 68
 - creating system table subscriptions 69
 - creating, after dropped 66
 - creating, problems with 65
 - drop route protocol in queue 72
 - drop route protocol not sent 71
 - drop route protocol waiting for delivery 72
 - dropping 70, 74
 - error messages while creating 66
 - incorrect login at desination 67
 - indirect, creating 66
 - procedure for checking problems 23
 - protocol message not sent 67
 - rs_helproute diagnostic messages 67, 69
 - RSSD inbound queue not processing 68
 - system table subscriptions, creating 69
 - system table subscriptions, dropping 71
 - troubleshooting overview 63
 - waiting for protocol message 68
 - rows in system table, swapped out 53
 - rs_helproute stored procedure
 - diagnostic messages 67, 69
 - protocol message not sent 67
 - rs_init
 - common errors 37
 - error logs 11
 - problems 6
 - rs_lastcommit system table 24
 - rs_marker stored procedure 67
 - rs_oqid system table 22
 - rs_recovery system table
 - recoverable actions 22
 - rs_subcmp
 - error message information file 11
 - mixed languages 156
 - rs_update_lastcommit 56
 - RSI
 - connection errors 103
 - data loss 55
 - locater inconsistencies 106
 - RSM
 - error logs 26
 - error message information file 10
 - problems 2
 - RSSD
 - deadlocks 45
 - inbound stable queue not processing 68
 - log device full 45, 47
 - restarted 48
 - routes, creating 69
 - routes, dropping 71
 - system table rows swapped out 53
 - system tables, checking 22
- ## S
- schema inconsistency 81
 - set autocorrection command, key violations 123
 - severity of RepAgent error messages 32
 - single-byte data 158
 - skipping transactions 40
 - sort orders
 - binary vs. non-binary 157
 - described 157
 - problems with 157
 - SQT thread, down 49
 - stable queues
 - dumping 137
 - dumps, using 23
 - inbound, requires more disk space 135
 - RSSD, not processing 68
 - stack traces on Windows NT 2
 - standard error output. See stderr
 - state of RepAgent error messages 32
 - stderr
 - description 27
 - role in troubleshooting 10
 - stored procedures
 - connection failures 123

Index

- messages in mixed languages 156
 - nested 57
 - subscriptions
 - connect source command, not executed 51
 - dematerialization 95, 100
 - dematerialization of 95
 - maintenance user permission 81
 - missing rs_select function string 82
 - monitoring 76
 - no database connection 50
 - overview 75
 - permissions, missing or incorrect 52, 80
 - primary subscribe required 52
 - problems 6
 - replication definition not found 50
 - RSSD log device full 47
 - status 100
 - update permission and maintenance user 81
 - verifying login account requirements for 80
 - syntax statements, conventions xii
 - sysadmin dump_file command 152
 - sysadmin dump_queue command 139
 - sysadmin purge_all_open command 153
 - sysadmin purge_first_open command 153
 - sysadmin sqt_dump_queue command 150
 - system tables
 - checking 22
 - rows swapped out of 53
 - rs_oqid system table, checking 22
 - rs_recovery system table, checking 22
- ## T
- tasks causing problems 2
 - text column
 - LTM shutdown 112
 - replication problem 22
 - thread-terminated messages 26
 - trace flags, using 23
 - transaction log. See log, transaction
 - transactions
 - asynchronous 123
 - large 4
 - nested 57
 - open 4
 - orphaned 134
 - skipping 40
 - viewing header information of 127
 - truncating
 - databases 130
 - problems caused by 133
 - transaction log 42, 47
 - truncation point
 - checking 130
 - setting 133
 - turning off 131
 - types of problems 5
- ## U
- unique-key violations 123
 - upgrading
 - Adaptive Server 3
 - Replication Server components 3
- ## V
- variables, incorrect for function string 7
- ## W
- warm standby, invalid enable marker 61
 - warnings 26
 - Windows NT operating system problems 2