



Heterogeneous Replication Guide

Replication Server®

12.6

DOCUMENT ID: DC36924-01-1260-01

LAST REVISED: October 2003

Copyright © 1992-2003 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, AccelaTrade, 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 Server IQ, Adaptive Warehouse, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, AvantGo, AvantGo Application Alerts, AvantGo Mobile Delivery, AvantGo Mobile Document Viewer, 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, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, 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, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, Mail Anywhere Studio, MainframeConnect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, My AvantGo, My AvantGo Media Channel, My AvantGo Mobile Marketing, 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 Net Library, PocketBuilder, Pocket PowerBuilder, Power++, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Rapport, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, S-Designor, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILS, smart.partners, smart.parts, smart.script, 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 MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TradeForce, Transact-SQL, Translation Toolkit, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, 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 and XP Server are trademarks of Sybase, Inc. 03/03

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	ix
-----------------------	----

PART 1 INTRODUCTION

CHAPTER 1	Sybase Replication System Overview	3
	Basic replication concepts	3
	Heterogeneous replication	4
	Replication system components	6
	Primary data server	6
	Replication Agent	6
	Replication Server	7
	Database gateway	8
	Replicate data server	8
	Replication Server Manager	9
 CHAPTER 2	 Sybase Replication Tools	 11
	Sybase replication technology	11
	Replication Server	12
	Replication Agents	15
	DirectConnect database gateways	18
	Replication Server Manager	19
 CHAPTER 3	 Heterogeneous Replication Issues	 21
	Overview of heterogeneous replication issues	21
	Heterogeneous replication limitations	23
	Stored procedure replication	24
	Owner-qualified object names	24
	Large object replication	25
	Setting up replicate databases	27
	Subscription materialization	28
	Replication Server rs_dump command	28
	Replication Server rs_marker command	29

Replication Server rs_dumptran command	29
Replication Server rs_subcmp utility	30

PART 2 COMPONENT-SPECIFIC ISSUES

CHAPTER 4	Data Server Issues.....	33
	Primary data servers	33
	DB2 Universal Database primary data servers	34
	Informix Dynamic Server primary data servers	49
	Microsoft SQL Server primary data servers	56
	Oracle primary data servers	63
	Replicate data servers	72
	DB2 Universal Database replicate data servers.....	72
	Informix Dynamic Server replicate data servers.....	83
	Microsoft SQL Server replicate data servers.....	88
	Oracle replicate data servers	94
	Component Integration Services	99
 CHAPTER 5	 Replication Server Issues	 101
	Relationship with other system components.....	101
	Relationship with Replication Agents	103
	Database connections.....	105
	Rep Agent User thread.....	106
	DSI thread	107
	Replication Agent connections	107
	DirectConnect database gateways.....	109
	Maintenance User	112
	Replication Server heterogeneous datatype support	115
	Datatype translation	115
	HDS issues and limitations.....	116
	Emulating rs_init activity for a non-Sybase database	122
	Case sensitivity in publications and subscriptions	125

PART 3 IMPLEMENTING HETEROGENEOUS REPLICATION

CHAPTER 6	Using RSM to Set Up Replication Systems.....	131
	Introduction to Replication Server Manager.....	131
	Setting up a heterogeneous replication system	133
	Replication system configuration options	133
	DB2 Universal Database for OS/390 replication environment.....	136

	Setting up Replication Server	138
	Setting up a replicate database with HDS scripts	141
	Setting up a Replication Agent	141
	Setting up a non-Sybase data server with RSM.....	143
	Creating a database connection.....	144
	Resuming a database connection	146
CHAPTER 7	Replication System Configuration Examples.....	147
	Non-Sybase primary to Adaptive Server replicate	147
	Replication system components.....	147
	Replication system issues	148
	Adaptive Server primary to non-Sybase replicate	148
	Replication system components.....	148
	Replication system issues	149
	Non-Sybase primary to non-Sybase replicate.....	150
	Replication system components.....	150
	Replication system issues	150
	Bidirectional non-Sybase to non-Sybase replication.....	152
	Replication system components.....	152
	Replication system issues	152
CHAPTER 8	Administering Heterogeneous Replication Systems.....	155
	Replication system maintenance	155
	Primary data servers and Replication Agents	155
	Replication Server	156
	Replicate data servers and DirectConnect gateways.....	156
	Replication system monitoring	157
	Replication system monitoring overview	157
	Monitoring non-Sybase heartbeat and latency.....	158
	Troubleshooting non-Sybase heartbeat and latency	160
CHAPTER 9	Troubleshooting Heterogeneous Replication Systems	163
	Troubleshooting overview	163
	Inbound and outbound queue problems	164
	Inbound queue problems.....	164
	Outbound queue problems.....	166
	Replication failure problems.....	168
	Trigger-based logging failures.....	168
	Replicate database is not updated.....	169
	Troubleshooting specific errors	171
	Date information does not include time values.....	171
	Updates to rs_lastcommit fail	171

Expected datatype translations do not occur	172
Replication Agent triggers do not fire	174
Determining log scan position	176
LTL generation and tracing	176

PART 4

APPENDIXES

APPENDIX A	Datatype Translation and Mapping	183
	DB2 datatypes.....	184
	DB2 class-level translations	184
	Replication Server datatype names for DB2	186
	Informix datatypes.....	187
	Informix class-level translations	187
	Replication Server datatype names for Informix	189
	Microsoft SQL Server datatypes	190
	Microsoft SQL Server class-level translations.....	190
	Replication Server datatype names for Microsoft SQL Server	192
	Oracle datatypes	192
	Oracle class-level translations.....	192
	Replication Server datatype names for Oracle.....	195
APPENDIX B	Materialization Issues	197
	Materialization overview	197
	Unloading data from a primary database	199
	Datatype translation issues	199
	Loading data into replicate databases	201
	Atomic bulk materialization	202
	Preparing for materialization	202
	Performing atomic bulk materialization	203
	Nonatomic bulk materialization	205
	Preparing for materialization	205
	Performing nonatomic bulk materialization	206
APPENDIX C	Heterogeneous Database Reconciliation	209
	Sybase rs_subcmp utility	209
	Database comparison application	210
APPENDIX D	Replication with Adaptive Server Anywhere.....	213
	ASA primary data servers	213
	ASA replicate data servers.....	216

Glossary	219
Index	235

About This Book

Sybase® 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 introduces heterogeneous replication concepts, and it addresses the issues peculiar to heterogeneous replication with Sybase replication technology.

In this book, the term *heterogeneous replication* refers to replication between different types of data servers (for example, replicating from an Informix database to a Sybase Adaptive Server® Enterprise database).

This book also addresses the issues involved with replication between non-Sybase data servers of the same type (for example, replicating from one Oracle database to another Oracle database).

Audience

The *Replication Server Heterogeneous Replication Guide* is for anyone who needs to plan, design, implement, or maintain a Sybase replication system with heterogeneous or non-Sybase data servers.

If you are new to Replication Server, refer to the *Replication Server Design Guide* for an introduction to basic data replication concepts and Sybase replication technology.

How to use this book

This book is divided into four parts:

- Part 1, “Introduction,” provides an overview of data replication concepts and Sybase replication technology, and it introduces the issues specific to replication systems with heterogeneous or non-Sybase data servers. It contains the following chapters:
 - Chapter 1, “Sybase Replication System Overview,” introduces replication system concepts, with a focus on heterogeneous replication using Sybase replication technology.
 - Chapter 2, “Sybase Replication Tools,” introduces the Sybase software products that you can use to implement a replication system with heterogeneous or non-Sybase data servers.

-
- Chapter 3, “Heterogeneous Replication Issues,” describes the issues and problems that are peculiar to heterogeneous replication, and that must be addressed in a successful replication system.
 - Part 2, “Component-Specific Issues,” describes the issues you must consider for each component of a heterogeneous replication system, and explains how to address those issues. It contains the following chapters:
 - Chapter 4, “Data Server Issues,” describes the issues and considerations for specific non-Sybase data servers in a Sybase replication system.
 - Chapter 5, “Replication Server Issues,” describes the issues related to using Sybase Replication Server in a replication system with heterogeneous or non-Sybase data servers.
 - Part 3, “Implementing Heterogeneous Replication,” describes how to set up and maintain a replication system using Replication Server with heterogeneous or non-Sybase data servers. It contains the following chapters:
 - Chapter 6, “Using RSM to Set Up Replication Systems,” introduces the Replication Server Manager™ (RSM) and describes how to use it to set up a replication system.
 - Chapter 7, “Replication System Configuration Examples,” describes several configuration options for a Sybase replication system with heterogeneous or non-Sybase data servers, and explains the issues involved with each configuration.
 - Chapter 8, “Administering Heterogeneous Replication Systems,” describes administration tasks for replication systems with heterogeneous or non-Sybase data servers.
 - Chapter 9, “Troubleshooting Heterogeneous Replication Systems,” describes common problems and troubleshooting procedures for replication systems with heterogeneous or non-Sybase data servers.
 - Part 4, “Appendixes,” contains supplemental information in the following appendixes:
 - Appendix A, “Datatype Translation and Mapping,” lists the class-level datatype translations for all non-Sybase data servers supported by Replication Server, and lists Replication Server datatype names for non-Sybase datatypes.

- Appendix B, “Materialization Issues,” describes the materialization issues that you need to consider when implementing a replication system with heterogeneous or non-Sybase data servers.
- Appendix C, “Heterogeneous Database Reconciliation,” describes the issues involved with reconciling data from different databases in a replication system with heterogeneous or non-Sybase data servers.
- Appendix D, “Replication with Adaptive Server Anywhere,” describes primary and replicate data server issues for Adaptive Server Anywhere in a Sybase replication system.

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 Sybase Technical Library CD and the Technical Library 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 Technical Library CD. It is included with your software. To read or print documents on the Getting Started CD you need Adobe Acrobat Reader (downloadable at no charge from the Adobe Web site, using a link provided on the CD).
- The Technical Library CD contains product manuals and is included with your software. The DynaText reader (included on the Technical Library CD) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

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

To access the Technical Library 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 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list and click Go.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

❖ **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. Enter user name and password information, if prompted (for existing Web accounts) or create a new account (a free service).
- 3 Select a product.
- 4 Specify a time frame and click Go.
- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.

Conventions

The following style conventions are used in this manual:

- In a sample screen display, commands that you should enter exactly as shown are in:

`this font`

- In a sample screen display, words that you should replace with the appropriate value for your installation are shown in:

this font

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

/usr/u/sybase

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

bcp

The conventions for syntax statements in this manual are as follows:

Table 1: SQL syntax conventions

Key	Definition
command	Command names, command option names, utility names, utility flags, and other keywords.
<i>variable</i>	Variables, or words that stand for values that you fill in.
{ }	Curly braces indicate that you must choose at least one of the enclosed options. Do not include braces in your option.
[]	Brackets mean that choosing one or more of the enclosed options is optional. Do not include brackets in your option.
()	Parentheses are to be typed as part of the command.
	The vertical bar means you can select only one of the options shown.
,	The comma means you can choose as many of the options shown as you like, separating your choices with commas that you type as part of the command.

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.

Introduction

Chapters in this part introduce some data replication concepts and the Sybase replication technology that supports replication systems with heterogeneous or non-Sybase data servers.

- Chapter 1, “Sybase Replication System Overview,” introduces basic replication system concepts, with a focus on heterogeneous replication using Sybase replication technology.
- Chapter 2, “Sybase Replication Tools,” introduces the Sybase software products that enable you to implement a heterogeneous replication system using Sybase replication technology.
- Chapter 3, “Heterogeneous Replication Issues,” describes the issues and problems that are specific to heterogeneous replication and that must be addressed in a successful heterogeneous replication system.

Sybase Replication System Overview

This chapter introduces some basic, replication system concepts, with a focus on heterogeneous replication using Sybase replication technology.

Topic	Page
Basic replication concepts	3
Heterogeneous replication	4
Replication system components	6

Basic replication concepts


A basic Sybase replication system consists of three components:

- **Primary database** – a database in which original data-changing operations (or transactions) are performed and captured for replication
- **Replication Server** – a Sybase Open Client/Server™ product that receives transactions to be replicated from a primary database, and delivers them to a replicate database
- **Replicate database** – a database that receives replicated transactions from a Replication Server and applies those transactions to its own “copy” of the primary data

If both primary and replicate data servers are Adaptive Server Enterprise, you can implement a replication system with only these three components. Adaptive Server Enterprise includes all the features necessary to support a Sybase replication system, with no additional components other than the Replication Server.

Figure 1-1 illustrates a basic Sybase replication system, showing the flow of data between two Adaptive Servers and a Replication Server.

Figure 1-1: Basic Sybase replication system



If you are not familiar with basic Sybase replication system concepts and Replication Server features, see the first two chapters of the *Replication Server Administration Guide*.

Heterogeneous replication

The term *heterogeneous replication* refers to replicating data-changing operations between heterogeneous data servers (that is, data servers that are supplied by different vendors). For example:

- A replication system in which Adaptive Server Enterprise is either the primary or the replicate data server, and a non-Sybase data server (such as DB2 Universal Database) is the other data server
- A replication system in which the primary and replicate data servers are different non-Sybase data servers (for example, Oracle is the primary data server and DB2 Universal Database is the replicate data server)

Note This book also addresses Sybase replication systems with no Adaptive Server, and just one type of non-Sybase data server (for example, a system in which both the primary and replicate data servers are Oracle). This is not a heterogeneous replication system, because only one type of data server is used, but the issues you must consider are the same as heterogeneous replication issues.

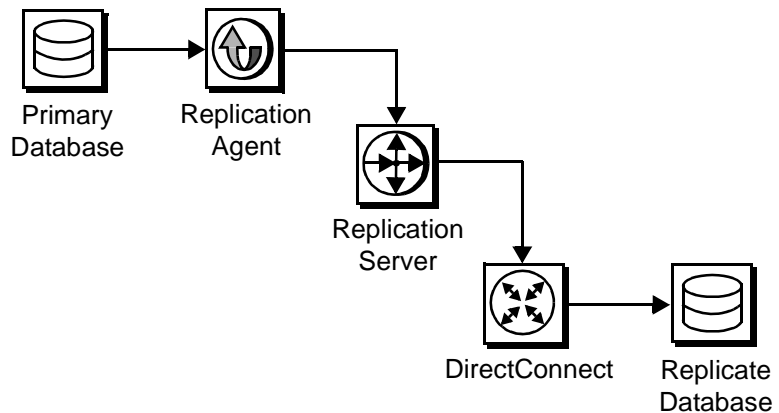
Adaptive Server Enterprise is the data server that Sybase Replication Server was designed to work with. All of the data server elements required to support Replication Server (that is, a data-change capture mechanism in the primary database, and system tables and stored procedures in the replicate database) are either built into Adaptive Server Enterprise, or enabled by utilities that are provided with the Replication Server or Adaptive Server software.

Two additional components are required to implement a Sybase replication system with non-Sybase data servers:

- A Replication Agent™
- A database gateway (such as Sybase DirectConnect™)

Figure 1-2 illustrates a Sybase replication system with non-Sybase data servers, showing the flow of data between the data servers, through the Replication Agent, Replication Server, and DirectConnect database gateway.

Figure 1-2: Sybase replication system with non-Sybase data servers



Replication Agents support non-Sybase data servers by capturing the data-changing operations in the primary database, and sending them to Replication Server for distribution.

DirectConnect database gateways support non-Sybase data servers by providing connectivity between Sybase Open Client/Server and either ODBC, or the native protocol of the data server, and by providing SQL transformation and other services.

Replication Server also includes features and utilities that support non-Sybase data servers. See “Replication Server heterogeneous datatype support” on page 115 for more details about Replication Server support for non-Sybase data servers.

Replication system components

This section describes the major components in a Sybase replication system and the Replication Server Manager™ (RSM) system management tool. The following components are described by their function and role in the replication system:

- Primary data server
- Replication Agent
- Replication Server
- Database gateway
- Replicate data server
- Replication Server Manager

For a more complete description of the Sybase software products (Replication Server, Replication Agents, and DirectConnect gateways), see Chapter 2, “Sybase Replication Tools.”

Primary data server

A **primary data server** is a server that manages one or more primary databases, which are the sources of the data-changing operations or transactions in a replication system.

Non-Sybase primary data servers are supported by Replication Agents, which access replication data in the primary database by logging in to the data server as a database user, or by reading the data server’s transaction log.

For more information about primary data servers, see “Primary data servers” on page 33.

Replication Agent

A **Replication Agent** transfers transaction information, which represents changes made to data, from a primary data server to a Replication Server, for distribution to other (replicate) databases.

For Sybase data servers (such as Adaptive Server Enterprise), an embedded Replication Agent is provided with the database management system software.

The Replication Agent for Adaptive Server Enterprise is called RepAgent, and it is an Adaptive Server thread.

For non-Sybase data servers, Sybase provides two Replication Agent products:

- Replication Agent for DB2 UDB for OS/390 – provides primary data server support for DB2 Universal Database servers that run on OS/390 platforms
- Sybase Replication Agent – provides primary data server support for DB2 Universal Database, Informix, Microsoft SQL Server, and Oracle data servers that run on UNIX or Microsoft Windows platforms

Replication Agents read a transaction log and transfer log records for replicated tables and replicated stored procedures to the Replication Server that manages the primary database.

Replication Agents for DB2 Universal Database (on all platforms) read the DB2 transaction log. Sybase Replication Agent (for Informix, Microsoft SQL Server, and Oracle) uses its own proprietary transaction log, which consists of database objects created in the primary database.

The primary Replication Server reconstructs the transaction and forwards it to replicate sites that have subscriptions for the data.

A Replication Agent is required for each database that contains primary data, or for each database where replicated stored procedures are executed. A replicate database that contains only replicated data, and that has no replicated stored procedures, does not require a Replication Agent.

For more information about Replication Agents, see “Replication Agents” on page 15 and “Primary data servers” on page 33.

Replication Server

A **Replication Server** at each primary or replicate site coordinates data replication activities for local data servers, and exchanges data with Replication Servers at other sites.

Replication Server performs the following major tasks:

- Receives data-changing operations from primary databases through a Replication Agent, and distributes them to replicate database sites that have subscriptions for the data
- Receives data-changing operations from other Replication Servers, and applies them to local replicate databases

- Provides guaranteed delivery of transactions to each replicate site

The information needed to accomplish these tasks is stored in Replication Server system tables. The system tables include descriptions of the replicated data and replication objects, such as replication definitions and subscriptions, security records for Replication Server users, routing information for other Replication Server sites, access methods for local databases, and other administrative information.

Replication Server system tables are stored in a database called the *Replication Server System Database* (RSSD). Each Replication Server has its own RSSD.

For more information about Replication Server, see “Replication Server” on page 7 and Chapter 5, “Replication Server Issues.”

Database gateway

A **database gateway** allows clients using one communication protocol to connect with data servers that use a different protocol.

The Sybase DirectConnect product line consists of database gateway servers that allow clients using the Sybase Open Client/Server protocol (such as Replication Server) to connect with non-Sybase data servers, using either the data server’s native communications protocol or the standard, ODBC protocol.

Sybase DirectConnect products also allow for the retrieval of metadata from non-Sybase replicate data servers.

For more information about DirectConnect database gateways and replicate data servers, see “Replicate data servers” on page 72.

Gatewayless
connections

Sybase also provides the MainframeConnect for DB2 UDB product that allows “gatewayless” connections to DB2 Universal Database data servers on OS/390 platforms. In some environments, a gatewayless connection to the DB2 data server may be more efficient than a gateway server connection.

Replicate data server

A **replicate data server** manages a database that contains replicate data, which is data that is a “copy” of the data in a primary database.

Replication Server maintains the data in a replicate data server by logging in as a database user. In the case of non-Sybase data servers, Replication Server logs in to the replicate data server through a database gateway server.

Replication Server can treat any system as a data server if it supports a set of required data operations and transaction processing directives, either directly (such as Adaptive Server Enterprise) or indirectly (such as a DirectConnect database gateway server).

For more information about replicate data servers and DirectConnect database gateway servers, see “Replicate data servers” on page 72.

Replication Server Manager

Sybase integrates system management tools into a desktop product called Sybase Central, which users can navigate using a central viewer. Each server product, such as Replication Server or Adaptive Server Enterprise, is managed by a service provider “plug-in” that coexists with other service providers in the Sybase Central framework.

Replication Server uses Replication Server Manager (RSM) as a Sybase Central plug-in. This tool provides:

- A graphical user interface
- The ability to manage, monitor, and troubleshoot most replication system components (including primary and replicate data servers, Replication Servers, Replication Agents, and database gateway servers)
- Extensive online help that provides a quick reference for Replication Server concepts and tasks

For more information about Replication Server Manager, see Chapter 6, “Using RSM to Set Up Replication Systems,” and Chapter 8, “Administering Heterogeneous Replication Systems.”

Sybase Replication Tools

This chapter introduces the Sybase software products that enable you to implement a Sybase replication system with heterogeneous or non-Sybase data servers.

Topic	Page
Sybase replication technology	11
Replication Server	12
Replication Agents	15
DirectConnect database gateways	18
Replication Server Manager	19

Sybase replication technology

As the key to developing cost-effective distributed systems that meet critical business requirements, Sybase replication technology solves several major problems that are not addressed by a traditional distributed database architecture:

- Corporate consolidation – Sybase replication technology lets you maintain a corporate overview of distributed operations that is very close to real time, even when distributed business units run on a variety of hardware and DBMS platforms.
- Decentralization – Sybase replication technology allows you to locate data where it is needed, making distributed business units much less vulnerable to central computer or network downtime while reducing overall communication costs. It also allows bidirectional data sharing with a safe approach for replicating remote updates.
- High availability – replication systems based on Sybase replication technology remain robust, despite typical hardware, software, and network failures, delivering applications with very high uptime at a reasonable cost.

Sybase replication products

- Live decision support – replication systems based on Sybase replication technology can replicate an OLTP database, allowing analysts to run complex decision-support queries on data that is within seconds of real time, without affecting OLTP system performance.
- Disaster recovery – Sybase replication technology lets you maintain a near-real-time “warm standby” database, to which applications can switch with virtually no downtime if the primary site fails.

Sybase offers three product lines that specifically support replication systems with heterogeneous or non-Sybase data servers, based on Sybase replication technology:

- Replication Server
- Replication Agent
- DirectConnect

Replication Server is the centerpiece of Sybase’s advanced replication technology, and it incorporates several features specifically to support non-Sybase data servers in a Sybase replication system.

Replication Agents support Replication Server by providing a means to obtain replication data from non-Sybase primary databases. Replication Agents provide this support for DB2 Universal Database, Informix, Microsoft SQL Server, and Oracle data servers.

DirectConnect database gateways support Replication Server by providing access to a variety of non-Sybase databases, allowing them to function as replicate databases in a Sybase replication system.

Replication Server

Replication Server keeps data up to date in multiple databases so that clients can access local data instead of remote, centralized databases. Compared to a centralized data system, a replication system provides improved system performance and data availability and reduces communication overhead. Replication Server provides a cost-effective, fault-tolerant system for replicating data.

Because it transfers transactions rather than rows, Replication Server maintains the integrity of replicated data across the system, while also increasing data

availability. Replication Server also allows you to replicate stored procedure invocations, further enhancing performance.

Replication Server works to distribute data over a network by:

- Providing application developers and system administrators with a flexible publish-and-subscribe model for marking data and stored procedures to be replicated
- Managing replicated transactions while retaining transaction integrity across the network

Because Replication Server replicates transactions—incremental changes instead of data copies—and stored procedure invocations, rather than the operations that result from execution of the stored procedures, it enables a high-performance distributed data environment while maintaining data integrity.

How Replication Server works

A Replication Server at each primary or replicate site coordinates the data replication activities for the local data servers and exchanges data with Replication Servers at other sites.

A Replication Server:

- Receives transactions from primary databases by way of Replication Agents, and distributes them to sites with subscriptions for the data
- Receives transactions from other Replication Servers and applies them to local databases

Replication Server system tables store the information needed to accomplish these tasks. The system tables include descriptions of the replicated data and replication objects, such as replication definitions and subscriptions, security records for Replication Server users, routing information for other sites, access methods for local databases, and other administrative information.

Replication Server system tables are stored in a database called the Replication Server System Database (RSSD).

You use Replication Command Language (RCL) or the Replication Server Manager (RSM) component of Sybase Central to manage information in Replication Server. RCL commands, which resemble SQL commands, can be executed on Replication Server using `isql`, the Sybase interactive SQL utility. The Replication Server *Reference Manual* is the complete reference for RCL. You can find information about Replication Server Manager in the Replication Server *Administration Guide* and in Replication Server online help.

Publish-and-subscribe model

Transactions that occur in a primary database are detected by a Replication Agent and transferred to the local Replication Server, which distributes the information across a network to Replication Servers at destination sites. These Replication Servers in turn update the replicate database according to the requirements of the remote client.

The primary data is the source of the data that Replication Server replicates in other databases. You publish data at primary sites, to which Replication Servers at other (replicate) sites subscribe. You first create a *replication definition* to designate the location of the primary data. The replication definition describes the structure of the table and names the database that contains the primary copy of the table. For easier management, you can collect replication definitions into *publications*.

Creating a replication definition or publication does not, by itself, cause Replication Server to replicate data. You must also create a *subscription* against the replication definition (or the publication) to instruct Replication Server to replicate the data in another database. A subscription resembles a SQL select statement. It can include a where clause to specify which rows of a table you want to replicate in the local database, allowing you to replicate only the necessary data.

You can have multiple replication definitions for a primary table. Replicate tables can subscribe to different replication definitions to obtain different views of the data.

Once you have created subscriptions to replication definitions or publications, Replication Server replicates transactions to databases with subscriptions for the data.

Replicated functions

With some data servers, Replication Server lets you replicate stored procedure invocations asynchronously between databases. This can improve performance over normal data replication by encapsulating many changes in a single replicated function. Because they are not associated with table replication definitions, replicated functions can execute stored procedures that may or may not modify data directly.

Note Replication Server does not support stored procedure replication on all types of data servers. For more information about replicating stored procedures on a particular data server, refer to the appropriate Replication Agent documentation.

You can replicate stored procedure invocations from a primary database to a replicate database, or from a replicate database to a primary database.

With replicated functions, you can execute a stored procedure in another database. A replicated function allows you to:

- Replicate the execution of a stored procedure to subscribing sites
- Improve performance by replicating only the name and parameters of the stored procedure rather than the actual database changes

Replication Server supports both applied functions and request functions:

- An *applied function* is replicated from a primary to a replicate database. You create subscriptions at replicate sites for the function replication definition and mark the stored procedure for replication in the primary database.
- A *request function* is replicated from a replicate to a primary database. There is no subscription for a request function. You mark the stored procedure for replication in the replicate database.

Transaction management

Replication Server depends on data servers to provide the transaction processing services needed to protect their stored data. To guarantee the integrity of distributed data, data servers must comply with such transaction-processing conventions as atomicity and consistency.

data servers that store primary data provide most of the concurrency control needed for the distributed database system. If a transaction fails to update a table with primary data, Replication Server does not distribute the transaction to other sites. When a transaction does update primary data, Replication Server distributes the changes and, unless a failure occurs, the update succeeds at all sites that have subscribed to the data.

Replication Agents

A Replication Agent is a Replication Server client that retrieves information from the transaction log for a primary database and formats it for the primary Replication Server.

Replication Agent detects changes to primary data and, using Log Transfer Language (LTL), which is a subset of Replication Control Language (RCL), sends changes in primary data to the primary Replication Server.

A Replication Agent:

- 1 Logs in to the Replication Server.

- 2 Sends a connect source command to identify the session as a log transfer source and to specify the database for which transaction information will be transferred.
- 3 Gets the name of the Maintenance User for the database from the Replication Server.
- 4 Requests the secondary truncation point for the database from the Replication Server.

This returns a value, called the *origin queue ID*, that the Replication Agent uses to find the location in the transaction log where it is to begin transferring transaction operations. The Replication Server has already received operations up to this location.

- 5 Retrieves records from the transaction log, beginning at the record following the secondary truncation point, and formats the information into LTL commands.

Transaction logs

A Replication Agent can use one of two types of transaction logs:

- The native transaction log of the primary database
- A transaction log created and maintained by the Replication Agent

Native transaction logs are used by the RepAgent thread (for Adaptive Server Enterprise) and the Replication Agents for DB2 Universal Database (the Replication Agent for DB2 UDB for OS/390, and Sybase Replication Agent).

A native database transaction log is the ideal source of information about changes to primary data because it is maintained by the database itself and it contains permanent, recoverable updates. If a Replication Agent cannot access the native transaction log of the primary database, the Replication Agent must use another source of reliable information.

Sybase Replication Agent (for Informix, Microsoft SQL Server, and Oracle) uses a transaction log that it creates and maintains in the primary database. The Sybase Replication Agent transaction log consists of database objects (triggers, tables, and procedures) in the primary database.

Sybase Replication Agent generates a SQL script that runs in the primary database. This script creates the Replication Agent transaction log components. See the Sybase Replication Agent *Administration Guide* for more information about the Sybase Replication Agent transaction log.

Replication Agent products

Replication Agent products extend the capabilities of Replication Server by supporting non-Sybase data servers as primary data servers in a Sybase replication system.

Sybase offers the following Replication Agent products for non-Sybase databases:

- Replication Agent for DB2 UDB for OS/390 – provides primary data server support for a DB2 Universal Database server running on OS/390 platforms.
- Sybase Replication Agent – provides primary data server support for DB2 Universal Database, Informix, Microsoft SQL Server, and Oracle data servers running on UNIX and Microsoft Windows platforms.

Replication Agent for DB2 UDB

Replication Agent for DB2 UDB for OS/390 fits into a replication system as follows:

- The primary data server is DB2 Universal Database, which runs as a subsystem in OS/390. The transaction logs are DB2 logs.
- Replication Agent for DB2 UDB for OS/390 provides a log extract, called Sybase Log Extract, that reads DB2 logs and retrieves the relevant DB2 active and archive log entries for the tables marked for replication.
- LTM for MVS receives the data marked for replication from Sybase Log Extract and transfers that data to Replication Server using the TCP/IP communications protocol.

The DB2 data server logs changes to rows in DB2 tables as they occur. The information written to the transaction log includes copies of the data before and after the changes. In DB2, these records are known as “undo” and “redo” records. Control records are written for commits and aborts. These records are translated to commit and rollback operations.

The DB2 log consists of a series of data sets. Sybase Log Extract uses these log data sets to identify DB2 data changes. Since DB2 writes change records to the active log as they occur, Sybase Log Extract can process the log records immediately after they are entered.

Sybase Replication Agent

Sybase Replication Agent is the component that captures transactions in DB2 Universal Database, Informix, Microsoft SQL Server, or Oracle primary databases on UNIX and Microsoft Windows platforms.

Sybase Replication Agent is implemented in the Java programming language. Therefore, you must have a Java Runtime Environment (JRE) installed on the computer that will act as the Replication Agent host machine.

Sybase Replication Agent uses the Java Database Connectivity (JDBC) protocol for all of its communications. It uses a single instance of the Sybase JDBC driver (jConnect™ for JDBC™) to manage all of its connections to Open Client/Server applications, including the primary Replication Server. In

the case of the primary data server, Sybase Replication Agent connects to the JDBC driver for the primary database.

DirectConnect database gateways

This section describes the Sybase DirectConnect products, which are a group of LAN-based middleware servers that provide access to non-Sybase data sources.

The DirectConnect products are Open Server-based software gateways that support DB-Library and CT-Library application program interfaces, and Java Database Connectivity (JDBC) and Open Database Connectivity (ODBC) protocols. They serve as fundamental building blocks for database middleware applications that allow you to access mainframe and LAN-based non-Sybase data sources.

DirectConnect products provide the following middleware services:

- DirectConnect access services that provide access to non-Sybase data sources
- DirectConnect administrative services (DirectConnect Manager) that provide server-side system management

DirectConnect can be used with other Sybase products, such as Replication Server, and the CIS feature of Adaptive Server Enterprise.

The following DirectConnect products support heterogeneous replication systems based on Sybase replication technology:

- DirectConnect Anywhere
- DirectConnect for OS/390
- DirectConnect for AS/400
- DirectConnect for DRDA
- DirectConnect for Informix
- DirectConnect for Microsoft SQL Server
- DirectConnect for Oracle

How DirectConnect works

All Sybase DirectConnect products provide basic connectivity to non-Sybase data services. In particular, they provide access management, copy management, and remote systems management.

Each DirectConnect product consists of a server and one or more service libraries. The server provides the framework in which the service libraries operate. From the server, each access service library accesses data from a particular target database, such as DB2 Universal Database, Informix, Microsoft SQL Server, or Oracle.

Each access service library contains one or more access services that are specific sets of configuration properties. An access service provides the following functionality:

- Transforms SQL
- Supports remote procedure calls (RPCs)
- Transfers data between target databases and other servers, if the servers are accessible through Open Client™
- Supports catalog stored procedures (CSPs) and system stored procedures from Adaptive Server/CIS

The DirectConnect server listens for, validates, and accepts incoming client connections, such as language events or remote procedure calls (RPCs). These events are routed to the target data source (replicate database) through DirectConnect access services, which provide target-specific connectivity features, including datatype conversion, network connectivity, and SQL transformation.

When a client connects to the target database (for example, a Replication Server connecting to a replicate database), the access service logs in using the client user ID, client password, and the ODBC-configured data source name (DSN).

Replication Server Manager

Sybase integrates its system management tools into a single desktop application, Sybase Central, which users navigate using a central viewer. Each server product, such as Replication Server or Adaptive Server Enterprise, is managed by a service provider “plug-in” that coexists with other service providers in the Sybase Central framework.

Replication Server uses Replication Server Manager (RSM) as a Sybase Central plug-in. The RSM interface provides the following features:

- A graphical user interface

- Allows you to manage, monitor, and troubleshoot most replication system components, including primary and replicate data servers
- Contains extensive online help that provides a quick reference for Replication Server concepts and tasks

For more detailed information about RSM features and using RSM, see Chapter 6, “Using RSM to Set Up Replication Systems,” and Chapter 8, “Administering Heterogeneous Replication Systems.”

Heterogeneous Replication Issues

This chapter describes the issues and problems that are specific to heterogeneous replication and that must be addressed in a successful heterogeneous replication system.

Topic	Page
Overview of heterogeneous replication issues	21
Heterogeneous replication limitations	23

Overview of heterogeneous replication issues

The biggest challenge in implementing a successful heterogeneous replication system is accommodating the unique characteristics of data servers that are supplied by different vendors. Regardless of the type (or brand) of a data server, there are issues that are specific to the data server's role in the replication system (that is, primary data server issues and replicate data server issues). And when a single data server acts as both a primary data server and a replicate data server (bidirectional replication), there are still more issues to consider.

Primary database issues

The following primary database issues must be addressed in a successful heterogeneous replication system:

- The requirements of the Replication Agent and the intrusions and impacts of the Replication Agent on the data server. For example, some Replication Agents create and use database objects in the primary database to maintain a transaction log.
- The access and permissions required in the data server for other replication system components. Both the primary Replication Server and the Replication Agent for a database must have user IDs and passwords defined in the database with appropriate permissions to access primary database objects.

- The connectivity required to support communications between the data server and other replication system components. Replication Agents use either the native communication protocol of the data server, ODBC protocols, or JDBC protocols to communicate with the primary database. Replication Server may require a database gateway to communicate with a data server.
- The specific limitations on replication from the particular data server. For example, some Replication Agents restrict the configuration options of some data servers. Replication Server may impose size limitations on some native datatypes in some databases.
- How replication definitions stored in the RSSD are used by the Replication Agent for the particular data server. For example, both Replication Server and Replication Agents are case sensitive in identifying database object names, but some databases are not.
- The datatype conversions that may be required when replicating transactions from the particular data server to another type of data server. For example, almost every type of data server has a unique way of representing temporal data. The `TIMESTAMP` datatype in one database may need to be “translated” to be stored as a `datetime` datatype in another database.
- The replication system management issues specific to the particular data server. For example, different data servers allow different system management options.

For more information about specific primary database issues for specific databases, see “Primary data servers” on page 33.

Replicate database issues

The following replicate database issues must be addressed in a successful heterogeneous replication system:

- The requirements of the DirectConnect database gateway for the particular database server. The DirectConnect Access Services must be configured to work with the replicate database server and Replication Server.
- The access and permissions required in the data server for the replication system to apply transactions to the replicate database. Both the replicate Replication Server and the DirectConnect gateway for a database must have user IDs and passwords defined in the database, with appropriate permissions to access replicate database objects.
- The connectivity required to support communication between the replicate data server and other replication system components. DirectConnect gateways use either the native communication protocol of a data server, or

standard ODBC or JDBC protocols to communicate with a replicate database. Replication Server generally requires a database gateway to communicate with a non-Sybase data server.

Note In the case of DB2 Universal Database on OS/390, Replication Server can use Sybase MainframeConnect for DB2 UDB to connect directly to the mainframe in a “gatewayless” system. This eliminates the need for a database gateway. Replication through a gatewayless connection requires a TCP/IP connection to the mainframe. For more information about DB2 for OS/390 gatewayless connections, see “DB2 Universal Database replicate data servers” on page 72, and the Sybase Open ServerConnect *Installation and Administration Guide*.

- The limitations on replication into the particular data server. For example, Replication Server imposes limitations on some native datatypes in some databases.
- The intrusion and impact of the database objects required to support Replication Server operations. Replication Server requires two tables and may require some stored procedures to manage a replicate database.
- The replication system management issues specific to the particular data server. For example, different data servers allow different system management options.

For more information about specific replicate database issues for specific databases, see “Replicate data servers” on page 72.

Heterogeneous replication limitations

This section describes some limitations of a heterogeneous replication system based on Sybase replication technology:

- Stored procedure replication
- Owner-qualified object names
- Large object replication
- Setting up replicate databases
- Subscription materialization
- Replication Server `rs_dump` command

- Replication Server `rs_marker` command
- Replication Server `rs_dumptran` command
- Replication Server `rs_subcmp` utility

Stored procedure replication

Stored procedure replication allows the execution call of a stored procedure to be replicated, including the parameter values passed as arguments to the primary stored procedure call.

The availability of stored procedure replication depends on the capabilities of the primary and replicate databases, as well as support from the associated Replication Agent and DirectConnect database gateway. Refer to the documentation for the specific Replication Agent and DirectConnect components to determine if stored procedure replication is available for your databases.

Owner-qualified object names

Access to replicate tables and stored procedures in a non-Sybase database often requires that the reference to the replicate table or stored procedure be owner qualified.

For example, suppose the Replication Server Maintenance User defined to apply transactions to an Oracle replicate database is `orauser`. A replicate insert command to table `table1` may fail with a “table not found” error if the owner of `table1` is `bob`. When attempting to find `table1`, Oracle looks for `orauser.table1`, not `bob.table1`. To properly identify the replicate table to be updated, you can:

- Create an alias at the Oracle replicate database that refers to the correct replicate table. For example, you could create a synonym object in Oracle named `table1`, which refers to the fully qualified name of `bob.table1`.
- When creating the replication definition, use the `with replicate table named [table_owner.[table_name]]` clause. To satisfy the example, the clause would be:

```
with replicate table named bob.table1
```

Note If you use RSM, you can also specify the replicate owner and table name to be used for a replication definition.

Owner qualifying with multiple replicate databases

The problem becomes a little more complicated when table1 is to be replicated to more than one replicate database (for example, Oracle replicate table bob.table1 and Informix replicate table sue.table1). The option of using the with replicate table named clause in the replication definition supports only one replicate table name.

To work around this issue, you must create multiple replication definitions, one for each unique replicate table name required. Then, each subscription must refer to the correct replication definition. Each replication definition must use the with replicate table named clause.

Large object replication

Primary database
LOB replication issues

Large object (LOB) datatypes (such as BLOB, CLOB, IMAGE, and TEXT) provide support for the longest streams of character and binary data in a single column. Their size poses unique challenges, both as primary and replicate data.

At the primary database, the impact of LOB datatypes is on the transaction logging function. For Replication Agents that read the transaction log (that is, Replication Agent for DB2 UDB for OS/390 and Sybase Replication Agent for UDB), the log resources must be adequate to support retention of the changes in LOB data, including before images and after images.

For Replication Agents that use a trigger-based method to capture data (that is, Sybase Replication Agent for Informix, Microsoft SQL Server, or Oracle), LOB datatypes are not allowed for trigger or procedure access, which is the normal logging process. The data involved in changes to primary LOB columns is captured separately from other data in a transaction. LOB data is sent to the Replication Server after all the other data in a transaction. Therefore, there may be a temporary loss of transactional integrity as the LOB data is retrieved when the other transaction data is sent to the Replication Server. For example, the LOB data could have changed between the time the transaction was committed and the time the LOB data was selected from the database.

For more information about the specific impacts and concerns for each primary database, see the appropriate Replication Agent documentation.

Replicate database
LOB replication issues

Adaptive Server Enterprise uses a text pointer to identify the location of text and image column data. The text pointer is passed to system functions that perform the actual updates to data in these large columns. The same technique is used internally in Replication Server to apply LOB datatypes. Replication Server obtains a text pointer and system function calls are made to apply the data to replicate databases.

Replication Server default function strings are designed for an Adaptive Server replicate database. Replication Server executes an `rs_textptr_init` or `rs_get_textptr` function string, followed by one or more `rs_writetext` function strings to apply the LOB data to a replicate database. The default function strings supplied in Replication Server for most non-Sybase data servers do not support LOB replication.

When a non-Sybase database is the replicate database, the database gateway used to communicate with the replicate database must be able to emulate the Adaptive Server text pointer processing. Some DirectConnect database gateways do not provide this feature.

DirectConnect for Microsoft SQL Server provides support for LOB replication into Microsoft SQL Server and SQL Server 2000 databases. See “Microsoft SQL Server replicate data servers” on page 88 for more information.

In the case of a replicate database in DB2 for OS/390, you can use the Sybase MainframeConnect for DB2 UDB to provide a “gatewayless” connection to the replicate database, and use modified `rs_get_textptr` and `rs_writetext` function strings to support LOB replication into DB2. See “DB2 Universal Database replicate data servers” on page 72 for more information.

You might be able to remove the dependency on text pointers from the Replication Server (or the DirectConnect database gateway) by modifying the Replication Server text pointer function strings and creating a stored procedure in the replicate database.

Each text (LOB) column to be processed by Replication Server has a unique `rs_writetext` function string created to issue a `writetext` function call to the replicate database. You can modify the `rs_writetext` function string to replace the default `writetext` function call with a remote procedure call (RPC) to a stored procedure you create in the replicate database.

Replicating LOB datatypes to DB2 Universal Database (on UNIX and Windows platforms), Informix, and Oracle requires the RPC method, with a separate stored procedure in the replicate database for each LOB column.

Note The Component Integration Services (CIS) feature of Adaptive Server provides text pointer handling for Oracle LONG and LONG RAW datatypes. See “Oracle replicate data servers” on page 94 for more information.

Replicating LOB datatypes to DB2 for OS/390 requires the `writetext` method, in which the MainframeConnect AMD2 language handler facilitates replication of LOB datatypes to DB2. See the *Replication Server Reference Manual* for more information about the `rs_writetext` function string.

To implement an RPC workaround, you must create one or more stored procedures in the replicate database (depending on the types of primary keys processed), and you must create one Replication Server function string for each text or image (LOB) column to be replicated.

The stored procedure must perform the following processing:

- The stored procedure must be able to identify the row where the text data should be applied. One way to accomplish this is to have the primary key values for the table passed as parameters in the RPC.
- The stored procedure must be able to append text data to the data in the column. Replication Server executes the RPC repeatedly, passing 16,384 bytes (16K) of data at each invocation, until all the data is sent.
- The stored procedure must handle both insert and update operations. During insert operations, the Replication Server inserts all other (non-LOB) column data into a row, leaving the LOB columns null. Then Replication Server invokes the RPC to insert LOB data. During update operations, the Replication Server updates all other (non-LOB) column data, but it does not update the LOB column—any existing data remains. The stored procedure must be able to replace the existing LOB data on an update transaction, and also append data from the multiple RPC invocations.

The DirectConnect gateway for the replicate database must support RPC processing. Refer to the appropriate DirectConnect documentation for more information about RPC handling.

For more information about Replication Server processing of text and image columns and the `rs_writetext` system function, see the *Replication Server Administration Guide* and the *Replication Server Reference Manual*.

Setting up replicate databases

Replication Server provides a utility named `rs_init`, which sets up an Adaptive Server database as a replicate database. This setup creates the Replication Server database connection, creates the required tables and stored procedures in the replicate database, and defines the Replication Server Maintenance User ID.

For non-Sybase databases, no `rs_init` utility is provided. Instead, Sybase provides an example script for each step in setting up a non-Sybase database as a replicate database. These sample scripts are designed for each supported database. You must modify them and apply them manually.

For more information on setting up replicate databases, see “Replicate data servers” on page 72.

Subscription materialization

Materialization is the process of creating and activating subscriptions and copying data from the primary database to the replicate database, thereby initializing the replicate database.

Before you can replicate data from a primary database, you must set up and populate each replicate database so that it is in a state consistent with that of the primary database. There are two types of subscription materialization supported by Replication Server:

- Bulk materialization – manually creating and activating a subscription and populating a replicate database using data unload and load utilities outside the control of the replication system.
- Automatic materialization – creating a subscription and populating a replicate database using Replication Server commands.

Because of the limitations of reading a transaction log (a complete picture of all table data is not available), Replication Agents do not support automatic materialization.

Bulk materialization methods are supported, with varying complexity based on the specific Replication Agent capabilities.

Refer to the Replication Server *Administration Guide* for a general discussion of subscription materialization, and refer to the appropriate Replication Agent documentation for details regarding a particular Replication Agent and its materialization support.

Replication Server rs_dump command

The Replication Server `rs_dump` command is typically used to coordinate database dump activities across a replication system. When a replicate connection receives an `rs_dump` transaction, Replication Server executes the `rs_dump` function string for that connection. You can customize the `rs_dump` function string to execute whatever commands are required.

For heterogeneous replication, some Replication Agents provide a method to invoke the `rs_dump` command from a non-Sybase primary database. Refer to

the appropriate Replication Agent documentation to determine if `rs_dump` execution from the primary database is supported.

For replicate databases, no default function string for `rs_dump` is provided.

For more information about the `rs_dump` command, its use, and function-string modifications, see the Replication Server *Reference Manual*.

Replication Server `rs_marker` command

The Replication Server `rs_marker` command is a primary database transaction log marker mechanism that is used to assist with the bulk materialization process. An `rs_marker` execution is used to pass activate subscription and validate subscription commands to a primary Replication Server. Most Replication Agents support an `rs_marker` invocation to assist with materialization.

For more information about `rs_marker` usage, see the Replication Server *Reference Manual*. For more information about the use and availability of `rs_marker` for a particular database, refer to the appropriate Replication Agent documentation.

Replication Server `rs_dumptran` command

The Replication Server `rs_dumptran` command is typically used to coordinate database transaction dump activities across a replication system. When a replicate connection receives an `rs_dumptran` transaction, the Replication Server executes the `rs_dumptran` function string for that connection. You can customize the `rs_dumptran` function string to execute whatever commands are required.

Support for `rs_dumptran` is not provided for non-Sybase primary databases.

For replicate databases, no default function string for `rs_dumptran` is provided.

For more information about the `rs_dumptran` command, its use, and function-string modifications, see the Replication Server *Reference Manual*.

Replication Server rs_subcmp utility

Replication Server provides an `rs_subcmp` executable program that you can use to compare primary and replicate tables, optionally reconciling any differences found.

There is no equivalent program to support non-Sybase data servers. For non-Sybase database support, your only options are to buy third-party tools that provide such functionality, or build your own application.

For more information about comparing and reconciling databases in a heterogeneous replication system, see Appendix C, “Heterogeneous Database Reconciliation.”

Component-Specific Issues

Chapters in this part describe component-specific issues and considerations peculiar to a replication system with heterogeneous or non-Sybase data servers.

- Chapter 4, “Data Server Issues,” describes the issues and problems that are peculiar to specific data servers in a heterogeneous replication system.
- Chapter 5, “Replication Server Issues,” describes the issues specific to Sybase Replication Server in a heterogeneous replication system.

This chapter describes the issues specific to non-Sybase primary and replicate data servers in a Sybase replication system.

Topic	Page
Primary data servers	33
DB2 Universal Database primary data servers	34
Informix Dynamic Server primary data servers	49
Microsoft SQL Server primary data servers	56
Oracle primary data servers	63
Replicate data servers	72
DB2 Universal Database replicate data servers	72
Informix Dynamic Server replicate data servers	83
Microsoft SQL Server replicate data servers	88
Oracle replicate data servers	94
Component Integration Services	99

Primary data servers

In addition to Adaptive Server Enterprise, Sybase replication technology supports transaction replication from the following relational database servers:

- DB2 Universal Database
- Informix Dynamic Server
- Microsoft SQL Server
- Oracle

To find out about the most current, supported versions of these data servers, refer to the documentation for the Replication Agent that supports a particular non-Sybase data server.

General issues for
non-Sybase primary
data servers

There are several issues related to non-Sybase primary data servers in a Sybase replication system. A successful replication system must address all of the issues for each primary data server.

Non-Sybase primary data server issues include:

- The requirements of the Replication Agent, including any limitations, intrusions, and impacts on the data server's operation
- The access and permissions necessary for the Replication Agent to obtain transactions from the primary database
- The connectivity requirements to support communications between the primary data server and other replication system components
- The limitations imposed on the replication system by the non-Sybase data server
- The datatype conversions (or translations) that may be required to replicate transactions from one type of data server to another
- The replication system management issues specific to the non-Sybase data server

The following sections describe the specific primary data server issues for each supported type of non-Sybase data server.

DB2 Universal Database primary data servers

Sybase provides two different Replication Agent software products to support primary databases in DB2 Universal Database:

- **Replication Agent for DB2 UDB for OS/390** – supports replication from DB2 Universal Database on OS/390 (mainframe) platforms.
- **Sybase Replication Agent** – supports replication from DB2 Universal Database on UNIX and Microsoft Windows platforms.

In the following sections, DB2 Universal Database is treated separately for the two platform types:

- DB2 Universal Database primary data servers on OS/390
- DB2 Universal Database primary data servers on UNIX or Windows

DB2 Universal Database primary data servers on OS/390

This section describes the primary database issues and considerations specific to the DB2 Universal Database server (on an OS/390 platform) in a Sybase replication system.

Replication Agent for DB2 UDB for OS/390

As a primary data server, DB2 interacts with the Replication Agent for DB2 UDB for OS/390 (Replication Agent for DB2 UDB).

The Replication Agent identifies and transfers information about data-changing operations (or transactions) from a DB2 primary data server to a primary Replication Server.

The Replication Agent interacts with the primary Replication Server, and with the RSSD of the primary Replication Server, if so configured.

Note This book describes how to use Replication Agent for DB2 UDB for OS/390 version 12.6. If you have a different version of Replication Agent for DB2 UDB, refer to the appropriate documentation for that product.

DB2 for OS/390 primary database permissions

Two user IDs must be created in DB2 to facilitate its operation as a primary data server:

- LTMADMIN user – a TSO user, optionally named LTMADMIN, to:
 - Install, start, and stop the Replication Agent for DB2 UDB
 - Manage the Replication Agent system tables on DB2

The LTMADMIN user must have ALTER TABLE authority on any DB2 table to be marked for replication. This user ID issues an ALTER TABLE DATA CAPTURE CHANGES command on a primary table that is marked for replication.

The LTMADMIN user must have READ permission on the DB2 log files.

- Replication Server Maintenance User – the user ID specified in the Replication Server create connection command for the primary database.

Any updates applied to the primary database by the Maintenance User are ignored for replication, unless the value of the LTM for MVS LTM_process_maint_uid_trans configuration parameter is Y. Because the Maintenance User ID must be a user ID that can perform updates on DB2, it must conform to OS/390 user ID restrictions (that is, eight characters maximum, all uppercase).

Note To avoid cyclic replication, the Maintenance User ID must not be identical to the value of the LTM for MVS RS_user configuration parameter (the RS user).

Replication intrusions
and impacts in DB2
for OS/390

The performance and operation of DB2 primary data servers in a Sybase replication system might be affected by the following:

- The DB2 transaction log is affected in the following ways:
 - Replication requires a *before* and *after* image of each row that is changed. When you mark a primary table for replication, the table is altered with the DATA CAPTURE CHANGES clause. As the number of tables marked for replication increases, so does the DASD space requirement for the DB2 active log data sets.
 - Using Replication Agent for DB2 UDB increases the amount of data stored in DB2 logs. The size of the increase depends on the number, type, and size of the primary tables, and the types of transactions replicated. For example, update transactions require both *before* and *after* images, and they include all of the columns in a row, even if those columns do not change. See the Replication Agent for DB2 UDB for OS/390 documentation for more detailed information.
- Two Replication Agent system tables are created in the DB2 primary database when the Replication Agent is installed:
 - LTMOBJECTS – contains a row for each primary table marked for replication. Its size depends on the number of tables marked for replication.
 - LTMLASTCOMMIT – stores information about the most recent replicated transaction successfully committed in the replicate database. Its size does not change.

Note Unless the DB2 primary database also serves as a replicate database (in bidirectional replication), you can remove the LTMLASTCOMMIT table from the primary database.

- Each Replication Agent for DB2 UDB started task can process only one DB2 log. Replicating transactions from multiple DB2 subsystems (or from multiple data-sharing members in a data-sharing environment) requires one Replication Agent started task for each DB2 subsystem (or data-sharing member).

- The following DB2 utilities may jeopardize replication integrity if they are executed:
 - LOAD LOG NO
 - RECOVER
 - REORG with RECOVER

Note The DB2 LOAD LOG YES utility is supported in Replication Agent for DB2 UDB.

Primary data server
connectivity

Replication Agent for DB2 UDB requires the following to connect to a DB2 primary data server in an OS/390 environment:

- A valid user ID must be defined to OS/390 and granted permissions to the correct DB2 packages. The Replication Agent for DB2 UDB uses this user ID to log in to DB2. This user ID must have the following permissions:
 - CREATE authority for the Replication Agent for DB2 UDB system tables
 - ALTER TABLE authority on all tables marked for replication
 - READ authority on the DB2 log data sets
- Replication Agent for DB2 UDB jobs must have their JCL modified to execute with the correct accounting, user_id, DB2 logs, and DB2 subsystem libraries identified.

Replication Server
connectivity

Replication Agent for DB2 UDB requires the following to connect to the primary Replication Server:

- To support a TCP/IP connection between the Replication Agent and the primary Replication Server, you must provide the OS/390 system with information about the Sybase servers and TCP/IP host addresses in your replication system. You do this by editing the SYGWHOST macro, located in the XCPHPING member in the *hlq*.JCL data set.

The SYGWHOST macro does the following:

- Serves as a directory containing network addresses and other information that controls how the Replication Agent connects to other servers
- Configures the Open Client component of LTM for MVS to use the MVS TCP/IP protocol for a connection between the Replication Agent and the primary Replication Server

Note You must create a SYGWHOST macro entry for each primary Replication Server that receives transactions from a DB2 subsystem in the OS/390 system. You can define any number of Replication Servers in the SYGWHOST macro.

- The values of the following LTM for MVS configuration parameters must be set as described so that the Replication Agent for DB2 UDB can connect to the primary Replication Server:
 - Communications_Protocol – must be set to IBMTCP for connectivity to the Replication Server.
 - RS – the name of the host machine on which the primary Replication Server resides.
 - RS_user – the user ID that the Replication Agent uses to log in to the primary Replication Server. This user ID must be defined and granted connect source permission in the Replication Server.

Note The RS_user must not be identical to the Maintenance User ID specified in the Replication Server create connection command for the primary database.

- RS_pw – the password for the user ID that the Replication Agent uses to log in to the primary Replication Server.
- RS_source_db – the logical database name of the primary database specified in the Replication Server database connection. Sybase recommends using the DB2 subsystem or group name.

Note By making RS_source_db the actual DB2 database name, you enable the Replication Server Manager (RSM) to assist in creating replication definitions for the primary database.

- RS_source_ds – the MVS system that contains the DB2 primary database. This value can be arbitrary, as long as it matches the *data_server* value specified in the Replication Server create connection command for the primary database.

RSSD connectivity

Replication Agent for DB2 UDB requires the following to connect to the Replication Server System Database (RSSD) of the primary Replication Server:

- If you set the value of the LTM for MVS Use_repdef configuration parameter to Y, you must provide values in the SYGWHOST macro for the data server that contains the RSSD for the primary Replication Server.
- The values of the following LTM for MVS configuration parameters must be set as described so that the Replication Agent can connect to the RSSD of the primary Replication Server:
 - RSSD_server – the name of the host machine for the data server that contains the RSSD.
 - RSSD_database – the database name of the RSSD.
 - RSSD_user – the user ID that the Replication Agent uses to log in to the RSSD. This user ID must be defined and granted select permission in the RSSD.
 - RSSD_pw – the password for the user ID that the Replication Agent uses to log in to the RSSD.

DB2 for OS/390
primary database
configuration issues

The Replication Agent for DB2 UDB is a mainframe OS/390 application. It consist of two tasks that run simultaneously in a single OS/390 address space:

- **Sybase Log Extract** – continuously scans the DB2 active and archive logs for data-changing operations on primary tables.
- **LTM for MVS** – receives replicated transactions from Sybase Log Extract, converts them to Log Transfer Language (LTL), and then sends them to the primary Replication Server.

One Replication Agent for DB2 UDB started task is required for each DB2 log from which transactions will be replicated.

All Replication Agent installation and configuration issues are described in the Replication Agent for DB2 UDB *Installation Guide* for OS/390. The following items, however, may need further attention in a heterogeneous replication system:

- All configuration parameter values in the LTM for MVS configuration file are case sensitive. Be careful when specifying the values of the RS_source_ds and RS_source_db parameters, as the Replication Server is also case sensitive. If the same case is not used in both Replication Agent and Replication Server parameters, the connection does not succeed.
- The LTM for MVS LTM_process_maint_uid_trans configuration parameter controls whether the Replication Agent sends transactions performed by the Maintenance User to the primary Replication Server. The Maintenance

User ID is defined in the Replication Server create connection command for the primary database.

In a bidirectional replication environment (replicating both into and out of the same DB2 region), the value of the LTM_process_maint_uid_trans parameter should be set to N. If it is not, transactions replicated to another site could return to be applied at the originating site, creating an endless loop.

Replication definitions
for primary tables in
DB2 for OS/390

The LTM for MVS Use_repdef configuration parameter controls whether the Replication Agent sends Log Transfer Language (LTL) that contains only the columns specified in a replication definition, or all of the columns in the DB2 primary table.

When the value of the Use_repdef parameter is set to N, the Replication Agent sends LTL with data for all of the columns in the DB2 primary table. When the value of the Use_repdef parameter is set to Y, the Replication Agent sends LTL with data for only the columns specified in the replication definition.

By sending data for only the columns needed for the replication definition, network traffic is reduced, which can improve performance.

If you set the value of Use_repdef to Y, you can use other LTM for MVS parameters, such as Minimal_cols and Suppress_col_names, to enhance Replication Agent performance. See the Replication Agent for DB2 UDB *Installation Guide* for OS/390 for more information.

DB2 for OS/390
primary datatype
translation issues

The LTM for MVS Date_in_char configuration parameter controls whether the Replication Agent sends values in DB2 DATE columns as default CHAR(10) character strings, or converts them to the Sybase datetime format.

Note Setting the Date_in_char parameter to convert DB2 DATE columns to Sybase datetime values will conflict with replication definition map to clauses that refer to Replication Server user-defined datatype definitions (UDDs).

In general, if the value of the Date_in_char parameter is set to N, all of the corresponding datatypes in replication definitions for DATE columns must specify the Sybase datetime datatype (that is, you cannot specify a map to clause in the replication definitions).

If you want to specify a UDD in the replication definition, you must set the value of the Date_in_char parameter to Y, and you cannot specify the Sybase datetime datatype as the source datatype for DATE columns in the replication definition.

The UDD specified in a replication definition must correspond to the default DATE format for the primary data server. The character (char or varchar) datatype specified in a replication definition for a DATE column must have sufficient length to accommodate the column's default display length.

See the Replication Agent for DB2 UDB for OS/390 *User's and Troubleshooting Guide* for a complete description of the Date_in_char parameter. See the Replication Server *Administration Guide* for a complete description of UDDs and their use.

Note If you use any date- or time-related UDDs in a replication definition, Sybase recommends that you set the Date_in_char parameter to Y to send the character representation of those values. This is a general recommendation for all Replication Agent configurations when UDDs are used.

In general, the Replication Agent for DB2 UDB should not perform datatype translations. The only condition under which you should consider having the Replication Agent perform datatype translation is when all of the replicate data servers require the same translation. In that case, from the point of view of saving processing time, it is probably better to perform the translation once at the Replication Agent, rather than at each replicate database DSI.

DB2 for OS/390
system management
issues

Replication Server Manager (RSM) does not manage or monitor the Replication Agent for DB2 UDB for OS/390. RSM can, however, manage a replication system in which DB2 for OS/390 is either the primary or replicate data server.

DB2 Universal Database primary data servers on UNIX or Windows

This section describes the primary database issues and considerations specific to the DB2 Universal Database server (on a UNIX or Microsoft Windows platform) in a Sybase replication system.

Replication Agent for
UDB

As a primary data server, DB2 Universal Database interacts with Sybase Replication Agent. An instance of the Sybase Replication Agent configured for DB2 Universal Database is referred to as a *Replication Agent for UDB*.

The Replication Agent for UDB identifies and transfers information about data-changing operations (or transactions) from a DB2 Universal Database primary data server to a primary Replication Server.

The Replication Agent interacts with the primary Replication Server, and with the RSSD of the primary Replication Server, if so configured.

DB2 Universal Database primary database permissions

The Replication Agent for UDB requires a DB2 Universal Database login that has permission to access data and create new objects in the primary database.

The DB2 login must have the following permissions in the primary database:

- SELECT and ALTER TABLE authority on any table that is to be replicated
- CREATE TABLE, ALTER TABLE, and DROP TABLE for Replication Agent system tables (or transaction log objects)

The DB2 login must also have READ authority on the DB2 log data sets.

Replication intrusions and impacts in DB2 Universal Database

The performance and operation of DB2 Universal Database primary data servers in a Sybase replication system might be affected by the following:

- The DB2 transaction log is affected in the following ways:
 - You must set the RECOVERY MODE parameter of the DB2 primary database to LOGRETAIN before you create any Replication Agent for UDB transaction log objects (by using the pdb_xlog command).
 - Replication requires a *before* and *after* image of each row that is changed. When you mark a primary table for replication, the Replication Agent for UDB sets the table's DATA CAPTURE option to DATA CAPTURE CHANGES. As the number of tables marked for replication increases, so does the space requirement for the DB2 transaction log.
- When you create Replication Agent for UDB transaction log objects using the pdb_xlog command, the following tables are created in the primary database:
 - Transaction log (xlog) system table
 - Marked objects table
 - BLOB columns table
 - Proc-Active table
 - RS_MARKER shadow table
 - RS_DUMP shadow table

Primary data server connectivity

Replication Agent for UDB requires the following to connect to a DB2 Universal Database primary data server:

- If the Replication Agent for UDB is installed on a different host machine from the DB2 Universal Database server, you must install the DB2 Universal Database Administration Client on the Replication Agent host machine.

If the Replication Agent for UDB software is installed on the same host machine as the DB2 Universal Database server, a separate client is not required.

In either case, you must configure an ODBC data source in the DB2 Administration Client, then use the database name and database alias specified for that ODBC data source when you configure connectivity for the Replication Agent.

- To configure the DB2 Universal Database Administration Client for the JDBC version required by Sybase Replication Agent 12.6, you must run the *usejdbc2* script in the *DB2DIR/sql/lib/java12* directory, where *DB2DIR* is the path where you installed the DB2 software.

Note Sybase Replication Agent is a Java program. Some operating systems may require patches to support Java. Refer to the Sybase Replication Agent *Administration Guide* and the Sybase Replication Agent release bulletin for more information.

- If the Replication Agent for UDB is installed on a UNIX machine, you must source the *db2cshrc* file before attempting to start the Replication Agent instance. The *db2cshrc* file is in the *DB2DIR/sql/lib* directory, where *DB2DIR* is the path where you installed the DB2 software.
- The values of the following Replication Agent for UDB configuration parameters must be set as described:
 - *pds_username* – the DB2 user ID that the Replication Agent uses to log in to the primary database. This user ID must be granted permissions as described in “DB2 Universal Database primary database permissions” on page 42.
 - *pds_password* – the password for the user ID that the Replication Agent uses to log in to the primary database.
 - *pds_connection_type* – must be set to UDBODBC for connectivity to a DB2 primary database.

Note The value of the *pds_connection_type* parameter is set automatically when the Replication Agent for UDB instance is created. Refer to the Sybase Replication Agent *Administration Guide* for more information about creating Replication Agent instances.

- *pds_host_name* – the name of the host machine on which the DB2 data server resides.

- `pds_port_number` – the client socket port number where the DB2 data server listens for connections.
- `pds_server_name` – the server name of the DB2 primary data server.
- `pds_database_name` – the name of the primary database on the DB2 data server, from which transactions will be replicated.
- `pds_datasource_name` – the ODBC database alias for the primary database in the DB2 data server.

Note A separate Replication Agent for UDB instance is required for each database from which transactions are replicated.

- A valid user ID must exist on the data server platform, and it must be granted permissions to the primary database. The Replication Agent for UDB uses this user ID to log in to DB2 Universal Database. This user ID must have the following permissions:
 - `SELECT` and `ALTER TABLE` authority on all tables marked for replication
 - `CREATE TABLE`, `ALTER TABLE`, and `DROP TABLE` authority for Replication Agent system tables
 - `READ` authority on the DB2 log data sets

Replication Server
connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent for UDB can connect to the primary Replication Server:

- `rs_host_name` – the name of the host machine on which the primary Replication Server resides.
- `rs_port_number` – the client socket port number where the primary Replication Server listens for connections.
- `rs_username` – the user ID that the Replication Agent uses to log in to the primary Replication Server. This user ID must be defined and granted connect source permission in the Replication Server.
- `rs_password` – the password for the user ID that the Replication Agent uses to log in to the primary Replication Server.
- `rs_source_ds` – the server name of the primary data server specified in the Replication Server database connection.
- `rs_source_db` – the database name of the primary database specified in the Replication Server database connection.

Note Sybase Replication Agent uses TCP/IP and the Sybase JDBC driver (jConnect for JDBC, which is included in Replication Agent installation) to communicate with other Sybase servers. The Replication Agent does not rely on the Sybase interfaces file for connectivity information.

RSSD connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent for UDB can connect to the RSSD of the primary Replication Server:

- `rssd_host_name` – the name of the host machine for the data server that contains the RSSD.
- `rssd_port_number` – the client socket port number where the RSSD data server listens for connections.
- `rssd_username` – the user ID that the Replication Agent uses to log in to the RSSD. This user ID must be defined and granted select permission in the RSSD.
- `rssd_password` – the password for the user ID that the Replication Agent uses to log in to the RSSD.
- `rssd_database_name` – the database name of the RSSD. When logging in to the data server specified in the `rssd_host_name` parameter, the Replication Agent invokes a use command for the database name specified in the `rssd_database_name` parameter.

Primary database limitations in DB2 Universal Database

Sybase Replication Agent does not support stored procedure replication for DB2 Universal Database.

Replication Server Manager has the following limitations with a DB2 Universal Database primary data server:

- Replication Server Manager cannot start or stop a Sybase Replication Agent instance.

See the Sybase Replication Agent *Administration Guide* for more information on starting and stopping a Replication Agent instance.

- Replication Server Manager does not provide a method to invoke the Sybase Replication Agent `pdb_xlog` command. You must log in to the Replication Agent for UDB instance to execute this command.

See the Sybase Replication Agent *Administration Guide* for more information on logging in to a Replication Agent instance and using the `pdb_xlog` command.

DB2 Universal
Database primary
database
configuration issues

All the installation issues and configuration parameter details for a DB2 Universal Database primary data server are provided in the Sybase Replication Agent *Installation Guide*. The following are a few items that may need additional attention:

- When you install Sybase Replication Agent, a Java Runtime Environment (JRE) that is compatible with the Replication Agent for UDB may be installed for you. For each operating system, you should download and install the most recent recommended patches specified by your operating system vendor for Java compatibility.
- Sybase Replication Agent requires a JDBC driver for connectivity to the DB2 Universal Database server. The JDBC driver for DB2 Universal Database is provided with the DB2 Universal Database Administration Client. See the Sybase Replication Agent *Administration Guide* for more information about the JDBC driver for DB2 Universal Database.
- All configuration parameter values in the Replication Agent configuration file are case sensitive. Be careful when specifying the values of the `rs_source_ds` and `rs_source_db` parameters, as Replication Server is also case sensitive. If the same case is not used in both Replication Agent and Replication Server parameters, no connection occurs.
- The Replication Agent `filter_maint_userid` configuration parameter controls whether the Replication Agent forwards transactions performed by the Maintenance User to the primary Replication Server. The Maintenance User name is defined in the Replication Server create connection command for the primary database.

In a bidirectional replication environment (replicating both into and out of the same database), the value of the `filter_maint_userid` parameter should be set to true. If it is not, transactions replicated to another site could return to be applied at the originating site, creating an endless loop.

- The Replication Agent `lcl_character_case` configuration parameter controls the character case in which the Replication Agent sends database object names to the Replication Server.

For example, if a replication definition is created for all tables named `testtab`, the table name sent by the Replication Agent must be `testtab`, or no match occurs. Because Replication Server is case sensitive, a value of `TESTTAB` does not match the value of `testtab`.

If you use Replication Server Manager to create replication definitions, set the value of the Replication Agent `lcl_character_case` configuration parameter to `asis`. This allows the Replication Agent to send database

object names in the same case as they appear in the primary database, which is the way RSM creates replication definitions, using the case returned by the primary data server. If you create replication definitions outside RSM, it may be easier to choose a default case (for example, create all replication definitions in either all uppercase or all lowercase), and change the value of the Replication Agent `ltl_character_case` parameter to match.

See the Sybase Replication Agent *Administration Guide* for more information on the `ltl_character_case` parameter.

- No Open Client interface application (such as `isql`) is provided as part of the Sybase Replication Agent installation. Use the Replication Server Manager, or use an Open Client application provided with another Sybase product (for example, the `isql` installed with Adaptive Server Enterprise or Replication Server).

Replication definitions for primary tables in DB2 Universal Database

The Replication Agent `use_rssd` configuration parameter controls whether the Replication Agent sends Log Transfer Language (LTL) that contains only the columns specified in a replication definition, or all of the columns in the primary table.

When the value of the `use_rssd` parameter is set to `false`, the Replication Agent sends LTL with data for all of the columns in the primary table. When the value of the `use_rssd` parameter is set to `true`, the Replication Agent sends LTL with data for only the columns specified in the replication definition for each primary table.

By sending data for only the columns specified in the replication definition, network traffic is reduced, which can improve performance.

In addition, column names and parameter names are removed from the LTL because the Replication Agent can send information in the order identified by the replication definition. The LTL `minimal columns` and `structured tokens` options are also available when the value of the `use_rssd` parameter is set to `true`. See the Sybase Replication Agent *Administration Guide* for more information.

DB2 Universal Database primary datatype translation issues

The Sybase Replication Agent allows you to control how it sends DB2 DATE, TIME, and TIMESTAMP column values to the Replication Server. There are two options:

- Send the value as a character string (the default)
- Send the value as a Sybase datetime value

The value of the Replication Agent `pdb_convert_datetime` configuration parameter determines how the Replication Agent handles temporal datatypes.

If you set the `pdb_convert_datetime` parameter to `true`, all corresponding datatypes in a replication definition for `DATE`, `TIME`, and `TIMESTAMP` columns are converted to the Sybase `datetime` datatype.

If you set the `pdb_convert_datetime` parameter to `false`, the datatype in a replication definition for a `DATE`, `TIME`, or `TIMESTAMP` column must be either:

- A Replication Server user-defined datatype (UDD), or
- A character datatype

The character (`char` or `varchar`) datatype specified in a replication definition for a `DATE`, `TIME`, or `TIMESTAMP` column must have sufficient length to accommodate the column's default display length.

See the Sybase Replication Agent *Administration Guide* for more information about the `pdb_convert_datetime` parameter and a complete list of datatype mapping for DB2 Universal Database datatypes. See the Replication Server *Administration Guide* for more information about UDDs and their use.

DB2 Universal
Database system
management issues

The Sybase Replication Agent provides a number of commands that return metadata information about the primary database (database names, table names, procedure names, column names, and so forth). It does this by issuing specific JDBC calls designed to return this information, or by querying the system tables directly.

Replication Server Manager (RSM) can be configured to use either the Replication Agent or a DirectConnect database gateway to return information about the primary database. Typically, RSM uses the Replication Agent to obtain information from the primary database, and uses a DirectConnect gateway to obtain information from a replicate database. You can specify which product is used to return metadata information to RSM when you add a non-Sybase data server using the RSM Add Server wizard.

Further considerations of using RSM are:

- RSM cannot create a Sybase Replication Agent instance, nor can it start and stop a Replication Agent instance.
- RSM can fully configure a Replication Agent for UDB instance, as well as invoke most Replication Agent administrative commands.
- RSM can monitor and display the contents of Replication Agent instance system logs.

Other primary database issues for DB2 Universal Database

- For RSM latency calculations to work correctly, the value of the Replication Agent `lrl_origin_time_required` parameter must be true.

The Sybase Replication Agent provides a set of sample scripts that you can use to set up simple, single-table replication from DB2 Universal Database to Adaptive Server. These scripts are located in the `$SYBASE/rax-12_6/scripts` directory of the Sybase Replication Agent installation.

See the Sybase Replication Agent *Administration Guide* for more information about the sample scripts and their use.

Informix Dynamic Server primary data servers

Replication Agent for Informix

This section describes the primary database issues and considerations specific to the Informix Dynamic Server data server in a Sybase replication system.

As a primary data server, Informix interacts with Sybase Replication Agent. The Replication Agent identifies and transfers information about data-changing operations (or transactions) from an Informix primary data server to a primary Replication Server.

The Replication Agent interacts with the primary Replication Server, and with the RSSD of the primary Replication Server, if so configured.

Informix primary database permissions

The Replication Agent requires an Informix login ID that has permission to access data and create new objects in the primary database.

The Informix login ID must have the following permissions in the primary database:

- SELECT on any table to be replicated
- CREATE TRIGGER on any trigger on any table to be replicated
- DROP TRIGGER on any trigger on any table to be replicated
- CREATE TABLE
- DROP TABLE
- CREATE PROCEDURE on any procedure to be replicated
- DROP PROCEDURE on any procedure to be replicated

Replication intrusions and impacts in Informix

The performance and operation of Informix primary data servers in a Sybase replication system might be affected by the following:

- The Sybase Replication Agent uses its own proprietary transaction log to capture and record transactions in the primary database. The Replication

Agent transaction log consists of shadow tables, stored procedures, and triggers that are created as objects in the primary database.

For each table operation or procedure invocation to be replicated, a sequence of triggers and/or stored procedures is executed to copy the operational data or parameters into a copy (shadow) table. The shadow tables are read periodically by the Replication Agent, which sends the transaction data to the primary Replication Server.

The impact of making these copies and scanning them depends on the capacity of the host machine and data server at the time of execution. A server with low transaction volumes or generous system resources may experience very little or no noticeable impact on performance. For a heavily loaded system, or one with limited resources, the impact on performance could be significant. In addition, the need to keep copies of primary transaction data may require significant additional disk space for systems with large transaction volumes.

Primary data server
connectivity

Sybase Replication Agent requires the following to connect to an Informix primary data server:

- A JDBC driver for the Informix data server. You should ask your Informix vendor for the most recent version of the Informix JDBC driver that supports your Informix data server version.

The JDBC driver must be installed and referenced in the *CLASSPATH* environment variable on the Replication Agent host machine. Java uses the contents of the *CLASSPATH* variable to identify search paths for Java classes. For the Informix JDBC driver, the full path and file name must be in the *CLASSPATH* variable, for example: *c:\informix\jdbc\ifxjdbc.jar*

Note Sybase Replication Agent is a Java program. Some operating systems may require patches to support Java. Refer to the Sybase Replication Agent *Administration Guide* and the Sybase Replication Agent release bulletin for more information.

- The values of the following Replication Agent configuration parameters must be set as described:
 - *pds_username* – the Informix user ID that the Replication Agent uses to log in to the primary database. This user ID must be granted permissions as described in “Informix primary database permissions” on page 49.
 - *pds_password* – the password for the Informix user ID that the Replication Agent uses to log in to the Informix primary database.

- `pds_connection_type` – must be set to `IFXJDBC` for connectivity to an Informix primary database.

Note The value of the `pds_connection_type` parameter is set automatically when the Replication Agent instance is created. Refer to the Sybase Replication Agent *Administration Guide* for more information about creating Replication Agent instances.

- `pds_host_name` – the name of the host machine on which the Informix data server resides.
- `pds_port_number` – the client socket port number where the Informix data server listens for connections.
- `pds_server_name` – the server name of the Informix data server.
- `pds_database_name` – the name of the primary database on the Informix data server, from which transactions will be replicated.

Note A separate Replication Agent instance is required for each database from which transactions are replicated.

Replication Server connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent can connect to the primary Replication Server:

- `rs_host_name` – the name of the host machine on which the primary Replication Server resides.
- `rs_port_number` – the client socket port number where the primary Replication Server listens for connections.
- `rs_username` – the user ID that the Replication Agent uses to log in to the primary Replication Server. This user ID must be defined and granted connect source permission in the Replication Server.
- `rs_password` – the password for the user ID that the Replication Agent uses to log in to the primary Replication Server.
- `rs_source_ds` – the server name of the primary data server specified in the Replication Server database connection.
- `rs_source_db` – the database name of the primary database specified in the Replication Server database connection.

Note Sybase Replication Agent uses TCP/IP and the Sybase JDBC driver (jConnect for JDBC, which is included in Replication Agent installation) to communicate with other Sybase servers. The Replication Agent does not rely on the Sybase interfaces file for connectivity information.

RSSD connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent can connect to the RSSD of the primary Replication Server:

- `rssd_host_name` – the name of the host machine for the data server that contains the RSSD.
- `rssd_port_number` – the client socket port number where the RSSD data server listens for connections.
- `rssd_username` – the user ID that the Replication Agent uses to log in to the RSSD. This user ID must be defined and granted `select` permission in the RSSD.
- `rssd_password` – the password for the user ID that the Replication Agent uses to log in to the RSSD.
- `rssd_database_name` – the database name of the RSSD. When logging in to the data server specified in the `rssd_host_name` parameter, the Replication Agent invokes a `use` command for the database name specified in the `rssd_database_name` parameter.

Primary database limitations in Informix

Sybase Replication Agent imposes the following limitations on an Informix primary data server:

- The default transaction isolation of the Informix database must be Read Committed. This ensures that the Replication Agent reads and processes only committed transactions.
- The Informix database must have the Informix Schema installed. The Replication Agent uses system tables created by the schema installation.
- No primary table marked for replication can have a *selective update* trigger on it. (A selective update trigger is an update trigger that acts based only on operations in specific columns.) This type of trigger can prevent the Replication Agent from capturing transaction data on updates.
- The maximum number of columns allowed in a primary table marked for replication is four columns less than the maximum number of columns allowed by the Informix data server instance. The shadow table that stores captured transaction data uses the primary table schema, plus four

additional columns for Replication Agent system information. If four more columns cannot be added to the number of columns in a primary table, the shadow table creation fails.

Replication Server Manager has the following limitations with an Informix primary data server:

- Replication Server Manager cannot start or stop a Sybase Replication Agent instance.

See the Sybase Replication Agent *Administration Guide* for more information on starting and stopping the Replication Agent instance.

- Replication Server Manager does not provide a method to invoke the Replication Agent `pdb_xlog` command. You must log in to the Replication Agent instance to execute this command.

See the Sybase Replication Agent *Administration Guide* for more information on logging in to a Replication Agent instance and using the `pdb_xlog` command.

Informix primary database configuration issues

All the installation issues and configuration parameter details for an Informix primary data server are provided in the Sybase Replication Agent *Installation Guide*. The following are a few items that may need additional attention:

- When you install Sybase Replication Agent, a Java Runtime Environment (JRE) that is compatible with the Replication Agent may be installed for you. For each operating system, you should download and install the most recent recommended patches specified by your operating system vendor for Java compatibility.
- Sybase Replication Agent requires a JDBC driver for connectivity to the primary data server. Sybase does not provide a JDBC driver for Informix data servers. Contact your Informix database vendor for more information about JDBC drivers for Informix data servers.
- All configuration parameter values in the Replication Agent configuration file are case sensitive. Be careful when specifying the values of the `rs_source_ds` and `rs_source_db` parameters, as Replication Server is also case sensitive. If the same case is not used in both Replication Agent and Replication Server parameters, no connection occurs.
- The Replication Agent `filter_maint_userid` configuration parameter controls whether the Replication Agent forwards transactions performed by the Maintenance User to the primary Replication Server. The Maintenance User name is defined in the Replication Server create connection command for the primary database.

In a bidirectional replication environment (that is, replicating both into and out of the same database), the value of the `filter_maint_userid` parameter should be set to true. If it is not, transactions replicated to another site could return to be applied at the originating site, creating an endless loop.

- The Replication Agent `ltl_character_case` configuration parameter controls the character case in which the Replication Agent sends database object names to the Replication Server.

For example, if a replication definition is created for all tables named `testtab`, the table name sent by the Replication Agent must be `testtab`, or no match occurs. Because Replication Server is case sensitive, a value of `TESTTAB` does not match a value of `testtab`.

If you use Replication Server Manager to create replication definitions, set the value of the Replication Agent `ltl_character_case` configuration parameter to `asis`. This allows the Replication Agent to send database object names in the same case as they appear in the primary database, which is the way RSM creates replication definitions, using the case returned by the primary data server. If you create replication definitions outside RSM, it may be easier to choose a default case (for example, create all replication definitions in either all uppercase or all lowercase), and change the value of the Replication Agent `ltl_character_case` parameter to match.

See the Sybase Replication Agent *Administration Guide* for more information on the `ltl_character_case` parameter.

- No Open Client interface application (such as `isql`) is provided as part of the Sybase Replication Agent installation. Use the Replication Server Manager, or use an Open Client application provided with another Sybase product (for example, the `isql` installed with Adaptive Server Enterprise or Replication Server).

Replication definitions for primary tables in Informix

The Replication Agent `use_rssd` configuration parameter controls whether the Replication Agent sends Log Transfer Language (LTL) that contains only the columns specified in a replication definition, or all of the columns in the primary table.

When the value of the `use_rssd` parameter is set to false, the Replication Agent sends LTL with data for all of the columns in the primary table. When the value of the `use_rssd` parameter is set to true, the Replication Agent sends LTL with data for only the columns specified in the replication definition for each primary table.

By sending data for only the columns specified in the replication definition, network traffic is reduced, which can improve performance.

In addition, column names and parameter names are removed from the LTL because the Replication Agent can send information in the order identified by the replication definition. The LTL minimal columns and structured tokens options are also available when the value of the `use_rssd` parameter is set to true. See the Sybase Replication Agent *Administration Guide* for more information.

Informix primary datatype translation issues

The Sybase Replication Agent allows you to control how it sends Informix DATE column values to the Replication Server. There are two options:

- Send the value as a character string (the default)
- Send the value as a Sybase datetime value

The value of the Replication Agent `pdb_convert_datetime` configuration parameter determines how the Replication Agent handles temporal datatypes.

If you set the `pdb_convert_datetime` parameter to true, all corresponding datatypes in a replication definition for DATE columns are converted to the Sybase datetime datatype.

If you set the `pdb_convert_datetime` parameter to false, the datatype in a replication definition for a DATE column must be either:

- A Replication Server user-defined datatype (UDD), or
- A character datatype

The character (char or varchar) datatype specified in a replication definition for a DATE column must have sufficient length to accommodate the column's default display length.

See the Sybase Replication Agent *Administration Guide* for more information about the `pdb_convert_datetime` parameter and a complete list of datatype mapping for Informix datatypes. See the Replication Server *Administration Guide* for more information about UDDs and their use.

Informix systems management issues

The Sybase Replication Agent provides a number of commands that return metadata information about the primary database (database names, table names, procedure names, column names, and so forth). It does this by issuing specific JDBC calls designed to return this information, or by querying the system tables directly.

Replication Server Manager (RSM) can be configured to use either the Replication Agent or a DirectConnect database gateway to return information about the primary database. Typically, RSM uses the Replication Agent to

obtain information from the primary database, and uses a DirectConnect gateway to obtain information from a replicate database. You can specify which product is used to return metadata information to RSM when you add a non-Sybase data server using the RSM Add Server wizard.

Further considerations of using RSM are:

- RSM cannot create a Sybase Replication Agent instance, nor can it start and stop a Replication Agent instance.
- RSM can fully configure a Replication Agent instance, as well as invoke most Replication Agent administrative commands.
- RSM can monitor and display the contents of Replication Agent instance system logs.
- For RSM latency calculations to work correctly, the value of the Replication Agent `lrl_origin_time_required` parameter must be true.

Other primary
database issues for
Informix

The Sybase Replication Agent provides a set of sample scripts that you can use to set up simple, single-table replication from Informix to Adaptive Server. These scripts are located in the `$SYBASE/rax-12_6/scripts` directory of the Sybase Replication Agent installation.

See the Sybase Replication Agent *Administration Guide* for more information about the sample scripts and their use.

Microsoft SQL Server primary data servers

This section describes the primary database issues and considerations specific to the Microsoft SQL Server data server in a Sybase replication system.

Replication Agent for
Microsoft SQL Server

As a primary data server, Microsoft SQL Server interacts with Sybase Replication Agent. The Replication Agent identifies and transfers information about data-changing operations (or transactions) from a Microsoft SQL Server primary data server to a primary Replication Server.

The Replication Agent interacts with the primary Replication Server, and with the RSSD of the primary Replication Server, if so configured.

Microsoft SQL Server
primary database
permissions

The Replication Agent requires a Microsoft SQL Server user login that has permission to access data and create new objects in the primary database.

The Microsoft SQL Server login must have the following permissions in the primary database:

- select on any table to be replicated
- create trigger on any trigger on any table to be replicated
- drop trigger on any trigger on any table to be replicated
- create table
- drop table
- create procedure on any procedure to be replicated
- drop procedure on any procedure to be replicated

Replication intrusions and impacts in Microsoft SQL Server

The performance and operation of Microsoft SQL Server primary data servers in a Sybase replication system might be affected by the following:

- The Sybase Replication Agent uses its own proprietary transaction log to capture and record transactions in the primary database. The Replication Agent transaction log consists of shadow tables, stored procedures, and triggers that are created as objects in the primary database.

For each table operation or procedure invocation to be replicated, a sequence of triggers and/or stored procedures is executed to copy the operational data or parameters into a copy (shadow) table. The shadow tables are read periodically by the Replication Agent, which sends the transaction data to the primary Replication Server.

The impact of making these copies and scanning them depends on the capacity of the host machine and data server at the time of execution. A server with low transaction volumes or generous system resources may experience very little or no noticeable impact on performance. For a heavily loaded system or one with limited resources, the impact on performance could be significant. In addition, the need to keep copies of primary transaction data may require significant additional disk space for systems with large transaction volumes.

- Sybase Replication Agent updates the @@IDENTITY global system variable in Microsoft SQL Server when either of the following occurs:
 - The Replication Agent trigger executes on a marked table.
 - A stored procedure marked for replication is invoked.

If a client application uses the @@IDENTITY column after an update to a replicate object, the value returned is the identity value from a Replication Agent system table, not the value from the client application table.

For example, if table_a has an identity column, an insert into table_a should set the @@IDENTITY global variable to the value set in the table_a

identity column. If table_a is marked for replication, then the Replication Agent trigger executes after the insert into table_a, which inserts a record into a Replication Agent system table. The value of the @@IDENTITY global variable becomes the value of the Replication Agent system table identity column, not the value of the identity column in table_a.

Note A workaround for the Microsoft SQL Server @@IDENTITY global variable issue is described in the Sybase Replication Agent *Administration Guide*.

Primary data server
connectivity

Sybase Replication Agent requires the following to connect to a Microsoft SQL Server primary data server:

- Sybase Replication Agent 12.6 requires a JDBC 2.0-compliant driver for connectivity to a Microsoft SQL Server primary database. Sybase provides the DataDirect JDBC 2.0 driver for Microsoft SQL Server with the Sybase Replication Agent 12.6 software.
- The values of the following Replication Agent configuration parameters must be set as described:
 - pds_username – the Microsoft SQL Server user login that the Replication Agent uses to log in to the primary database. (This user login must be granted permissions as described in “Microsoft SQL Server primary database permissions” on page 56.)
 - pds_password – the password for the user login that the Replication Agent uses to log in to the primary database.
 - pds_connection_type – must be set to MSMERJDBC for connectivity to a Microsoft SQL Server primary database.

Note The value of the pds_connection_type parameter is set automatically when the Replication Agent instance is created. Refer to the Sybase Replication Agent *Administration Guide* for more information about Replication Agent instances.

- pds_host_name – the name of the host machine on which the Microsoft SQL Server data server resides.
- pds_port_number – the client socket port number where the Microsoft SQL Server data server listens for connections.
- pds_server_name – the server name of the Microsoft SQL Server data server.

- `pds_database_name` – the name of the primary database on the Microsoft SQL Server data server, from which transactions will be replicated.

Note A separate Replication Agent instance is required for each database from which transactions are replicated.

Replication Server connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent can connect to the primary Replication Server:

- `rs_host_name` – the name of the host machine on which the primary Replication Server resides.
- `rs_port_number` – the client socket port number where the primary Replication Server listens for connections.
- `rs_username` – the user ID that the Replication Agent uses to log in to the primary Replication Server. This user ID must be defined and granted connect source permission in the Replication Server.
- `rs_password` – the password for the user ID that the Replication Agent uses to log in to the primary Replication Server.
- `rs_source_ds` – the server name of the primary data server specified in the Replication Server database connection.
- `rs_source_db` – the database name of the primary database specified in the Replication Server database connection.

Note Sybase Replication Agent uses TCP/IP and the Sybase JDBC driver (jConnect for JDBC, which is included in Replication Agent installation) to communicate with other Sybase servers. The Replication Agent does not rely on the Sybase interfaces file for connectivity information.

RSSD connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent can connect to the RSSD of the primary Replication Server:

- `rssd_host_name` – the name of the host machine for the data server that contains the RSSD.
- `rssd_port_number` – the client socket port number where the RSSD data server listens for connections.

Primary database
limitations in Microsoft
SQL Server

- `rssd_username` – the user ID that the Replication Agent uses to log in to the RSSD. This user ID must be defined and granted select permission in the RSSD.
- `rssd_password` – the password for the user ID that the Replication Agent uses to log in to the RSSD.
- `rssd_database_name` – the database name of the RSSD. When logging in to the data server specified in the `rssd_host_name` parameter, the Replication Agent invokes a `use` command for the database name specified in the `rssd_database_name` parameter.

The Sybase Replication Agent imposes the following limitations on a Microsoft SQL Server primary data server:

- The maximum number of columns allowed in a primary table marked for replication is three columns less than the maximum number of columns allowed by the Microsoft SQL Server instance. The shadow table that stores captured transaction data uses the primary table schema, plus three additional columns for Replication Agent system information. If three more columns cannot be added to the number of columns in a primary table, the shadow table creation fails.
- The Microsoft SQL Server `@@IDENTITY` global system variable is updated whenever a Replication Agent trigger executes on a table to be replicated, or whenever a stored procedure to be replicated is executed.

See “Replication intrusions and impacts in Microsoft SQL Server” on page 57 for more information about this issue.

Note A workaround for the `@@IDENTITY` global variable issue is described in the Sybase Replication Agent *Administration Guide*.

Replication Server Manager has the following limitations with a Microsoft SQL Server primary data server:

- Replication Server Manager cannot start or stop a Sybase Replication Agent instance.

See the Sybase Replication Agent *Administration Guide* for more information on starting and stopping the Replication Agent instance.

- Replication Server Manager does not provide a method to invoke the Replication Agent `pdb_xlog` command. You must log in to the Replication Agent instance to execute this command.

See the Sybase Replication Agent *Administration Guide* for more information on logging in to a Replication Agent instance and using the `pdb_xlog` command.

Microsoft SQL Server primary database configuration issues

All the installation issues and configuration parameter details for a Microsoft SQL Server primary data server are provided in the Sybase Replication Agent *Installation Guide*. The following are a few items that may need additional attention:

- All configuration parameter values in the Replication Agent configuration file are case sensitive. Be careful when specifying the values of the `rs_source_ds` and `rs_source_db` parameters, as Replication Server is also case sensitive. If the same case is not used in both Replication Agent and Replication Server parameters, no connection occurs.
- The Replication Agent `filter_maint_userid` configuration parameter controls whether the Replication Agent forwards transactions performed by the Maintenance User to the primary Replication Server. The Maintenance User name is defined in the Replication Server `create` connection command for the primary database.

In a bidirectional replication environment (replicating both into and out of the same database), the value of the `filter_maint_userid` parameter should be set to `true`. If it is not, transactions replicated to another site could return to be applied at the originating site, creating an endless loop.

- The Replication Agent `ltl_character_case` configuration parameter controls the character case in which the Replication Agent sends database object names to the primary Replication Server.

For example, if a replication definition is created for all tables named `testtab`, the table name sent by the Replication Agent must be `testtab`, or no match occurs. Because Replication Server is case sensitive, a value of `TESTTAB` does not match a value of `testtab`.

If you use Replication Server Manager to create replication definitions, set the value of the Replication Agent `ltl_character_case` configuration parameter to `asis`. This allows the Replication Agent to send database object names in the same case as they appear in the primary database, which is the way RSM creates replication definitions, using the case returned by the primary data server. If you create replication definitions outside RSM, it may be easier to choose a default case (for example, create all replication definitions in either all uppercase or all lowercase), and change the value of the Replication Agent `ltl_character_case` parameter to match.

In a Microsoft SQL Server database, object names are stored in lowercase by default, if no case was forced when the object was created. That means the Replication Agent sends object names in lowercase to the primary Replication Server, unless configured to do otherwise.

See the Sybase Replication Agent *Administration Guide* for more information on the `ltl_character_case` parameter.

- No Open Client interface application (such as `isql`) is provided as part of the Sybase Replication Agent installation. Use the Replication Server Manager, or use an Open Client application provided with another Sybase product (for example, the `isql` installed with Adaptive Server Enterprise or Replication Server).

Replication definitions for primary tables in Microsoft SQL Server

The Replication Agent `use_rssd` configuration parameter controls whether the Replication Agent sends Log Transfer Language (LTL) that contains only the columns specified in a replication definition, or all of the columns in the primary table.

When the value of the `use_rssd` parameter is set to `false`, the Replication Agent sends LTL with data for all of the columns in the primary table. When the value of the `use_rssd` parameter is set to `true`, the Replication Agent sends LTL with data for only the columns specified in the replication definition for each primary table.

By sending data for only the columns specified in the replication definition, network traffic is reduced, which can improve performance.

In addition, column names and parameter names are removed from the LTL because the Replication Agent can send information in the order identified by the replication definition. The LTL minimal columns and structured tokens options are also available when the value of the `use_rssd` parameter is set to `true`. See the Sybase Replication Agent *Administration Guide* for more information.

Microsoft SQL Server primary datatype translation issues

The Microsoft SQL Server `datetime` datatype is compatible with the Sybase `datetime` datatype, so the Replication Agent `pdb_convert_datetime` parameter is not significant for a Microsoft SQL Server primary data server.

All Microsoft SQL Server datatypes are compatible with the corresponding Adaptive Server datatypes.

Microsoft SQL Server system management issues

The Sybase Replication Agent provides a number of commands that return metadata information about the primary database (database names, table names, procedure names, column names, and so forth). It does this by issuing specific JDBC calls designed to return this information, or by querying the system tables directly.

Replication Server Manager (RSM) can be configured to use either the Replication Agent or a DirectConnect database gateway to return information about the primary database. Typically, RSM uses the Replication Agent to obtain information from the primary database, and uses a DirectConnect gateway to obtain information from a replicate database. You can specify which product is used to return metadata information to RSM when you add a non-Sybase data server using the RSM Add Server wizard.

Further considerations of using RSM are:

- RSM cannot create a Sybase Replication Agent instance, nor can it start and stop a Replication Agent instance.
- RSM can fully configure a Replication Agent instance, as well as invoke most Replication Agent administrative commands.
- RSM can monitor and display the contents of Replication Agent instance system logs.
- For RSM latency calculations to work correctly, the value of the Replication Agent `lrl_origin_time_required` parameter must be true.

Other primary
database issues for
Microsoft SQL Server

The Sybase Replication Agent provides a set of sample scripts that you can use to set up simple, single-table replication from Microsoft SQL Server to Adaptive Server. These scripts can be found in the `$SYBASE/rax-12_6/scripts` directory.

See the Sybase Replication Agent *Administration Guide* for more information about the sample scripts and their use.

Oracle primary data servers

This section describes the primary database issues and considerations specific to the Oracle data server in a Sybase replication system.

Replication Agent for
Oracle

As a primary data server, Oracle interacts with Sybase Replication Agent. The Replication Agent identifies and transfers information about data-changing operations (or transactions) from an Oracle primary data server to a primary Replication Server.

The Replication Agent interacts with the primary Replication Server, and with the RSSD of the primary Replication Server, if so configured.

Note Sybase Replication Agent is a Java program and it requires a Java Runtime Environment (JRE) to run. Some operating systems require system patches to properly support Java. See the Sybase Replication Agent release bulletin for more information on your specific operating system version and the patches required.

Oracle primary
database permissions

The Replication Agent requires an Oracle login ID that has permission to access data and create new objects in the primary database.

The Oracle login ID must have the following permissions:

- CREATE SESSION
- RESOURCE
- SELECT on any table to be replicated
- CREATE TABLE
- DROP TABLE
- CREATE TRIGGER on any trigger on any table to be replicated
- DROP TRIGGER on any trigger on any table to be replicated
- CREATE PROCEDURE on any procedure to be replicated
- DROP PROCEDURE on any procedure to be replicated

Stored procedures that are created by the Replication Agent require access to two Oracle system views (dynamic performance views) that do not have SELECT authority granted by default. Any user that updates a table or procedure marked for replication must have SELECT authority on the following views:

- V_\$TRANSACTION
- V_\$SESSION

Only the Oracle SYS user (or a user with SYSDBA ROLE) can grant SELECT authority on these views. To simplify implementation, you can grant SELECT authority on these views to PUBLIC.

Replication intrusions
and impacts in Oracle

The performance and operation of Oracle primary data servers in a Sybase replication system might be affected by the following:

- The Sybase Replication Agent uses its own proprietary transaction log to capture and record transactions in the primary database. The Replication

Agent transaction log consists of shadow tables, stored procedures, and triggers that are created as objects in the primary database.

For each table operation or procedure invocation to be replicated, a sequence of triggers and/or stored procedures is executed to copy the operational data or parameters into a copy (shadow) table. The shadow tables are read periodically by the Replication Agent, which sends the transaction data to the primary Replication Server.

The impact of making these copies and scanning them depends on the capacity of the host machine and data server at the time of execution. A server with low transaction volumes or generous system resources may experience very little or no noticeable impact on performance. For a heavily loaded system, or one with limited resources, the impact on performance could be significant. In addition, the need to keep copies of primary transaction data may require significant additional disk space for systems with large transaction volumes.

Primary data server connectivity

Sybase Replication Agent requires the following to connect to an Oracle primary data server:

- A JDBC driver for the Oracle data server. You should ask your Oracle vendor for the most recent version of the JDBC driver that supports your Oracle database version.

Note Sybase Replication Agent version 12.1 or earlier requires a JDBC 1.0-compliant driver. Sybase Replication Agent version 12.5 or later requires a JDBC 2.0-compliant driver.

The JDBC driver must be installed and referenced in the *CLASSPATH* system variable of the Replication Agent host machine. Java uses the contents of the *CLASSPATH* system variable to identify the search locations for Java classes. For the Oracle JDBC driver, the full path and file name must be included in the *CLASSPATH* variable, for example:

c:\oracle\jdbc\classes12.zip

- For JDBC connectivity, the TNS Listener process for the Oracle primary data server must be running.
- The values of the following Replication Agent configuration parameters must be set as described:
 - *pds_username* – the Oracle user ID that the Replication Agent uses to log in to the primary database. This user ID must be granted permissions as described in “Oracle primary database permissions” on page 64.

- `pds_password` – the password for the Oracle user ID that the Replication Agent uses to log in to the primary database.
- `pds_connection_type` – must be set to `ORAJDBC` for connectivity to an Oracle primary database.

Note The value of the `pds_connection_type` parameter is set automatically when the Replication Agent instance is created. Refer to the Sybase Replication Agent *Administration Guide* for more information about Replication Agent instances.

- `pds_host_name` – the name of the host machine on which the Oracle data server resides.
- `pds_port_number` – the client socket port number where the Oracle data server listens for connections.
- `pds_server_name` – the server name of the Oracle data server.
- `pds_database_name` – the name of the Oracle SID, from which transactions will be replicated.

Note A separate Replication Agent instance is required for each Oracle instance from which transactions are replicated.

Replication Server connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent can connect to the primary Replication Server:

- `rs_host_name` – the name of the host machine on which the primary Replication Server resides.
- `rs_port_number` – the client socket port number where the primary Replication Server listens for connections.
- `rs_username` – the user ID that the Replication Agent uses to log in to the primary Replication Server. This user ID must be defined and granted connect source permission in the Replication Server.
- `rs_password` – the password for the user ID that the Replication Agent uses to log in to the primary Replication Server.
- `rs_source_ds` – the server name of the primary data server specified in the Replication Server database connection.
- `rs_source_db` – the database name of the primary database specified in the Replication Server database connection.

Note Sybase Replication Agent uses TCP/IP and the Sybase JDBC driver (jConnect for JDBC, which is included in Replication Agent installation) to communicate with other Sybase servers. The Replication Agent does not rely on the Sybase interfaces file for connectivity information.

RSSD connectivity

The values of the following Replication Agent configuration parameters must be set as described so that the Replication Agent can connect to the RSSD of the primary Replication Server:

- `rssd_host_name` – the name of the host machine for the data server that contains the RSSD.
- `rssd_port_number` – the client socket port number where the RSSD data server listens for connections.
- `rssd_username` – the user ID that the Replication Agent uses to log in to the RSSD. This user ID must be defined and granted select permission in the RSSD.
- `rssd_password` – the password for the user ID that the Replication Agent uses to log in to the RSSD.
- `rssd_database_name` – the database name of the RSSD. When logging in to the data server specified in the `rssd_host_name` parameter, the Replication Agent invokes a use for the database name specified in the `rssd_database_name` parameter.

Primary database limitations in Oracle

Sybase Replication Agent imposes the following limitations on an Oracle primary data server:

- No primary table marked for replication can have a *selective update* trigger on it. (A selective update trigger is an update trigger that acts based only on operations in specific columns.) This type of trigger can prevent the Replication Agent from capturing transaction data on updates.
- Oracle allows multiple AFTER row triggers that fire for the same statement on the same table, however, the order in which these triggers fire is indeterminate. Multiple AFTER row triggers on a table can prevent the Replication Agent from capturing transaction data on updates.
- The maximum number of columns allowed in a primary table marked for replication is three columns less than the maximum number of columns allowed by the Oracle data server instance. The shadow table that stores captured transaction data uses the primary table schema, plus three additional columns for Replication Agent system information. If three

more columns cannot be added to the number of columns in a primary table, the shadow table creation fails.

Replication Server Manager has the following limitations with an Oracle primary data server:

- Replication Server Manager cannot start or stop a Sybase Replication Agent instance.

See the Sybase Replication Agent *Administration Guide* for more information on starting and stopping the Replication Agent instance.

- Replication Server Manager does not provide a method to invoke the Replication Agent `pdb_xlog` command. You must log in to the Replication Agent instance to execute this command.

See the Sybase Replication Agent *Administration Guide* for more information on logging in to the Replication Agent instance and using the `pdb_xlog` command.

Oracle primary database configuration issues

All the installation issues and configuration parameter details for an Oracle primary data server are provided in the Sybase Replication Agent *Installation Guide*. The following are a few items that may need additional attention:

- When you install Sybase Replication Agent, a Java Runtime Environment (JRE) that is compatible with the Replication Agent may be installed for you. For each operating system, you should download and install the most recent recommended patches specified by your operating system vendor for Java compatibility.
- Sybase Replication Agent requires a JDBC driver for connectivity to the primary data server. Sybase does not provide a JDBC driver for Oracle data servers. You must contact your database vendor for more information about JDBC drivers for Oracle data servers.
- All configuration parameter values in the Replication Agent configuration file are case sensitive. Be careful when specifying the values of the `rs_source_ds` and `rs_source_db` parameters, as Replication Server is also case sensitive. If the same case is not used in both Replication Agent and Replication Server parameters, no connection occurs.
- The Replication Agent `filter_maint_userid` configuration parameter controls whether the Replication Agent forwards transactions performed by the Maintenance User to the primary Replication Server. The Maintenance User name is defined in the Replication Server `create connection` command for the primary database.

In a bidirectional replication environment (replicating both into and out of the same database), the value of the `filter_maint_userid` parameter should be set to true. If it is not, transactions replicated to another site could return to be applied at the originating site, creating an endless loop.

- The Replication Agent `ltl_character_case` configuration parameter controls the case in which the Replication Agent sends database object names to the primary Replication Server.

For example, if a replication definition is created for all tables named `testtab`, the table name sent by the Replication Agent must be `testtab`, or no match occurs. Because Replication Server is case sensitive, a value of `TESTTAB` does not match a value of `testtab`.

If you use Replication Server Manager to create replication definitions, set the value of the Replication Agent `ltl_character_case` configuration parameter to `asis`. This allows the Replication Agent to send database object names in the same case as they appear in the primary database, which is the way RSM creates replication definitions, using the case returned by the primary data server. If you create replication definitions outside RSM, it may be easier to choose a default case (for example, create all replication definitions in either all uppercase or all lowercase), and change the value of the Replication Agent `ltl_character_case` parameter to match.

In an Oracle database, object names are stored in all uppercase by default, if no case was forced when the object was created. That means the Replication Agent sends object names in uppercase to the primary Replication Server, unless configured to do otherwise.

See the Sybase Replication Agent *Administration Guide* for more information on the `ltl_character_case` parameter.

- No Open Client interface application (such as `isql`) is provided as part of the Sybase Replication Agent installation. Use the Replication Server Manager, or use an Open Client application provided with another Sybase product (for example, the `isql` installed with Adaptive Server or Replication Server).

Replication definitions
for primary tables in
Oracle

The Replication Agent `use_rssd` configuration parameter controls whether the Replication Agent sends Log Transfer Language (LTL) that contains only the columns specified in a replication definition, or all of the columns in the primary table.

When the value of the `use_rssd` parameter is set to false, the Replication Agent sends LTL with data for all of the columns in the primary table. When the value

of the `use_rssd` parameter is set to `true`, the Replication Agent sends LTL with data for only the columns specified in the replication definition for each primary table.

By sending data for only the columns specified in the replication definition, network traffic is reduced, which can improve performance.

In addition, column names and parameter names are removed from the LTL because the Replication Agent can send information in the order identified by the replication definition. The LTL minimal columns and structured tokens options are also available when the value of the `use_rssd` parameter is set to `true`. See the Sybase Replication Agent *Administration Guide* for more information.

Oracle primary datatype translation issues

The Sybase Replication Agent allows you to control how it sends Oracle DATE column values to the Replication Server. There are two options:

- Send the value as a character string (the default)
- Send the value as a Sybase datetime value

The value of the Replication Agent `pdb_convert_datetime` configuration parameter determines how the Replication Agent handles temporal datatypes.

If you set the `pdb_convert_datetime` parameter to `true`, all corresponding datatypes in a replication definition for DATE columns are converted to the Sybase datetime datatype.

If you set the `pdb_convert_datetime` parameter to `false`, the datatype in a replication definition for a DATE column must be either:

- A Replication Server user-defined datatype (UDD), or
- A character datatype

The character (`char` or `varchar`) datatype specified in a replication definition for a DATE column must have sufficient length to accommodate the column's default display length.

See the Sybase Replication Agent *Administration Guide* for a full description of the `pdb_convert_datetime` parameter and a complete list of datatype mapping for Oracle datatypes. See the Replication Server *Administration Guide* for more information on UDDs and their use.

Oracle system management issues

The Sybase Replication Agent provides a number of commands that return metadata information about the primary database (database names, table names, procedure names, column names, and so forth). It does this by issuing specific JDBC calls designed to return this information, or by querying the Oracle system tables directly.

Replication Server Manager (RSM) can be configured to use either the Replication Agent or a DirectConnect database gateway to return information about a primary database. Typically, RSM uses a Replication Agent to obtain information from a primary database, and uses a DirectConnect gateway to obtain information from a replicate database. You can specify which product is used to return metadata information to RSM when you add a non-Sybase data server using the RSM Add Server wizard.)

Further considerations of using RSM are:

- RSM cannot create a Sybase Replication Agent instance, nor can it start and stop a Replication Agent instance.
- RSM can fully configure a Replication Agent instance, as well as invoke most Replication Agent administrative commands.
- RSM can monitor and display the contents of Replication Agent instance system logs.
- For RSM latency calculations to work correctly, the value of the Replication Agent `lrl_origin_time_required` parameter must be true.

Note Oracle does not support multiple databases within a single server instance as Adaptive Server Enterprise does. Because RSM is designed primarily to work with Adaptive Server, an Oracle database represented in RSM must have a database name as well as a server name. To accommodate this, RSM uses the data server and database name specified in the Replication Server primary database connection to identify the Oracle data server.

Other primary
database issues for
Oracle

The Sybase Replication Agent provides a set of sample scripts which can be used to set up a simple, single-table replication from Oracle to Adaptive Server. These scripts can be found in the `$SYBASE/rax-12_6/scripts` directory of the Sybase Replication Agent installation.

See the Sybase Replication Agent *Administration Guide* for more information about the sample scripts and their use.

Replicate data servers

In addition to Adaptive Server Enterprise, Sybase replication technology supports transaction replication into the following relational database servers:

- DB2 Universal Database
- Informix Dynamic Server
- Microsoft SQL Server
- Oracle

To find out more about the current supported versions of these data servers, refer to the documentation for the DirectConnect database gateway associated with a particular non-Sybase data server.

General issues for
non-Sybase replicate
data servers

There are several issues related to non-Sybase replicate data servers in a Sybase replication system. A successful replication system must address all of the issues for each replicate data server.

Non-Sybase replicate data server issues include:

- The requirements of the DirectConnect database gateway for the non-Sybase data server
- The access and permissions required in the replicate data server for the replication system to apply transactions to the replicate database
- The connectivity requirements for communication between the replicate data server and other components of the replication system
- The limitations on replication into the non-Sybase data server
- The intrusion and impact of the database objects required to support Replication Server operations
- The replication system management issues specific to the non-Sybase data server

DB2 Universal Database replicate data servers

In the following sections describing replicate database issues, the DB2 Universal Database server is treated separately for the OS/390 platform, and for UNIX and Microsoft Windows platforms:

- DB2 Universal Database replicate data servers on OS/390
- DB2 Universal Database replicate data servers on UNIX or Windows

DB2 Universal Database replicate data servers on OS/390

This section describes the replicate database issues and considerations specific to the IBM DB2 Universal Database on the OS/390 platform in a Sybase replication system.

As a replicate data server in a gateway environment, DB2 for OS/390 interacts with the DirectConnect for OS/390 database gateway. The DirectConnect for OS/390 gateway is responsible for accepting commands from the replicate Replication Server and applying those commands to a DB2 replicate database.

As a replicate data server in a “gatewayless” environment, DB2 for OS/390 interacts with MainframeConnect for DB2 UDB through the AMD2 CICS transaction. AMD2 accepts commands from Replication Server and applies those commands to a DB2 database. AMD2 retrieves the results from those commands and returns the results to Replication Server.

Note The gatewayless environment requires a TCP/IP connection to the mainframe. See the Sybase Open ServerConnect *Installation and Administration Guide* for more information about gatewayless connections.

DB2 for OS/390 replicate database permissions

To apply transactions in a replicate database, Replication Server requires a Maintenance User ID that you specify in the Replication Server create connection command. The Maintenance User ID must be defined to the DB2 for OS/390 data server and granted authority to apply transactions in the replicate database. The Maintenance User ID must have the following permissions in the DB2 replicate database:

- CREATE TABLE authority to create tables used for Replication Server processing
- UPDATE authority to all replicate tables and EXECUTE authority to all replicate stored procedures

Replication intrusions and impacts in DB2 for OS/390

The only significant intrusions or impacts to the DB2 replicate database are the database objects created by the *hds_db2_setup_for_replicate.sql* script to support Replication Server replicate database operations.

The *hds_db2_setup_for_replicate.sql* script creates two tables in the replicate database to support Replication Server operations:

- RS_INFO, which contains information about the sort order and character set used by the replicate database.

Note You should confirm that the INSERT statements for this table (in the *hds_db2_setup_for_replicate.sql* script) specify the proper character set and sort order for your data server.

When using Replication Server version 12.0 or later, the replicate database sort order and character set must be recorded in the RS_INFO table.

The Replication Server `rs_get_charset` and `rs_get_sortorder` functions retrieve the character set and sort order from the RS_INFO table in the replicate database.

- RS_LASTCOMMIT, which contains information about replicated transactions applied to the replicate database.

Each row in the RS_LASTCOMMIT table identifies the most recent committed transaction that was distributed to the replicate database from a primary database. Replication Server uses this information to ensure that all transactions are distributed.

The Replication Server `rs_get_lastcommit` function retrieves information about the last transaction committed in the replicate database. For non-Sybase replicate databases, the `rs_get_lastcommit` function is replaced in the database-specific function-string class by the query required to access the RS_LASTCOMMIT table in the replicate database.

Note Replication Server version 12.0 or later must use the RS_LASTCOMMIT table instead of the LTMLASTCOMMIT table that may have been installed if Replication Agent for DB2 UDB for OS/390 was installed.

Replicate database connectivity for DB2 for OS/390

A Replication Server database connection name is made up of two parts: a data server name (`server_name`) and a database name (`db_name`).

When using the DirectConnect for OS/390 database gateway, the `server_name` is the name of the database gateway server, and the `db_name` is the name of the DB2 replicate database.

The replicate Replication Server looks for an interfaces file entry for the database gateway `server_name` specified in the Replication Server database connection. The replicate Replication Server logs in to the replicate data server using the `user_name` and password specified in the database connection.

You must make an entry in the Replication Server interfaces file to identify the host and port where the DirectConnect for OS/390 database gateway server is

Gatewayless connections for DB2 for OS/390

listening. The interfaces file entry name must match the `server_name` portion of the Replication Server database connection.

With a gatewayless connection from Replication Server to DB2 for OS/390 using the MainframeConnect for DB2 UDB, the `server_name` is the mainframe host name, and the `db_name` is the name of the DB2 replicate database. The interfaces file entry for the `server_name` maps to the mainframe IP address and port.

Gatewayless replication to DB2 for OS/390 requires that you:

- Use the Mainframe Connect IPS version 12.5 or later.
- Define the `rs_get_textptr` and `rs_writetext` function strings using the `writetext` method. See the Replication Server *Reference Manual* for more information.
- Run the following Replication Server HDS scripts:
 - `hds_clt_XXX_to_db2.sql` (class-level translations for the primary database)
 - `hds_db2_funcstrings.sql`
 - `hds_db2_udds.sql`

Replicate database limitations in DB2 for OS/390

The following replication limitations exist with a DB2 for OS/390 replicate data server:

- Replication of large object (LOB) datatypes (BLOB, CLOB, and DBCLOB) is supported over a gatewayless connection using MainframeConnect for DB2 UDB version 12.5 or later.
- Replication of large object (LOB) datatypes (BLOB, CLOB, and DBCLOB) is not supported directly by DirectConnect for OS/390.
- Replication Server cannot send a DB2 binary value as a binary string because the DirectConnect for OS/390 database gateway performs an ASCII to EBCDIC translation on the value. Therefore, all binary or varbinary datatypes replicated to DB2 for OS/390 must be mapped to the `rs_db2_char_for_bit` or `rs_db2_varchar_for_bit` datatype.

DB2 for OS/390 replicate database configuration issues

The heterogeneous datatype support (HDS) feature of Replication Server provides a number of sample SQL scripts that help you set up the HDS feature in the replicate Replication Server and the DB2 for OS/390 replicate database. These include:

- `hds_db2_setup_for_replicate.sql` – a script to be applied to the DB2 for OS/390 replicate database.

The *hds_db2_setup_for_replicate.sql* script creates the RS_INFO and RS_LASTCOMMIT tables in the replicate database. The script includes GRANT statements that need to be changed before execution to reference the correct user ID name of the Maintenance User defined in the Replication Server database connection for the replicate database. The Maintenance User ID must have UPDATE authority on the tables.

Note Sybase recommends using the Maintenance User ID to execute this script.

- *hds_db2_funcstrings.sql* and *hds_db2_udds.sql* – scripts to be applied to Replication Server System Database (RSSD).

The *hds_db2_udds.sql* script adds the user-defined datatypes (UDDs) that define the attributes of DB2 native datatypes to the RSSD. The UDDs are required to ensure datatypes received from primary transactions are properly formatted for application to the DB2 for OS/390 replicate database.

Note You may need to modify the *hds_db2_udds.sql* script to reference the correct RSSD database name.

The *hds_db2_funcstrings.sql* script replaces several default Replication Server function strings with custom function strings designed to communicate with DB2 for OS/390 and access the tables and procedures created by the *hds_db2_setup_for_replicate.sql* script. These function strings are added to the Replication Server default rs_db2_function_class.

Note You may need to modify the script to reference the correct RSSD database name.

If you are using a DirectConnect for OS/390 database gateway for replication to DB2 for OS/390, you must set the following properties in the DirectConnect *db2.cfg* Access Service configuration file:

```
SQLTransformation=passthrough
TransactionMode=long
```

Note If you use function strings from the rs_db2_function_class in a Replication Server version prior to 12.0, do *not* run the *hds_db2_funcstrings.sql* script, as it will overwrite the function strings you are currently using.

- Class-level datatype translation scripts to be applied to RSSD.

Class-level translations identify primary datatypes and the replicate datatypes the data should be translated into (for example, Oracle DATE should be translated to DB2 TIMESTAMP).

Note These translations can affect Replication Server performance. Only the translations needed for your specific primary database and replicate database should be applied to the RSSD.

Class-level translation scripts supplied for the DB2 for OS/390 replicate database are:

- *hds_clt_ase_to_db2.sql* – translates Adaptive Server datatypes to DB2 datatypes.
- *hds_clt_odb_to_db2.sql* – translates DB2 Universal Database (for UNIX and Windows) datatypes to DB2 for OS/390 datatypes.
- *hds_clt_informix_to_db2.sql* – translates Informix datatypes to DB2 datatypes.
- *hds_clt_msss_to_db2.sql* – translates Microsoft SQL Server datatypes to DB2 datatypes.
- *hds_clt_oracle_to_db2.sql* – translates Oracle datatypes to DB2 datatypes.
- *hds_db2_connection_sample.sql* – a script to be applied to the replicate Replication Server.

The *hds_db2_connection_sample.sql* script provides a template for creating the Replication Server database connection for a DB2 for OS/390 replicate database using the pre-defined DB2 function-string class provided with Replication Server.

You must modify the *hds_db2_connection_sample.sql* template to include your actual server, database, and Maintenance User names before executing the script.

Note You need not execute the *hds_db2_connection_sample.sql* script if you use Replication Server Manager (RSM) to define the Replication Server replicate database connection.

DB2 for OS/390
system management
issues

To replicate large-object (LOB) datatypes, you must create `rs_writetext` and `rs_get_textptr` function strings. These are not included in the Replication Server HDS sample scripts.

See the Replication Server *Reference Manual* for descriptions of the `rs_writetext` and `rs_get_textptr` function strings.

The following system management issues are specific to a DB2 for OS/390 replicate data server:

- When using Replication Server Manager (RSM), the value of the `DirectConnect DateTimeResults` configuration parameter must be set to `datetime` for latency calculations to work correctly.
- The `create connection` command's `dsi_sql_data_style` parameter was used in previous versions of Replication Server to provide some data translations for the DB2 for OS/390 database.

With the introduction of heterogeneous datatype support (HDS) in Replication Server version 12.0, the `create connection` command's `dsi_sql_data_style` parameter is no longer valid. Do *not* use this parameter with Replication Server version 12.0 or later. The default setting should be " " (blank space).

- The Component Integration Services (CIS) feature allows users to access both Adaptive Server databases and non-Sybase databases on different servers, through a local Adaptive Server connection.

CIS allows Replication Server to connect to an Adaptive Server database, using default Adaptive Server function strings and connection properties, while CIS handles the data conversion and manipulation, and all communications with the DirectConnect database gateway.

See “Component Integration Services” on page 99 for a description of the advantages and disadvantages of this approach.

DB2 Universal Database replicate data servers on UNIX or Windows

This section describes the replicate database issues and considerations specific to the IBM DB2 Universal Database server on a UNIX or Microsoft Windows platform.

As a replicate data server in a replication system, DB2 Universal Database interacts with the DirectConnect for DRDA database gateway. DirectConnect for DRDA is responsible for accepting commands from the replicate Replication Server, and applying those commands to a database residing in a DB2 Universal Database server.

DB2 Universal Database replicate database permissions

To apply transactions in a replicate database, Replication Server requires a Maintenance User ID that you specify in the Replication Server create connection command. The Maintenance User ID must be defined at the DB2 Universal Database server and granted authority to apply transactions in the replicate database. The Maintenance User ID must have the following permissions in the DB2 Universal Database replicate database:

- CREATE TABLE authority to create tables used for Replication Server processing.
- UPDATE authority on all replicate tables.

Replication intrusions and impacts in DB2 Universal Database

The only significant intrusions or impacts to the DB2 Universal Database replicate database are the database objects created by the *hds_udb_setup_for_replicate.sql* script to support Replication Server replicate database operations.

The *hds_udb_setup_for_replicate.sql* script creates two tables in the replicate database to support Replication Server operations:

- RS_INFO, which contains information about the sort order and character set used by the replicate database.

Note You should confirm that the INSERT statements for this table (in the *hds_udb_setup_for_replicate.sql* script) specify the proper character set and sort order for your DB2 Universal Database server.

When using Replication Server version 12.0 or later, the replicate database sort order and character set must be recorded in the RS_INFO table.

The Replication Server *rs_get_charset* and *rs_get_sortorder* functions retrieve the character set and sort order from the RS_INFO table in the replicate database.

- RS_LASTCOMMIT, which contains information about replicated transactions applied to the replicate database.

Each row in the RS_LASTCOMMIT table identifies the most recent committed transaction that was distributed to the replicate database from a primary database. Replication Server uses this information to ensure that all transactions are distributed.

- The Replication Server *rs_get_lastcommit* function retrieves information about the last transaction committed in the replicate database. For non-Sybase replicate databases, the *rs_get_lastcommit* function is replaced in the database-specific function-string class by the query required to access the RS_LASTCOMMIT table in the replicate database.

Replicate database connectivity for DB2 Universal Database

A Replication Server database connection name is made up of two parts: a data server name (*server_name*) and a database name (*db_name*). The *server_name* is the name of the DirectConnect for DRDA database gateway server, and the *db_name* is the name of the replicate database residing in DB2 Universal Database.

The replicate Replication Server looks for an interfaces file entry for the database gateway *server_name* specified in the Replication Server database connection. The replicate Replication Server logs in to the replicate data server using the *user_name* and password specified in the database connection.

You must make an entry in the Replication Server interfaces file to identify the host and port where the DirectConnect for DRDA database gateway server is listening. The interfaces file entry name must match the *server_name* portion of the Replication Server database connection.

Replicate database limitations in DB2 Universal Database

Replication of large object (LOB) datatypes (BLOB, CLOB, and LVARCHAR) is not supported directly from Replication Server to DirectConnect for DRDA.

Note For information about a possible workaround for this LOB data replication limitation, see “Large object replication” on page 25.

DB2 Universal Database replicate database configuration issues

The heterogeneous datatype support (HDS) feature of Replication Server provides a number of sample SQL scripts that help you set up the HDS feature in the replicate Replication Server and the DB2 Universal Database replicate database. These include:

- *hds_udb_setup_for_replicate.sql* – a script to be applied to the DB2 Universal Database replicate database.

The *hds_udb_setup_for_replicate.sql* script creates the RS_INFO and RS_LASTCOMMIT tables in the replicate database. The script includes GRANT statements that need to be changed before execution to reference the correct user ID name of the Maintenance User defined in the Replication Server database connection for the replicate database. The Maintenance User ID must have UPDATE authority on the tables.

Note Sybase recommends using the Maintenance User ID to execute this script.

- *hds_udb_funcstrings.sql* and *hds_udb_udds.sql* – scripts to be applied to Replication Server System Database (RSSD).

The *hds_udb_udds.sql* script adds the user-defined datatypes (UDDs) that define the attributes of DB2 Universal Database native datatypes to the RSSD. The UDDs are required to ensure datatypes received from primary transactions are properly formatted for application to the DB2 Universal Database replicate database.

Note You might need to modify the *hds_udb_udds.sql* script to reference the correct RSSD database name.

The *hds_udb_funcstrings.sql* script replaces several default Replication Server function strings with custom function strings designed to communicate with DB2 Universal Database and access the tables and procedures created by the *hds_udb_setup_for_replicate.sql* script. These function strings are added to the Replication Server default `rs_udb_function_class`.

Note You might need to modify the script to reference the correct RSSD database name.

- Class-level datatype translation scripts to be applied to RSSD.

Class-level translations identify primary datatypes and the replicate datatypes the data should be translated into (for example, Informix BYTE should be translated to DB2 Universal Database CHAR FOR BIT DATA).

Note These translations can affect Replication Server performance. Only the translations needed for your specific primary database and replicate database should be applied to the RSSD.

Class-level translation scripts supplied for the DB2 Universal Database replicate database are:

- *hds_clt_ase_to_udb.sql* – translates Adaptive Server datatypes to DB2 Universal Database datatypes.
- *hds_clt_db2_to_udb.sql* – translates DB2 for OS/390 datatypes to DB2 Universal Database (for UNIX and Windows) datatypes.
- *hds_clt_informix_to_udb.sql* – translates Informix datatypes to DB2 Universal Database datatypes.
- *hds_clt_msss_to_udb.sql* – translates Microsoft SQL Server datatypes to DB2 Universal Database datatypes.

- *hds_clt_oracle_to_udb.sql* – translates Oracle datatypes to DB2 Universal Database datatypes.
- *hds_udb_connection_sample.sql* – a script to be applied to the replicate Replication Server.

The *hds_udb_connection_sample.sql* script provides a template for creating the Replication Server database connection for a DB2 Universal Database replicate database using the pre-defined DB2 Universal Database function-string class provided with Replication Server.

You must modify the *hds_udb_connection_sample.sql* template to include your actual server, database, and Maintenance User names before executing the script.

Note You need not execute the *hds_udb_connection_sample.sql* script if you use Replication Server Manager (RSM) to define the Replication Server replicate database connection.

DB2 Universal Database system management issues

The following system management issues are specific to a DB2 Universal Database replicate data server:

- When using Replication Server Manager (RSM), the value of the DirectConnect DateTimeResults configuration parameter must be set to datetime for latency calculations to work correctly.
- The create connection command's *dsi_sql_data_style* parameter was used in previous versions of Replication Server to provide some data translations for the replicate database.

With the introduction of the heterogeneous datatype support (HDS) feature in Replication Server version 12.0, the create connection command's *dsi_sql_data_style* parameter is no longer valid. Do *not* use this parameter with Replication Server version 12.0 or later. The default setting should be " " (blank space).

- The Component Integration Services (CIS) feature allows users to access both Adaptive Server databases and non-Sybase databases on different servers, through a local Adaptive Server connection.

CIS allows Replication Server to connect to an Adaptive Server database, using default Adaptive Server function strings and connection properties, while CIS handles the data conversion and manipulation, and all communications with the DirectConnect database gateway.

See “Component Integration Services” on page 99 for a description of the advantages and disadvantages of this approach.

Informix Dynamic Server replicate data servers

This section describes the replicate database issues and considerations specific to the Informix Dynamic Server data server in a Sybase replication system.

As a replicate data server, Informix interacts with the DirectConnect for Informix database gateway. The DirectConnect for Informix is responsible for accepting commands from the replicate Replication Server, and applying those commands to an Informix database.

Informix replicate database permissions

To apply transactions in a replicate database, Replication Server requires a Maintenance User ID that you specify in the Replication Server create connection command. The Maintenance User ID must be defined at the Informix data server and granted authority to apply transactions in the replicate database. The Maintenance User ID must have the following permissions in the Informix replicate database:

- create table authority to create tables used for Replication Server processing.
- update authority on all replicate tables and execute authority on all replicate stored procedures.

Replication intrusions and impacts in Informix

The only significant intrusions or impacts to the Informix replicate database are the database objects created by the *hds_informix_setup_for_replicate.sql* script to support Replication Server replicate database operations.

The *hds_informix_setup_for_replicate.sql* script creates two tables in the replicate database to support Replication Server operations:

- *rs_info*, which contains information about the sort order and character set used by the replicate database.

Note You should confirm that the insert statements for this table (in the *hds_informix_setup_for_replicate.sql* script) specify the proper character set and sort order for your Informix data server.

When using Replication Server version 12.0 or later, the replicate database sort order and character set must be recorded in the *rs_info* table.

The Replication Server *rs_get_charset* and *rs_get_sortorder* functions retrieve the character set and sort order from the *rs_info* table in the replicate database.

- *rs_lastcommit*, which contains information about replicated transactions applied to the replicate database.

Each row in the `rs_lastcommit` table identifies the most recent committed transaction that was distributed to the replicate database from a primary database. Replication Server uses this information to ensure that all transactions are distributed.

The Replication Server `rs_get_lastcommit` function retrieves information about the last transaction committed in the replicate database. For non-Sybase replicate databases, the `rs_get_lastcommit` function is replaced in the database-specific function-string class by the query required to access the `rs_lastcommit` table in the replicate database.

Replicate database connectivity for Informix

A Replication Server database connection name is made up of two parts: a data server name (`server_name`) and a database name (`db_name`). The `server_name` is the name of the DirectConnect for Informix database gateway server, and the `db_name` is the name of the Informix replicate database.

The replicate Replication Server looks for an interfaces file entry for the database gateway `server_name` specified in the Replication Server database connection. The replicate Replication Server logs in to the replicate data server using the `user_name` and password specified in the database connection.

You must make an entry in the Replication Server interfaces file to identify the host and port where the DirectConnect for Informix database gateway server is listening. The interfaces file entry name must match the `server_name` portion of the Replication Server database connection.

Replicate database limitations in Informix

The following replication limitations exist with an Informix replicate data server:

- Replication of large object (LOB) datatypes (BLOB, CLOB, and LVARCHAR) is not supported directly from Replication Server to DirectConnect for Informix.

Note For information about a possible work-around for this LOB data replication limitation, see “Large object replication” on page 25.

- The Informix INT datatype allows values ranging from -2147483647 to 2147483647 and the Informix SMALLINT datatype allows values ranging from -32767 to 32767. Most data servers support integer datatypes with values ranging from -2147483648 to 2147483647 and small integer datatypes with values ranging from -32768 to 32767. Replication Server does not provide user-defined datatypes for Informix INT and SMALLINT datatypes to restrict the lower boundary.
- Informix supports up to 32 significant digits of precision/scale in exact numeric datatypes (DECIMAL, and NUMERIC), not 38 significant digits as

Sybase Adaptive Server Enterprise does. Replication Server does not provide user-defined datatypes to support the Informix numeric datatype restriction.

- Replication Server does not provide a class-level translation for Adaptive Server to Informix to handle the Adaptive Server bit datatype. If you use class-level translations to handle bit data replicated to an Informix replicate database, Replication Server returns the following error message:

```
E. 2001/12/14 15:19:57. ERROR #1028 DSI EXEC(136(1)
dock723_dev_db.devdb)
- dsigmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
'[VENDORLIB] Vendor Library Error: [[Message
Iteration=1|SQLState=37000|Native Error=
-201|Message=[Visigenic] [ODBC Informix 7.2
Driver] [Informix]-201: A syntax error has
occurred.] <DCA>'
```

The problem is that without a translation, the bit value goes out as 0x00 or 0x01 – syntax that Informix does not understand.

- Using Replication Server user-defined datatypes and default translations, you can not directly insert or update the BYTE binary datatype in Informix. If you need to replicate this type of data to an Informix database, you must capture the binary literal value as a char or varchar datatype, or use the Component Integration Services (CIS) feature of Adaptive Server, rather than a DirectConnect database gateway, to connect the replicate Replication Server to the Informix replicate database.

Informix replicate database configuration issues

The heterogeneous datatype support (HDS) feature of Replication Server provides a number of sample SQL scripts that help you set up the HDS feature in the replicate Replication Server and the Informix replicate database. These include:

- *hds_informix_setup_for_replicate.sql* – a script to be applied to the Informix replicate database.

The *hds_informix_setup_for_replicate.sql* script creates the *rs_info* and *rs_lastcommit* tables in the replicate database. The script includes grant statements that need to be changed before execution to reference the correct user ID name of the Maintenance User defined in the Replication Server database connection for the replicate database. The Maintenance User ID must have update authority on the tables.

Note Sybase recommends using the Maintenance User ID to execute this script.

- *hds_informix_funcstrings.sql* and *hds_informix_udds.sql* – scripts to be applied to Replication Server System Database (RSSD).

The *hds_informix_udds.sql* script adds the user-defined datatypes (UDDs) that define the attributes of Informix native datatypes to the RSSD. The UDDs are required to ensure datatypes received from primary transactions are properly formatted for application to the Informix replicate database.

Note You might need to modify the *hds_informix_udds.sql* script to reference the correct RSSD database name.

The *hds_informix_funcstrings.sql* script replaces several default Replication Server function strings with custom function strings designed to communicate with Informix and access the tables and procedures created by the *hds_informix_setup_for_replicate.sql* script. These function strings are added to the Replication Server default `rs_informix_function_class`.

Note You might need to modify the script to reference the correct RSSD database name.

- Class-level datatype translation scripts to be applied to RSSD.

Class-level translations identify primary datatypes and the replicate datatypes the data should be translated into (for example, DB2 CHAR FOR BIT DATA should be translated to Informix BYTE).

Note These translations can affect Replication Server performance. Only the translations needed for your specific primary database and replicate database should be applied to the RSSD.

Class-level translation scripts supplied for the Informix replicate database are:

- *hds_clt_ase_to_informix.sql* – translates Adaptive Server datatypes to Informix datatypes.
- *hds_clt_db2_to_informix.sql* – translates DB2 for OS/390 datatypes to Informix datatypes.

- *hds_clt_udb_to_informix.sql* – translates DB2 Universal Database (for UNIX and Windows) datatypes to Informix datatypes.
- *hds_clt_msss_to_informix.sql* – translates Microsoft SQL Server datatypes to Informix datatypes.
- *hds_clt_oracle_to_informix.sql* – translates Oracle datatypes to Informix datatypes.
- *hds_informix_connection_sample.sql* – a script to be applied to the replicate Replication Server.

The *hds_informix_connection_sample.sql* script provides a template for creating the Replication Server database connection for an Informix replicate database using the pre-defined Informix function-string class provided with Replication Server.

You must modify the *hds_informix_connection_sample.sql* template to include your actual server, database, and Maintenance User names before executing the script.

Note You need not execute the *hds_informix_connection_sample.sql* script if you use Replication Server Manager (RSM) to define the Replication Server replicate database connection.)

Informix system management issues

The following system management issues are specific to an Informix replicate data server:

- When using Replication Server Manager (RSM), the value of the `DirectConnect DateTimeResults` configuration parameter must be set to `datetime` for latency calculations to work correctly.
- The create connection command's `dsi_sql_data_style` parameter was used in previous versions of Replication Server to provide some data translations for the replicate database.

With the introduction of the heterogeneous datatype support (HDS) feature in Replication Server version 12.0, the create connection command's `dsi_sql_data_style` parameter is no longer valid. Do *not* use this parameter with Replication Server version 12.0 or later. The default setting should be " " (blank space).

- The Component Integration Services (CIS) feature allows users to access both Adaptive Server databases and non-Sybase databases on different servers, through a local Adaptive Server connection.

CIS allows Replication Server to connect to an Adaptive Server database, using default Adaptive Server function strings and connection properties, while CIS handles the data conversion and manipulation, and all communications with the DirectConnect database gateway.

See “Component Integration Services” on page 99 for a description of the advantages and disadvantages of this approach.

Microsoft SQL Server replicate data servers

This section describes the replicate database issues and considerations specific to the Microsoft SQL Server data server in a Sybase replication system.

As a replicate data server, Microsoft SQL Server interacts with the DirectConnect for Microsoft SQL Server database gateway. The DirectConnect for Microsoft SQL Server is responsible for accepting commands from the replicate Replication Server, and applying those commands to a Microsoft SQL Server database.

Note DirectConnect for Microsoft SQL Server supports replication of large object (LOB) datatypes (image, ntext, and text) from Replication Server directly to a Microsoft SQL Server database.

Microsoft SQL Server replicate database permissions

To apply transactions in a replicate database, Replication Server requires a Maintenance User ID that you specify in the Replication Server create connection command. The Maintenance User ID must be defined at the Microsoft SQL Server data server and granted authority to apply transactions in the replicate database. The Maintenance User ID must have the following permissions in the Microsoft SQL Server replicate database:

- create table authority to create tables used for Replication Server processing.
- update authority on all replicate tables.
- execute authority on all replicate stored procedures.

Replication intrusions and impacts in Microsoft SQL Server

The only significant intrusions or impacts to the Microsoft SQL Server replicate database are the database objects created by the *hds_msss_setup_for_replicate.sql* script to support Replication Server replicate database operations.

The *hds_msss_setup_for_replicate.sql* script creates two tables in the replicate database to support Replication Server operations:

- `rs_info`, which contains information about the sort order and character set used by the replicate database.

Note Verify that the insert statements for the `rs_info` table (in the *hds_msss_setup_for_replicate.sql* script) specify the proper character set and sort order for your Microsoft SQL Server data server.

When using Replication Server version 12.0 or later, the replicate database sort order and character set must be recorded in the `rs_info` table.

The Replication Server `rs_get_charset` and `rs_get_sortorder` functions retrieve the character set and sort order from the `rs_info` table in the replicate database.

- `rs_lastcommit`, which contains information about replicated transactions applied to the replicate database.

Each row in the `rs_lastcommit` table identifies the most recent committed transaction that was distributed to the replicate database from a primary database. Replication Server uses this information to ensure that all transactions are distributed.

- The Replication Server `rs_get_lastcommit` function retrieves information about the last transaction committed in the replicate database. For non-Sybase replicate databases, the `rs_get_lastcommit` function is replaced in the database-specific function-string class by the query required to access the `rs_lastcommit` table in the replicate database.

Replicate database connectivity for Microsoft SQL Server

A Replication Server database connection name is made up of two parts: a data server name (`server_name`) and a database name (`db_name`). The `server_name` is the name of the DirectConnect for Microsoft SQL Server database gateway server, and the `db_name` is the name of the Microsoft SQL Server replicate database.

The replicate Replication Server looks for an interfaces file entry for the database gateway `server_name` specified in the Replication Server database connection. The replicate Replication Server logs in to the replicate data server using the `user_name` and password specified in the database connection.

You must make an entry in the Replication Server interfaces file to identify the host and port where the DirectConnect for Microsoft SQL Server database gateway server is listening. The interfaces file entry name must match the `server_name` portion of the Replication Server database connection.

Replicate database limitations in Microsoft SQL Server

The following replication limitations exist with a Microsoft SQL Server replicate data server:

- Microsoft SQL Server supports either 28 digits of precision or 38 digits of precision, depending on the server's start-up options. The default precision is 28 digits. Replication Server does not provide user-defined datatypes (UDDs) to support the default 28 digits of precision.

If you attempt to replicate numeric data to a Microsoft SQL Server database in excess of the server's configured precision, Replication Server returns the following error:

```
E. 2001/12/14 11:14:58. ERROR #1028 DSI EXEC(134(1)
dcm_gabeat70_devdb.devdb)
- dsigmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
'[VENDORLIB] Vendor Library Error: [[Message
Iteration=1|Data Source Name=mssql70_devdb|
SQLState=22003|Native Error=1007|Message=
[Microsoft] [ODBC SQL Server Driver] [SQL
Server]The number
'999999999999999999.999999999999999999' is out
of the range for numeric representation (maximum
precision 28).
[Message Iteration=2|SQLState=22003|Native
Error=|Message=[Microsoft] [ODBC SQL Server
Driver] [SQL Server]The number
'0.999999999999999999999999999999999999' is out
of the range for numeric representation (maximum
precision 28).] <DCA>'
```

- Microsoft SQL Server appears to support identity columns in the same manner as Adaptive Server Enterprise, so the Replication Server function strings that set identity insert off and on work correctly with Microsoft SQL Server. However, to support 28-digit numeric precision, the Sybase native numeric datatype must be translated to the `rs_mssql_numeric` datatype, and as a result of this translation, the identity characteristic is lost.

If you choose to use the numeric to rs_mssql_numeric datatype translation to support 28-digit precision in a Microsoft SQL Server replicate database, the replicate table cannot declare the numeric column receiving that data as an identity.

If a replicate Microsoft SQL Server table declares a numeric column receiving translated data as an identity, Replication Server returns the following error:


```

E. 2001/12/14 12:05:39. ERROR #1028 DSI EXEC(134(1)
    dcm_gabeat70_devdb.devdb)
    - dsigmint.c(2888)
    Message from server: Message: 30291, State 0,
    Severity 19 --
    ' [VENDORLIB] Vendor Library Error: [[Message
    Iteration=1|Data Source Name=mssql70_devdb|SQL
    Function=INSERT|SQLState=23000|Native
    Error=544|Message=[Microsoft] [ODBC SQL Server
    Driver] [SQL Server] Cannot insert explicit value
    for identity column in table 'ase_alltypes' when
    IDENTITY_INSERT is set to OFF.] <DCA>'

```

Microsoft SQL Server
replicate database
configuration issues

The heterogeneous datatype support (HDS) feature of Replication Server provides a number of sample SQL scripts that help you set up the HDS feature in the replicate Replication Server and the Microsoft SQL Server replicate database. These include:

- *hds_msss_setup_for_replicate.sql* – a script to be applied to the Microsoft SQL Server replicate database.

The *hds_msss_setup_for_replicate.sql* script creates the *rs_info* and *rs_lastcommit* tables in the replicate database. The script includes grant statements that need to be changed before execution to reference the correct user ID name of the Maintenance User defined in the Replication Server database connection for the replicate database. The Maintenance User ID must have update authority on the tables.

Note Sybase recommends using the Maintenance User ID to execute this script.

- *hds_msss_funcstrings.sql* and *hds_msss_udds.sql* – scripts to be applied to Replication Server System Database (RSSD).

The *hds_msss_udds.sql* script adds the user-defined datatypes (UDDs) that define the attributes of Microsoft SQL Server native datatypes to the RSSD. The UDDs are required to ensure datatypes received from primary transactions are properly formatted for application to the Microsoft SQL Server replicate database.

Note You might need to modify the *hds_msss_udds.sql* script to reference the correct RSSD database name.

The *hds_msss_funcstrings.sql* script replaces several default Replication Server function strings with custom function strings designed to

communicate with Microsoft SQL Server and access the tables and procedures created by the *hds_msss_setup_for_replicate.sql* script. These function strings are added to the Replication Server default *rs_msss_function_class*.

Note You might need to modify the script to reference the correct RSSD database name.

- Class-level datatype translation scripts to be applied to RSSD.
Class-level translations identify primary datatypes and the replicate datatypes the data should be translated into (for example, DB2 TIMESTAMP should be translated to Microsoft SQL Server datetime).

Note These translations can affect Replication Server performance. Only the translations needed for your specific primary database and replicate database should be applied to the RSSD.

Class-level translation scripts supplied for the Microsoft SQL Server replicate database are:

- *hds_clt_db2_to_msss.sql* – translates DB2 for OS/390 datatypes to Microsoft SQL Server datatypes.
- *hds_clt_udb_to_msss.sql* – translates DB2 Universal Database (for UNIX and Windows) datatypes to Microsoft SQL Server datatypes.
- *hds_clt_informix_to_msss.sql* – translates Informix datatypes to Microsoft SQL Server datatypes.
- *hds_clt_oracle_to_msss.sql* – translates Oracle datatypes to Microsoft SQL Server datatypes.

Note Class-level translations are not supplied for Adaptive Server datatypes to Microsoft SQL Server datatypes (or Microsoft SQL Server datatypes to Adaptive Server datatypes) because Microsoft SQL Server datatypes are directly compatible with Adaptive Server datatypes and they require no translation.

- *hds_msss_connection_sample.sql* – a script to be applied to the replicate Replication Server.

The *hds_msss_connection_sample.sql* script provides a template for creating the Replication Server database connection for a Microsoft SQL

Server replicate database using the pre-defined Microsoft SQL Server function-string class provided with Replication Server.

You must modify the *hds_msss_connection_sample.sql* template to include your actual server, database, and Maintenance User names before executing the script.

Note You need not execute the *hds_msss_connection_sample.sql* script if you use Replication Server Manager (RSM) to define the Replication Server replicate database connection.

Microsoft SQL Server system management issues

The following system management issues are specific to a Microsoft SQL Server replicate data server:

- When using Replication Server Manager (RSM), the value of the DirectConnect DateTimeResults configuration parameter must be set to datetime for latency calculations to work correctly.
- The create connection command's *dsi_sql_data_style* parameter was used in previous versions of Replication Server to provide some data translations for the replicate database.

With the introduction of the heterogeneous datatype support (HDS) feature in Replication Server version 12.0, the create connection command's *dsi_sql_data_style* parameter is no longer valid. Do *not* use this parameter with Replication Server version 12.0 or later. The default setting should be " " (blank space).

- The Component Integration Services (CIS) feature allows users to access both Adaptive Server databases and non-Sybase databases on different servers, through a local Adaptive Server connection.

CIS allows Replication Server to connect to an Adaptive Server database, using default Adaptive Server function strings and connection properties, while CIS handles the data conversion and manipulation, and all communications with the DirectConnect database gateway.

See “Component Integration Services” on page 99 for a description of the advantages and disadvantages of this approach.

Oracle replicate data servers

This section describes the replicate database issues and considerations specific to the Oracle data server in a Sybase replication system.

As a replicate data server, Oracle interacts with the DirectConnect for Oracle database gateway. The DirectConnect for Oracle is responsible for accepting commands from the replicate Replication Server, and applying those commands to an Oracle database.

Oracle replicate database permissions

To apply transactions in a replicate database, Replication Server requires a Maintenance User ID that you specify in the Replication Server create connection command. The Maintenance User ID must be defined at the Oracle data server and granted authority to apply transactions in the replicate database. The Maintenance User ID must have the following permissions in the Oracle replicate database:

- CREATE TABLE authority to create tables used for Replication Server processing.
- UPDATE authority on all replicate tables and EXECUTE authority on all replicate stored procedures.

Replication intrusions and impacts in Oracle

The only significant intrusions or impacts to the Oracle replicate database are the database objects created by the *hds_oracle_setup_for_replicate.sql* script to support Replication Server replicate database operations.

The *hds_oracle_setup_for_replicate.sql* script creates two tables in the replicate database to support Replication Server operations:

- RS_INFO, which contains information about the sort order and character set used by the replicate database.

When using Replication Server version 12.0 or later, the replicate database sort order and character set must be recorded in the RS_INFO table.

Note You should confirm that the INSERT statements for this table (in the *hds_oracle_setup_for_replicate.sql* script) specify the proper character set and sort order for your Oracle data server.

The Replication Server *rs_get_charset* and *rs_get_sortorder* functions retrieve the character set and sort order from the RS_INFO table in the replicate database.

- RS_LASTCOMMIT, which contains information about replicated transactions applied to the replicate database.

Each row in the RS_LASTCOMMIT table identifies the most recent committed transaction that was distributed to the replicate database from a primary database. Replication Server uses this information to ensure that all transactions are distributed.

The Replication Server `rs_get_lastcommit` function retrieves information about the last transaction committed in the replicate database. For non-Sybase replicate databases, the `rs_get_lastcommit` function is replaced in the database-specific function-string class by the query required to access the RS_LASTCOMMIT table in the replicate database.

Replicate database connectivity for Oracle

A Replication Server database connection name is made up of two parts: a data server name (`server_name`) and a database name (`db_name`). The `server_name` is the name of the DirectConnect for Oracle database gateway server, and the `db_name` is the name of the Oracle SID for the replicate database.

The replicate Replication Server looks for an interfaces file entry for the database gateway `server_name` specified in the Replication Server database connection. The replicate Replication Server logs in to the replicate data server using the `user_name` and password specified in the database connection.

You must make an entry in the Replication Server interfaces file to identify the host and port where the DirectConnect for Oracle database gateway server is listening. The interfaces file entry name must match the `server_name` portion of the Replication Server database connection.

Replicate database limitations in Oracle

Replication of large object (LOB) datatypes to the Oracle LONG and LONG RAW datatypes is supported directly from Replication Server to Oracle through the Component Integration Services (CIS) feature of Adaptive Server. CIS sends the LOB data to an Oracle replicate database through the DirectConnect for Oracle gateway.

Replication of LOB datatypes to the Oracle BLOB, CLOB, and NCLOB datatypes is *not* supported directly from Replication Server.

Note For information about a possible workaround for this LOB data replication limitation, see “Large object replication” on page 25.

Oracle replicate database configuration issues

The heterogeneous datatype support (HDS) feature of Replication Server provides a number of sample SQL scripts that help you set up the HDS feature in the replicate Replication Server and the Oracle replicate database. These include:

- `hds_oracle_setup_for_replicate.sql` – a script to be applied to the Oracle replicate database.

The *hds_oracle_setup_for_replicate.sql* script creates the RS_INFO and RS_LASTCOMMIT tables in the replicate database. The script includes GRANT statements that need to be changed before execution to reference the correct user ID name of the Maintenance User defined in the Replication Server database connection for the replicate database. The Maintenance User ID must have UPDATE authority on the tables.

Note Sybase recommends using the Maintenance User ID to execute this script.

- *hds_oracle_funcstrings.sql* and *hds_oracle_udds.sql* – scripts to be applied to Replication Server System Database (RSSD).

The *hds_oracle_udds.sql* script adds the user-defined datatypes (UDDs) that define the attributes of Oracle native datatypes to the RSSD. The UDDs are required to ensure datatypes received from primary transactions are properly formatted for application to the Oracle replicate database.

Note You might need to modify the *hds_oracle_udds.sql* script to reference the correct RSSD database name.

The *hds_oracle_funcstrings.sql* script replaces several default Replication Server function strings with custom function strings designed to communicate with an Oracle data server and access the tables and procedures created by the *hds_oracle_setup_for_replicate.sql* script. These function strings are added to the Replication Server default rs_oracle_function_class.

Note You might need to modify the script to reference the correct RSSD database name.

- Class-level datatype translation scripts to be applied to RSSD.

Class-level translations identify primary datatypes and the replicate datatypes the data should be translated into (for example, DB2 TIMESTAMP should be translated to Oracle DATE).

Note These translations can affect Replication Server performance. Only the translations needed for your specific primary database and replicate database should be applied to the RSSD.

Class-level translation scripts supplied for the Oracle replicate database are:

- *hds_clt_ase_to_oracle.sql* – translates Adaptive Server datatypes to Oracle datatypes.
- *hds_clt_db2_to_oracle.sql* – translates DB2 for OS/390 datatypes to Oracle datatypes.
- *hds_clt_udb_to_oracle.sql* – translates DB2 Universal Database (for UNIX and Windows) datatypes to Oracle datatypes.
- *hds_clt_informix_to_oracle.sql* – translates Informix datatypes to Oracle datatypes.
- *hds_clt_msss_to_oracle.sql* – translates Microsoft SQL Server datatypes to Oracle datatypes.
- *hds_oracle_connection_sample.sql* – a script to be applied to the replicate Replication Server.

The *hds_oracle_connection_sample.sql* script provides a template for creating the Replication Server database connection for an Oracle replicate database using the pre-defined Oracle function-string class provided with Replication Server.

You must modify the *hds_oracle_connection_sample.sql* template to include your actual server, database, and Maintenance User names before executing the script.

Note You need not execute the *hds_oracle_connection_sample.sql* script if you use Replication Server Manager (RSM) to define the Replication Server replicate database connection.

Oracle system management issues

The following system management issues are specific to an Oracle replicate data server:

- When using Replication Server Manager (RSM), the value of the DirectConnect DateTimeResults configuration parameter must be set to datetime for latency calculations to work correctly.
- The create connection command's *dsi_sql_data_style* parameter was used in previous versions of Replication Server to provide some data translations for the replicate database.

With the introduction of the heterogeneous datatype support (HDS) feature in Replication Server version 12.0, the create connection command's *dsi_sql_data_style* parameter is no longer valid. Do *not* use

this parameter with Replication Server version 12.0 or later. The default setting should be " " (blank space).

- The Component Integration Services (CIS) feature allows users to access both Adaptive Server databases and non-Sybase databases on different servers, through a local Adaptive Server connection.

CIS allows Replication Server to connect to an Adaptive Server database, using default Adaptive Server function strings and connection properties, while CIS handles the data conversion and manipulation, and all communications with the DirectConnect database gateway.

See “Component Integration Services” on page 99 for a description of the advantages and disadvantages of this approach.

Other replicate database issues for Oracle

The following issues must be considered when using an Oracle replicate data server:

- To capture the century and time components of a temporal value sent to an Oracle replicate database, you must map temporal datatypes to the `rs_oracle_datetime` datatype instead of the `rs_oracle_date` datatype. However, you should use the Replication Server `datetime` datatype to declare replication definition columns with the `DATE` datatype in an Oracle primary database rather than mapping to the `rs_oracle_datetime` datatype.

The Oracle default date display format is `DD-MON-YY`. The Replication Agent presents this format without time values, unless the value of the Replication Agent `pdb_convert_datetime` configuration parameter is set to `true`.

- Since the Oracle data server does not support multiple databases within the same server instance, the effect of executing the Replication Server `rs_usedb` function is null. (The function string created by the `hds_oracle_funcstrings.sql` script performs no operation against the Oracle database for this function call.)
- In DirectConnect for Oracle version 12.0 or later, an additional trace flag allows the replicate Replication Server to control transaction commit boundaries when applying transactions to an Oracle replicate database.

Setting the value of the DirectConnect autocommit trace flag to 0 (zero) in the DirectConnect for Oracle configuration file allows Replication Server to control when a `COMMIT` command should be sent to Oracle. When the value of the autocommit trace flag is not set, DirectConnect for Oracle commits each individual operation (`INSERT`, `UPDATE`, and `DELETE`) sent by the replicate Replication Server.

Having DirectConnect commit each operation can cause a problem at the replicate database if an error occurs in the middle of a multiple operation transaction. In that event, the replicate Replication Server may attempt to re-send the entire transaction, while DirectConnect has already committed each individual operation.

To avoid this problem, set the value of the DirectConnect autocommit trace flag to 0 (zero).

Component Integration Services

Component Integration Services (CIS) is a feature of Adaptive Server Enterprise. CIS allows Adaptive Server to present a uniform view of enterprise data to client applications and provides location transparency to enterprise-wide data sources.

CIS allows users to access both Sybase and non-Sybase data servers. These external data sources include host data files and tables, views, and remote procedure calls (RPCs) in data servers such as Adaptive Server Enterprise, DB2, and Oracle.

Replication with CIS

The CIS feature allows an Adaptive Server to act as a proxy for a non-Sybase data server. Replication Server connects to the Adaptive Server proxy database as if it were a normal Adaptive Server database. CIS then takes responsibility, along with a DirectConnect gateway, for translating the replication data and commands to the replicate data server's native formats.

Replication without CIS

The heterogeneous datatype support (HDS) feature of Replication Server can provide datatype translation that was previously available only with CIS. Datatype translation is accomplished using translation definitions installed in the RSSD. Command translations are provided with custom function strings and function-string classes.

Advantages and disadvantages of CIS

If you are a current user of CIS, there might not be an advantage to using the HDS feature of Replication Server. Both CIS and Replication Server HDS facilitate replication into a non-Sybase replicate database.

If you are not currently using CIS to communicate with a non-Sybase database, using the HDS feature of Replication Server provides some advantages over using CIS:

- Replication system topography is simpler without CIS, and there is no CIS overhead.

- You have some greater freedom of component location if you do not use CIS; you are not tied to an existing Adaptive Server location or required to install and maintain another Adaptive Server to decentralize the processing.

The HDS feature of Replication Server does not support the following items, which require the use of CIS:

- Replicating to a non-Sybase data server that is not supported by HDS. The HDS feature in Replication Server supports replicate databases only in DB2 Universal Database, Informix, Microsoft SQL Server, and Oracle data servers. Any other non-Sybase replicate database requires the use of CIS.
- Replicating large object (LOB) data to the Oracle LONG and LONG RAW datatypes. You must use CIS and DirectConnect for Oracle to replicate LOB data to the Oracle LONG and LONG RAW datatypes.

Note For more information about using CIS, refer to the *Component Integration Services User's Guide*, which is part of the Adaptive Server Enterprise documentation set.

Replication Server Issues

This chapter describes the issues specific to Sybase Replication Server in a replication system with heterogeneous or non-Sybase data servers.

Topic	Page
Relationship with other system components	101
Database connections	105
Replication Server heterogeneous datatype support	115
Emulating rs_init activity for a non-Sybase database	122
Case sensitivity in publications and subscriptions	125

Relationship with other system components

Replication Server interacts with other components of a replication system as either a *server* or a *client*.

As a server, Replication Server supports connections from:

- Replication Agents, across which database commands are sent from primary databases
- Other Replication Servers, thus distributing the processing involved in message delivery and providing a degree of scalability in a replication system
- Users or management tools for administration, data server identification, message publication and subscription, and so on

As a client, Replication Server connects to:

- Replicate Replication Servers for message delivery
- Other Replication Servers, thus distributing the processing involved in message delivery and providing a degree of scalability in a replication system

Replication Server
communication
protocols

Replication Server also connects to an external Adaptive Server Enterprise database that acts as the Replication Server System Database (RSSD), if it is not configured to use an embedded RSSD.

Replication Server is an Open Client/Server application that uses Sybase TDS as the underlying communication protocol. As such, any clients that request services from Replication Server must implement an Open Client interface. This includes Replication Agents, system management tools, and user interface tools such as isql.

As a client distributing messages to other Replication Servers or to replicate data servers, Replication Server speaks with an Open Client interface. Therefore, when Replication Server needs to send a message to a data server, either that data server must support an Open Server interface running on TDS, or there must be an Open Server/TDS bridge or “gateway” application between Replication Server and the replicate data server.

Typically, this gateway software is in the form of a Sybase DirectConnect database gateway. Some DirectConnect gateways bridge from Open Server/TDS to the native interface of the replicate data server (for example, DirectConnect for Oracle) while others bridge from Open Server/TDS to an ODBC or JDBC driver for the data server. Replication Server configurations vary, depending on the gateway used.

Replication Server
user IDs and
permissions

Replication Server requires several different user IDs. Some user IDs are required for other components (or users) to access the Replication Server, and others are required for the Replication Server to have access to other components in a replication system.

User IDs are defined in the Replication Server using the Replication Server create connection command.

Note Depending on how your replication system is configured, some of the user IDs in the following list might not be required. For example, if you have separate Replication Servers for primary and replicate databases, the primary Replication Server does not require a user ID for access to a replicate database.

The following user IDs are defined in a Replication Server:

- Replication Agent user – used by a Replication Agent to log in to a primary Replication Server. This user ID must have connect source permission to deliver database commands through the LTL interface.

- Replication Server user – used by other Replication Servers to log in to a Replication Server and forward messages. This user ID must have connect source permission to forward database commands through the RCL interface.
- SysAdmin user – used by System Administrators or system administration tools to perform administration activities. Depending on the task, this user ID must have sa, create object, or primary subscribe permission.
- Maintenance User – used by a replicate Replication Server to deliver messages to a replicate data server. This user ID must have the necessary permissions in the replicate data server to execute the commands to which messages to be delivered are mapped.
- Replicate user – used by a replicate Replication Server to deliver messages to a primary data server. For delivery of “request” messages, that is, messages from a replicate data server that are selected for delivery to the primary data server, Replication Server uses the user ID of the user who executes the command in the replicate database. This user ID must have the necessary permissions in the primary data server to execute the commands to which messages to be delivered are mapped.
- RSI user – used by Replication Server to log in to other Replication Servers to forward messages to be delivered. This user ID must have connect source permission in the replicate Replication Server.
- RSSD user – used by Replication Server to log in to the Replication Server System Database (RSSD) that manages its operational data. This user ID must have full control in the RSSD to create and drop objects, execute procedures, and query and update tables.

Relationship with Replication Agents

While Replication Server is extensible to meet the needs of replicate data servers (customizable function strings and error handling, custom datatype definitions, and translations between datatypes), Replication Server support of primary data servers is limited.

The Replication Server interface for primary data servers is its proprietary Log Transfer Language (LTL). Transactions from a primary data server must be translated to LTL to be delivered to a primary Replication Server. Therefore, primary data servers are limited to those for which Sybase provides a Replication Agent to perform the translation to LTL for primary database operations.

Replication Server interfaces on both the primary and replicate sides are supported by the underlying Open Client/Open Server interface running on TDS.

LTM locator updates

The primary Replication Server maintains a “locator” value (LTM locator), which is used to identify the last point in a transaction log from which all data has been successfully received by the primary Replication Server. The Replication Agent periodically requests this value from the Replication Server connection to identify a position in the transaction log, which can then be used to identify where older data can be released or removed from the log.

There is a performance trade-off in determining how often to request an LTM locator update. Frequent queries of the LTM locator value from a Replication Server can slow down replication (the Replication Agent must stop sending LTL commands long enough to request and receive the LTM locator value) while it provides more frequent opportunities to release data from the log. When restarting, the Replication Agent must re-send all data in the log that exists since the last LTM locator value was received from Replication Server.

Generally, if replication throughput performance is a priority, you should acquire enough log resource to allow less frequent log truncation and less frequent retrieval of the LTM locator value. If log resources are scarce, more frequent retrieval of the LTM locator value and more frequent truncation may be necessary.

For more information on using the LTM Locator, refer to the appropriate Replication Agent *Administration Guide*.

LTL generation

The number of bytes of information sent to Replication Server has a direct impact on the performance of the replication system; more data and commands received by Replication Server require more work and time to process. In addition, more data also requires more resources from the network environment. There are several configuration options available for the Replication Agent that you can use to minimize this impact. Among these are:

- Using the RSSD. By reading replication definitions from the RSSD, the Replication Agent can send the column data in the same column order as specified by the replication definition. This allows Replication Server to bypass sorting the column information before processing. Furthermore, column names are not sent with the data, which reduces the number of bytes of information required.
- Sending minimal columns. When an update operation occurs on a table, only a portion of the columns may have been altered. By sending the before and after images of only those columns that changed, the Replication Agent sends less information.

- Batch mode. A Replication Agent must “wrap” transactions in a limited amount of administrative LTL for the Replication Server. In batch mode, the Replication Agent can wrap multiple commands in the same set of administrative commands, which reduces the overall LTL generated and processed by the network and the Replication Server.

In addition to batch mode, most Replication Agents have a “batch timeout” parameter, which allows a partial batch to be sent to the Replication Server after the Replication Agent waits a specified period of time and no additional transactions are received to fill the batch.

Note Replication Agent batch mode should not be used if any Replication Server user-defined datatype (UDD) translations are used, either column-level or class-level.

- Origin time. Each transaction sent to Replication Server has an *origin queue ID*. The origin queue ID may include the time that the transaction was committed at the primary database. If the origin time is not sent by the Replication Agent, its processing effort is reduced somewhat, but the quantity of LTL sent to the Replication Server is the same.

Note The origin time must be sent if you want to use Replication Server Manager to calculate replication latency.

For a complete description of the Replication Agent configuration parameters that affect LTL output, refer to the appropriate Replication Agent *Administration Guide*.

Database connections

Replication Server keeps track of other components in a replication system using *connections* that identify primary and replicate databases and *routes* that identify other Replication Servers.

Since Replication Server was originally designed for Adaptive Server Enterprise database replication, the definition of a connection in Replication Server follows the Sybase standard of a *server name* followed by a *database name*. For example, a Replication Server connection to an Adaptive Server named ASE1 and database PUBS would be ASE1.PUBS.

When you create a Replication Server connection using Replication Server Manager (RSM), RSM builds the connection using the server name and database name of the database you are browsing.

Replication Server cannot connect directly to a non-Sybase data server. In the case of a primary database, Replication Server allows a connection from a Replication Agent on behalf of the non-Sybase primary database. In the case of a replicate database, Replication Server connects to a DirectConnect database gateway, which in turn connects to the non-Sybase replicate data server. Since Replication Agents and DirectConnect gateways are not data servers themselves, the Replication Server connection properties for those components may have different meanings than they would for a database server connection.

Note If Replication Server uses CIS to connect to a replicate data server, the server name and database name are the Adaptive Server name and the Adaptive Server database name that represent the CIS connection.

A single Replication Server connection can support data flow in either one or two directions. Data flows *in* through a Replication Server connection by way of the Rep Agent User thread. Data flows *out* through a Replication Server connection by way of the Data Server Interface (DSI) thread. Each Replication Server connection can support either outbound data flow only (through the DSI thread), or both inbound and outbound data flow (through the Rep Agent User and DSI threads).

Rep Agent User thread

Replication Server receives all data-change operations or transactions to be replicated from a primary data server through the Rep Agent User thread of the database connection for that data server. Every primary database that supplies transactions to be replicated must be represented by a Replication Server database connection with an enabled Rep Agent User thread.

Replication Server establishes a connection directly with the primary database, if it resides in an Adaptive Server. If the primary database resides in a non-Sybase data server, a separate Replication Agent component communicates with the Replication Server, using a Rep Agent User thread connection, on behalf of the primary database.

Note Replication Server never attempts to connect to the Rep Agent User thread of a connection. The only entity that can initiate communication to a Rep Agent User thread is the primary data server or the Replication Agent.

On a Rep Agent User thread, the primary data server or Replication Agent is the client, and the primary Replication Server is the server.

DSI thread

The DSI thread of a Replication Server connection is where the replicated transaction is delivered by Replication Server. Every replicate database expected to receive replicated transactions must be represented by a Replication Server connection with an enabled DSI thread.

Replication Server establishes a connection directly with the replicate database, if it resides in an Adaptive Server. If the replicate database resides in a non-Sybase data server, Replication Server communicates with a DirectConnect database gateway (or MainframeConnect in a gatewayless environment), by way of the connection's DSI thread.

Note A replicate data server or database gateway never attempts to connect to the DSI thread of a connection. The only entity that can initiate communication to a DSI thread is the Replication Server.

On a DSI thread, the Replication Server is the client, and the replicate data server or database gateway is the server.

Replication Agent connections

A Replication Agent sends data one-way to Replication Server. It logs in to the Replication Server, connects as an LTM thread to the Rep Agent User thread of a Replication Server connection, and communicates with Replication Server over that connection. This has the following implications:

- A valid user ID, which the Replication Agent uses to log in to the Replication Server, must be defined at the Replication Server.

- The Replication Agent user ID must be granted connect source permission in Replication Server. connect source permission allows the Replication Agent to send commands that are valid only on a Rep Agent User thread.
- The Replication Agent must record this user ID and associated password.
- The Replication Agent must record the *server* and *database* portions of the Replication Server connection definition to identify and connect to the correct Rep Agent User thread.
- The DSI thread for this connection is not used. Replication Agents do not support sending transactions or other operations back to a primary database.

Note Adding the with `dsi_suspended` clause to the Replication Server create connection command suspends the DSI thread for the Replication Agent connection.

- The `user_name` and password defined in the Replication Server create connection command must be a valid user ID and password for the primary database represented by the Replication Agent.

The Replication Agent validates that the connection `user_name` exists in the primary database. However, Replication Server does not know if (or when) a DSI thread will be used. Therefore, the user ID and password must be valid in case the DSI thread is active.

Note The requirement for a valid primary database user ID varies by Replication Agent. Some Replication Agents do not require (nor do they check for) a valid user ID on the Replication Server connection.

- The Replication Server resume connection command initiates activity with the DSI thread of the specified connection. The Replication Agent will not respond to an attempt to connect the DSI thread, since the Replication Agent uses only the Rep Agent User thread. Since Replication Server does not initiate communication with Rep Agent User threads, no activity takes place.

Attempts by Replication Server to communicate with the DSI thread on a Replication Agent connection are recorded as failures in the Replication Server log.

Interfaces file

For the interaction between a Replication Agent and a Replication Server, the only interfaces file entry that may be required is one identifying the Replication Server.

The Replication Agent for DB2 UDB for OS/390 requires an interfaces file entry for the Replication Server.

The Sybase Replication Agent (for DB2 Universal Database on UNIX and Windows platforms, Informix, Microsoft SQL Server, and Oracle) does not require an interfaces file entry, as it records the Replication Server host name and port number in configuration parameters.

Note If you use RSM to manage a Replication Agent, an interfaces file entry must exist for the Replication Agent in the RSM interfaces file.

The values of the server and database names are not significant in a Replication Server connection definition for a Rep Agent User thread. However, they must be properly configured in the Replication Agent.

The Replication Agent `rs_source_ds` and `rs_source_db` configuration parameters must be configured to match the `ds.db` values of the Replication Server connection name. If they do not match, the Replication Agent will not find a Replication Server connection to send information to.

Since Replication Server never attempts to contact the Rep Agent User thread of a connection, if the connection is only for a Rep Agent User thread, there is no need for the server name portion of the connection to exist in an interfaces file. The database name portion of the Replication Server connection is also arbitrary, if the connection is only for a Rep Agent User thread.

Any unique combination of two strings (`namex.namey`) can be used to identify a Rep Agent User thread connection. For convenience, you should choose names that have some meaning in your particular environment (for example, `host.db2_subsystem`, `host.oracle_sid`, and so forth).

Note If DB2 Universal Database is the primary data server, you should use the actual DB2 database name so RSM can assist in creating replication definitions for the primary database.

DirectConnect database gateways

In a Sybase replication system, the purpose of a DirectConnect database gateway is to apply transactions from a Replication Server to a non-Sybase replicate database.

To accomplish this, Replication Server logs in to the DirectConnect gateway using the information specified for a Replication Server connection. Replication Server logs in to the server using the `user_name` and password, and issues a `use database` command for the database defined in the connection.

For Replication Server, there is nothing to distinguish a DirectConnect gateway from an Adaptive Server replicate database. Replication Server delivers the same commands, and expects the same results, from any DSI thread it communicates with.

This has the following implications:

- A valid user ID, which the Replication Server uses to log in to the replicate database, must be defined in a Replication Server connection.
- This user ID must be granted permissions to update replicate tables and execute replicate procedures.
- The replicate database must be able to maintain a `rs_lastcommit` table and support `rs_get_lastcommit` functionality.

Replication Server provides sample scripts to set up the tables and functions required for a replicate database in DB2 Universal Database, Informix, Microsoft SQL Server, and Oracle databases.

For an overview of the expectations of a replicate data server and gateway, see Chapter 6, “Replicating Data into Foreign Data Servers,” in the Replication Server *Design Guide*.

- Datatype representations must be translated to match the native datatypes of the replicate database. Replication Server provides sample scripts to set up the function strings, function-string classes, and base datatype definitions and translations necessary to support replication into DB2 Universal Database, Informix, Microsoft SQL Server, and Oracle data servers.
- The Rep Agent User thread for this connection is not used (the DirectConnect gateway does not send transactions as a primary database). Remove the `with log transfer on keywords` from the Replication Server `create connection` command to bypass creation of a Rep Agent User thread for this connection.

In RSM, when you create the Replication Server connection, select the Replicate Only check box in the Add Connection wizard.

- The Replication Server `resume connection` attempts to initiate activity with the DSI thread of the specified connection. For a DirectConnect, this is a sequence of logging in to the DirectConnect

server, accessing the `rs_lastcommit` table in the replicate database, and then applying transactions to the replicate database. Any failure in this sequence is recorded as a failure in the Replication Server log.

Interfaces file

Replication Server is an Open Server application. In an Open Server application, the preferred method for determining the location (host and port number) of another Open Server is to look up the location in an interfaces file. The interfaces file contains a list of labels, typically server names, that each have a corresponding host name and port number, where the identified server should be “listening” for login requests.

In the interaction between a DirectConnect database gateway and a Replication Server, the interfaces file is important. Because the Replication Server does attempt to log in to the server identified by the server name in the Replication Server connection, that server name must exist in the Replication Server interfaces file. In addition, the interfaces file entry must also exist as a service name in the DirectConnect gateway configuration file entries.

A single DirectConnect can act as a gateway for one or many different database instances. In the DirectConnect configuration, each database to be accessed by the DirectConnect is configured as a unique *service name*. For DirectConnect to know which configured service name a client wants to connect to, it uses the server name passed at login time and expects to find a matching service name to use to complete the connection. Not only must the Replication Server `server_name` in the connection match an interfaces file entry, but that interfaces file entry must match the DirectConnect service name describing the database you want to connect to. Refer to the *DirectConnect Access Service User's Guide* for more information on the role of service names and their configurations.

Connection shared by Replication Agent and DirectConnect

A single Replication Server connection can support both a DirectConnect gateway and a Replication Agent, since each of these components connects to the Replication Server on a different thread. If you replicate information both into and out of the same database, having a common connection for both a database gateway and a Replication Agent may make the replication system network topology simpler. On the other hand, since the Replication Agent and the DirectConnect gateway are two separate servers, it might be more intuitive to have a separate connection for each one.

To share a Replication Server connection between a DirectConnect gateway and a Replication Agent, you must define the connection to correctly support the replicate database, then configure the Replication Agent appropriately:

- In the Replication Server, use the `create connection` command to define the `server_name` and `database_name` for the connection. The `server_name` value must match a configured service name in the DirectConnect.
- In the Replication Agent, set the value of the `rs_source_ds` parameter to that `server_name` and the value of the `rs_source_db` parameter to the desired `database_name`.

In RSM, create the connection as primary and replicate in the Add Connection wizard. When configuring the Replication Agent, set the values of the `rs_source_ds` and `rs_source_db` parameters to the same values as the `server_name` and `database_name` created by the Add Connection wizard.

Maintenance User

Maintenance User purpose

This section describes the role and usage of the Replication Server Maintenance User.

To update replicated data, Replication Server logs in to the replicate data server as the Maintenance User. The Database Owner (or the System Administrator) must grant to the Maintenance User the permissions required to insert, delete, and update rows in replicated tables, and to execute replicated stored procedures. In an Adaptive Server replicate database, Sybase Central or `rs_init` creates the user ID for the Replication Server Maintenance User and adds the user to the replicate database.

The Maintenance User ID and password are defined to Replication Server with the Replication Server `create connection` command for the replicate database. In Adaptive Server, Sybase Central or the `rs_init` program executes this command automatically. If you change the password for the Maintenance User ID in the data server, you can use Sybase Central or the Replication Server `alter connection` command to change the password for the Replication Server connection.

The Replication Server Maintenance User must also have permission to access the `rs_lastcommit` and `rs_info` system tables in the replicate database, and any stored procedures that use those tables.

Neither Sybase Central nor `rs_init` grants database permissions to the Maintenance User for user tables and stored procedures. You must grant database permissions on replicated tables and stored procedures before you can replicate transactions for replicated tables or replicate executions of the replicated stored procedures. For each table replicated in the database, and for

each stored procedure executed due to replication, you must execute the following grant command:

```
grant all on table_name to maint_user
```

Alternatively, you can assign the Maintenance User ID (maint_user) to a Database Administrator role, if that role has the required authority on all replicate objects.

Replication Agent Maintenance User processing

When the Replication Agent connects to a Replication Server connection, the Replication Agent requests the user ID of the Maintenance User and validates that the user ID exists in the primary database. This validation requires that the Maintenance User ID defined in any Replication Server connection be valid for the database the connection represents, regardless of whether that connection is for primary transactions only, replicate transactions only, or both.

The Replication Agent does not log in to the primary database using the Maintenance User ID. Other than validating that the user ID exists, the only reference the Replication Agent makes to the Maintenance User ID is to filter out primary database transactions created by the Maintenance User.

The Replication Agent filters out Maintenance User transactions to avoid having a transaction applied more than once to the primary database. In a bidirectional replication scheme, replication can occur both to and from the same database (a single database having both a primary and a replicate role). When a primary transaction is applied to a replicate database, the applying user ID is the Maintenance User for the replicate database. A Replication Agent scanning transactions at the replicate database must ignore the transactions applied by the Replication Server Maintenance User to prevent those transactions from being sent back and applied to the primary database.

The Replication Agent accesses the database using a user ID defined at the primary database (or for DB2, a user ID that can access the DB2 log files). This user ID is not the same as the Maintenance User defined in the Replication Server connection. The Replication Agent user ID used to access the primary database has a different role and purpose than the Maintenance User defined to apply replicated transactions.

There may also be another user ID defined to the Replication Agent that is used to administer the Replication Agent. This user ID is also separate from the Replication Server Maintenance User that applies replicate transactions.

A Replication Agent can have knowledge of three different users:

- A user ID defined at the primary database, which the Replication Agent uses to log in to the primary data server and manipulate primary replication objects or read the database transaction log.

- A user ID that can log in to the Replication Agent and issue Replication Agent commands and configure Replication Agent parameters.
- A Maintenance User ID, defined at the primary database and recorded in the primary Replication Server connection. The Replication Agent validates this user ID on behalf of the Replication Server, and the Replication Agent can be configured to ignore transactions that are created by this user ID.

DirectConnect
Maintenance User
processing

In general, using a DirectConnect database gateway does not change the use or purpose of the Replication Server Maintenance User. The only difference is that Adaptive Server provides some utilities that automatically create and assign permissions to the Maintenance User (for example, `rs_init`). With non-Sybase databases, the Maintenance User ID must be defined and granted the appropriate permissions.

Different types of
users

In a Sybase replication system, there are several unique user IDs required. Sybase recommends that you provide a unique user ID for each role and component combination in your environment.

Each component (database, Replication Server, Replication Agent, and DirectConnect) must have some user ID defined to maintain or administer that component. Typically, maintenance tasks involve changing configuration parameters and access for the component “server.”

For Sybase products, the default component System Administrator user ID is typically `sa`. While it may be tempting to leave all of the components with an `sa` user defined, you may grant access to more components and more authority than intended or prudent. Do not confuse a component System Administrator user with a Replication Server connection Maintenance User. The Replication Server Maintenance User has authority to apply database transactions at a replicate database. It does not require authority to maintain the configuration of a replication system component.

Each Sybase Replication Agent requires a user ID in the primary database that can modify primary tables and procedures marked for replication and scan log files. This user ID does not perform the same function as the Replication Server Maintenance User ID, which is responsible for applying replicate transactions in a replicate database.

Maintenance User
checklist

The following are requirements for the Replication Server Maintenance User:

- A Maintenance User ID must be created in every primary and replicate database in a replication system.

- The Maintenance User ID must be identified in the Replication Server create connection command for every primary and replicate database in a replication system.
- For replicate databases, the Maintenance User ID must have authority to apply replicate transactions and update any Replication Server tables or functions installed at the replicate database (for example, rs_info and rs_lastcommit).

Replication Server heterogeneous datatype support

Reconciling incompatible data representations between different types of data servers is one of the major issues that must be addressed to implement a successful heterogeneous replication system.

Heterogeneous datatype support (HDS) enables you to specify automatic datatype translation when primary and replicate data servers use datatypes that are not directly compatible.

The HDS feature provides datatype definitions for the following non-Sybase data servers:

- DB2 Universal Database
- Informix Dynamic Server
- Microsoft SQL Server
- Oracle

Datatype translation

In general, datatype translation can be divided into two categories: trivial and nontrivial. Trivial datatype translations can be reduced to a translation between two Replication Server native datatypes (for example, int to char, datetime to char, or decimal to real). These translations are supported by the Open Client/Server function `cs_convert`. Refer to the Open Client/Server documentation on `cs_convert` for the kinds of translations supported and any limitations.

Nontrivial datatype translations are those in which the target datatype is a user-defined datatype, and where the user-altered attributes include more than just

delimiters. For example, translation from the Sybase datetime datatype to the DB2 TIMESTAMP datatype is considered nontrivial, because these two temporal datatypes support different ranges of values. However, a translation from binary to the DB2 CHAR FOR BIT DATA datatype is trivial, since the two sets of attributes are identical except for their delimiters.

In Replication Server, datatype translations can be specified at:

- The function-string class level (or class-level translations)
- The replication definition level (or column-level translations)

Each Replication Server connection is associated with exactly one function-string class. Therefore, the datatype translations defined for a function-string class are associated with each connection that is defined with that function-string class. Class-level translations define default translations for the connection by specifying a set of source and target datatypes so that any value sent over that connection with a datatype that matches the source datatype is translated into the associated target datatype.

Unlike function strings, class-level translations are not inherited. Moreover, class-level translations offer the only means to translate values associated with system-defined variables.

In Replication Server, there is no user interface for creating class-level translations. Therefore, Sybase supplies SQL scripts to create class-level translations that support some of the more common datatype translations.

Column-level translations can be specified in a replication definition for a given column by using the map to clause of the create replication definition or alter replication definition command. You can use column-level translations to fine-tune datatype translation on a column-by-column basis, per replication definition.

Column-level translations provide a performance advantage over class-level translations. However, only class-level translations can be used to translate values of system-defined variables.

HDS issues and limitations

This section describes some known issues and limitations with the HDS feature in Replication Server.

Command batching
on the DSI thread

Turn off “command batching” when you create the DSI thread connection to a non-Sybase replicate data server.

Command batching from Replication Server is not supported by DirectConnect gateways for all non-Sybase databases. Command batching specifies how Replication Server sends commands to replicate data servers. When command batching is enabled, Replication Server may send multiple commands to the replicate data server as a single command batch. When command batching is not enabled, Replication Server sends commands to the replicate data server one at a time.

Most DirectConnect gateways perform better when processing one command at a time.

Source value exceeds
target datatype
bounds

The datatype translations provided by Sybase specify that the thread attempting a translation where the source value exceeds the bounds of the target datatype should be stopped with the following error message:

```
E. 2001/12/14 11:14:54. ERROR #32055 DSI EXEC(135(1)
snickers_dco.ora805) -
/nrm/nrm.c(7023)
Class Level translation for column/parameter
'datetimecol' failed.
Source DTID is 'datetime'.
Target DTID is 'rs_oracle_datetime'.
Function String Class ID 'rs_oracle_function_class'.
Value length is '21'; Maximum target length is '20';
The value is '99991231 23:59:59:010'
```

These are typically the most difficult translation problems to diagnose because there appears to be no problem with either the pairing of source/target datatypes or the value to be translated.

To diagnose this type of problem, you have to be familiar with the datatype value bound limits of all the translated target datatypes. For example, to diagnose the error shown, you have to know that the upper bound of an Oracle DATE value is 12/31/4712.

There are other options:

- Use the maximum value of the datatype definition
- Use the minimum value for the datatype definition
- Use the default value for the datatype definition

The only way you can specify these options is to edit the minimum and maximum boundary error action column values in each *hds_XXX_udds.sql* script and re-install the datatype definitions by re-executing the *hds_XXX_udds.sql* scripts in the RSSD.

Microsoft SQL Server supports either 28 or 38 digits of precision, depending on how the server is started. By default, Microsoft SQL Server supports 28 digits of precision.

When a number exceeds numeric precision of the Microsoft SQL Server replicate database, Replication Server returns the following error:

The most difficult numeric datatype issues involve precision and scale. Replication Server does not allow the precision and scale of a decimal datatype to be specified. A datatype definition can specify the maximum precision and maximum scale to be supported. However, if this does not equate to the specified precision and scale of an individual replicate column, then as the data approaches values near or at the boundaries, you may encounter problems that are reported differently, depending on the replicate data server.

Replication Server

you specify the translation for the primary data server decimal datatype and for which there is a class-level translation to the replicate data server decimal datatype. Both datatype definitions specify the associated data server's maximum precision and scale.

If the value 999.99999 comes from the primary database, and the replicate data server's datatype definition specifies that rounding should be attempted, Replication Server attempts to apply a value of 1000.000. Even though this value satisfies the replicate database requirements for maximum precision and scale, it fails the precision and scale specified for this particular column. And if you specify for the replicate database's datatype definition that it should replace the value with the specified maximum value for the datatype definition, Replication Server attempts to apply a value of 9999999, which also fails the specified precision and scale for this particular column.

Error messages you might see from various data servers in this case include:

- The following DB2 error:

```
E. 2001/12/14 15:03:11. ERROR #1028 DSI EXEC(129(1)
dwm5_via_rct.dwmdbas)
- dsigmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
' [VENDORLIB] Vendor Library Error: [[Message
Iteration=1|SQLState=22003|Native Error=
-413|Message=[Sybase] [ClearConnect ODBC] [DB2] The
decimal or numeric value had an incorrect wire
length compared to its specified FDOCA length
1000000000000000000.0000000000] <DCA>'.
```

- The following Informix error:

```
E. 2001/12/14 16:07:31. ERROR #1028 DSI EXEC(136(1)
dock723_dev_db.devdb)
- dsigmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
' [VENDORLIB] Vendor Library Error: [[Message
Iteration=1|SQLState=S1000|Native Error=
-1226|Message=[Visigenic] [ODBC Informix 7.2
Driver] [Informix]-1226: Decimal or money value
exceeds maximum precision.] <DCA>'
```

- The following Microsoft SQL Server error:

```
E. 2001/12/14 12:29:16. ERROR #1028 DSI EXEC(134(1)
dcm_gabeat70_devdb.devdb)
```

```
- dsigmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
' [VENDORLIB] Vendor Library Error: [[Message
Iteration=1|Data Source Name=mssql70_devdb|SQL
Function=INSERT|SQLState=22003|Native Error=
8115|Message=[Microsoft][ODBC SQL Server Driver]
[SQL Server]Arithmetic overflow error converting
numeric to data type numeric.[Message Iteration=
2|SQLState=01000|Native Error=|Message=
[Microsoft][SQL Server]The statement has been
terminated.] <DCA>'
```

Informix integer and numeric datatypes

Most data servers support integer datatypes with values ranging from -2147483648 to 2147483647 and small integer datatypes with values ranging from -32768 to 32767. Informix supports INT datatypes with values ranging from -2147483647 to 2147483647 and SMALLINT datatypes with values ranging from -32767 to 32767.

Informix supports only up to 32 digits of precision/scale in exact numerics, and Adaptive Server supports up to 38 digits of precision/scale.

Sybase does not provide datatype definitions that:

- Limit the lower bound of Informix integer datatypes
- Support the restriction of precision/scale in Informix numeric datatypes

Informix binary datatypes

You cannot directly insert or update binary datatypes in Informix. If you need to replicate binary data to an Informix database, you must capture the binary literal value as a char or varchar datatype. However, there is no class-level translation from Adaptive Server datatypes to Informix datatypes that handles the bit binary datatype.

The problem is that, without a translation, the bit value goes out as 0x00 or 0x01, which is syntax that Informix does not understand. In this situation, you receive an error similar to this:

```
E. 2001/12/14 15:11:14. ERROR #1028 DSI EXEC(136(1)
dock723_dev_db.devdb)
- dsigmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
' [VENDORLIB] Vendor Library Error: [[Message
Iteration=1|SQLState=S1000|Native Error=
-280|Message=[Visigenic][ODBC Informix 7.2 Driver]
[Informix]-280: A quoted string exceeds 256 bytes.]
<DCA>'
```

If you attempt to rely on class-level translations to handle commands with bit data, you receive an error similar to this:

```
E. 2001/12/14 15:19:57. ERROR #1028 DSI EXEC(136(1)
dock723_dev_db.devdb)
- dsiqmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
' [VENDORLIB] Vendor Library Error: [[Message
Iteration=1|SQLState=37000|Native Error=
-201|Message=[Visigenic] [ODBC Informix 7.2 Driver]
[Informix]-201: A syntax error has occurred.] <DCA>'.
```

Numeric translation
and identity columns
in Microsoft SQL
Server

Replication Server function strings to set identity insert off and on work in Microsoft SQL Server because it supports identity columns in the same manner as Adaptive Server. However, to support 28-digit precision in a Microsoft SQL Server database, the numeric datatype must be translated to the `rs_mssql_numeric` datatype, and as a result, the identity characteristic is lost. To avoid this problem, the Microsoft SQL Server replicate table must not declare a translated numeric column as an identity.

If you attempt to replicate a translated numeric datatype into an identity column in Microsoft SQL Server, you receive an error similar to this:

```
E. 2001/12/14 12:05:39. ERROR #1028 DSI EXEC(134(1)
dcm_gabeat70_devdb.devdb)
- dsiqmint.c(2888)
Message from server: Message: 30291, State 0,
Severity 19 --
' [VENDORLIB] Vendor Library Error: [[Message
Iteration=1 |Data Source Name=mssql70_devdb|SQL
Function=INSERT|SQLState=23000|Native Error=544
|Message=[Microsoft] [ODBC SQL Server Driver] [SQL
Server]Cannot insert explicit value for identity
column in table 'ase_alltypes' when IDENTITY_INSERT
is set to OFF.] <DCA>'
```

Datatypes and
datatype definitions

Datatype definitions for a particular data server datatype are grouped in a *datatype class*. This enables Replication Server Manager (RSM) to provide a list of available datatypes for a given data server type when replication definitions are created or altered.

For more information about datatype definitions (user-defined datatypes), see the description of the RSSD `rs_datatype` table in the Replication Server *Reference Manual*.

Restricted datatypes

The `text`, `image`, `raw` and `rs_address` datatypes cannot be used as either the source or target of column-level or class-level translations.

Error and function-string classes for non-Sybase data servers

Sybase provides function-string classes and associated function strings for all supported non-Sybase replicate data servers. Sybase does not provide unique error classes for non-Sybase data servers.

The following limitations apply to HDS function strings:

- The function-string classes are loaded into the RSSD when the RSSD is installed. However, the individual function strings must be loaded separately.
- The function strings are designed to work with the replicate setup scripts provided by Sybase. Using any other technique to set up the replicate data server, in particular creating the Replication Server support objects (tables), may result in the function strings not working properly.
- After installing the function strings, the Replication Server must be restarted. This is because, with the exception of “input template” function strings, Replication Server caches all function strings at startup.

Object publication and subscriptions

The following limitations apply to object publications and subscriptions in a Sybase replication system:

- When declaring columns in a replication definition for a non-Sybase primary database, you must use the Replication Server datatype that matches the datatype of the column in the primary database. If there is no matching native Replication Server datatype, you must find a datatype definition that matches the primary database datatype.
- When creating subscriptions with where clauses predicated on a column involved in column-level translation, you must specify the predicate value in “declared” format (that is, before translation).

Emulating rs_init activity for a non-Sybase database

The rs_init utility is a Replication Server utility that can add an Adaptive Server database to a replication system, as either a primary or replicate database. When adding a replicate database, rs_init does the following:

- Creates a Replication Server DSI thread connection for the replicate database
- Creates a Replication Server Maintenance User ID for the replicate database

- Creates tables and procedures in the replicate database to support replication

To support the same functionality for a non-Sybase replicate database, Replication Server provides sample scripts. These scripts are located in the `$SYBASE/$SYBASE_RS/scripts` directory.

To use these scripts, you must perform the following steps manually:

1 At the replicate database:

- Create a Maintenance User ID, with permissions required to update and execute all replicate transactions.
- Execute the `hds_XXX_setup_for_replicate.sql` script.

This script creates the `rs_info` and `rs_lastcommit` tables in the replicate database.

This script includes grant statements that need to be changed before execution to reference the correct name of the Maintenance User ID defined in the Replication Server connection. The Maintenance User ID must have update authority to the `rs_info` and `rs_lastcommit` tables.

Note Sybase recommends using the Maintenance User ID to execute this script.

2 At the Replication Server System Database (RSSD):

- Execute the `hds_XXX_udds.sql` script.

This script adds the user-defined datatypes (UDDs), which define the attributes of the replicate database native datatypes to the Replication Server RSSD. These datatype definitions are required to ensure that datatypes received from primary database transactions are properly formatted for application to the replicate database. You may need to modify the script to reference the correct RSSD database name.

- Execute the `hds_XXX_funcstrings.sql` script.

This script replaces several default Replication Server function strings with custom function strings designed to communicate with the replicate database and access the tables and procedures created by the `hds_XXX_setup_for_replicate.sql` script. These function strings are added to the Replication Server `rs_XXX_function_class`. You may need to modify the script to reference the correct RSSD database name.

Note If you use older function strings that were supplied for DB2 in a version of Replication Server prior to 12.0, do *not* run the *hds_db2_setup_for_replicate.sql* script because the existing *rs_db2_function_class* function strings will be replaced.

- Execute the class-level translation scripts.

Class-level translations identify primary database native datatypes and the replicate database native datatypes the data should be translated into (e.g. Oracle DATE should be translated to DB2 TIMESTAMP).

Class-level translations impact Replication Server performance. Only the translations that are needed should be applied to the RSSD. The following scripts are an example for DB2:

- *hds_clt_ase_to_db2.sql* – translations for Adaptive Server datatypes to DB2 datatypes.
- *hds_clt_informix_to_db2.sql* – translations for Informix datatypes to DB2 datatypes.
- *hds_clt_mssql_to_db2.sql* – translations for Microsoft SQL Server datatypes to DB2 datatypes.
- *hds_clt_oracle_to_db2.sql* – translations for Oracle datatypes to DB2 datatypes.

3 At the replicate Replication Server:

- Execute the *hds_XXX_connection_sample.sql* script.

This script provides a template for creating the Replication Server connection for a replicate database, using the default database-specific function-string class provided with Replication Server. Before executing the script, the template must be modified to include your actual replicate data server name, database name, and Maintenance User ID.

Note The Replication Server create connection command uses the *dsi_sql_data_style* option, even for DB2. If you use this option, you effectively turn off datatype translations for that connection. The *dsi_sql_data_type* option remains available in Replication Server for backward compatibility only.

Case sensitivity in publications and subscriptions

In a replication system with only Sybase databases, case sensitivity is not a major issue because Sybase databases all use the same object name case conventions. In Sybase databases, uppercase and lowercase are both allowed, with the default being all lowercase.

Non-Sybase data servers have different object name case conventions. Like Sybase, both Informix and Microsoft SQL Server default to all lowercase. Both DB2 and Oracle default to all uppercase.

Since Replication Server follows the case conventions of Sybase databases, the default for database object names (for example, in replication definitions and subscriptions) is all lowercase.

Replication Server
object names are
case sensitive

Replication Server object names (connections, tables, procedures, columns, subscriptions, and replication definitions) are *case sensitive*. The exact case must be provided whenever you identify or reference those objects.

In a heterogeneous replication system, you must be aware of object name case conventions in all of the components to ensure that the connection attempts from the Replication Agent to the Replication Server, and the LTL generated by Replication Agent, matches the connection names and replication definition object names known to Replication Server.

Sybase Replication Agent uses a configuration parameter (`ltl_character_case`) to specify the case used when communicating with Replication Server. For more information, see the Sybase Replication Agent *Administration Guide*.

RSM object names
are case insensitive

Replication Server Manager (RSM) is case insensitive when processing object names. RSM displays and processes object names in the case presented by the server object being managed. RSM does not attempt to identify or resolve any case sensitivity variances.

In a heterogeneous replication system, it is possible to have a variety of servers with different case sensitivity, or different default display formats. If the primary and replicate data servers store database object names in different cases, you may have to explicitly specify the proper “replicate name” to be used for table and column names when you create replication definitions in Replication Server. The following example illustrates this.

Oracle case sensitivity
example

By default, Oracle processing of object names is case insensitive, however, object names are physically stored in all uppercase. The statement `create table test1 (col1 int)` creates a table named `test1`, which is identified in the Oracle system catalog with name `TABLE1`. The SQL statements `select * from test1` and `select * from TEST1` both return the same result from the same table.

When RSM displays metadata from an Oracle database, all object names are in uppercase (Oracle's default storage format). RSM uses the uppercase object names when it creates a replication definition from an Oracle database.

As Replication Server processes information, it requires the LTL from the Replication Agent to match the character case of the table and column names specified in the replication definition. As the data moves to the replicate site, the commands created by Replication Server to be applied to the replicate database use the case specified in the replication definition for table and column names (for example, insert into TEST1 (COL1) values (1)).

If the replicate database is case sensitive, the commands presented by Replication Server may be rejected with an "object not found" error. In this example, if the replicate database is Adaptive Server (and the Adaptive Server is installed as case-sensitive), and the replicate table was created with the statement create table test1 (col1 int), the insert into TEST1 (COL1) values (1) statement would fail.

RSM allows you to use the replicate as clause in replication definitions so that table and column names are specified in the proper case for the replicate database when you create a replication definition.

Note There is no other option in RSM to control or force the character case of object names.

The LTL generated by the Replication Agent must refer to table TEST1 and column COL1 to be accepted by Replication Server. For the Sybase Replication Agent, the value of the ltl_character_case configuration parameter must be set to asis to send LTL object names in the same case as they are returned by the primary database.

If no other changes are made, an insert into table test1 with a value of "1" results in the following command sent to the replicate database:

```
insert into TEST1 (COL1) values (1)
```

This statement may fail if the replicate table name is case sensitive and the name of the table is not TEST1 with a column named COL1.

Object name case
sensitivity solutions

Ideally, if all the data servers in your replication system are case insensitive, no special consideration is required.

The following items can be used to help in a situation where some of the servers are case sensitive:

- Use the `replicate as` clause in replication definitions. When you create a replication definition, either manually or with RSM, you can specify a separate replicate name for table and column names. Use this feature to specify the case required by your replicate database.

Note If your primary data is replicated to more than one replicate site, you might have to create a separate replication definition for each unique replicate table name.

- Configure the Replication Agent accordingly. If the option is available, you can configure the Replication Agent to send object names in the same case as the primary database stores them. This is particularly important for RSM users. If you are not using RSM to create replication definitions, you can force the case to be anything you like, as long as the LTL generated by the Replication Agent matches the case used in the replication definition.
- Use an alias at the replicate site. As an alternative to creating replication definitions with the `replicate as` clause, you can create an alias at the replicate database that matches the case in the primary database and points to the desired replicate table name.

Implementing Heterogeneous Replication

Chapters in this part describe how to create and maintain a replication system with heterogeneous or non-Sybase data servers, using Sybase replication technology.

- Chapter 6, “Using RSM to Set Up Replication Systems,” describes how to use Replication Server Manager (RSM) to set up and configure replication system components in a heterogeneous replication system.
- Chapter 7, “Replication System Configuration Examples,” gives examples of typical replication system configurations with heterogeneous or non-Sybase data servers.
- Chapter 8, “Administering Heterogeneous Replication Systems,” describes administration tasks that are specific to replication systems with heterogeneous or non-Sybase data servers.
- Chapter 9, “Troubleshooting Heterogeneous Replication Systems,” describes common problems and troubleshooting procedures for replication systems with heterogeneous or non-Sybase data servers.

Using RSM to Set Up Replication Systems

This chapter introduces the Replication Server Manager (RSM) and describes how to use the RSM system management interface to set up the components in a replication system with heterogeneous or non-Sybase data servers.

Topic	Page
Introduction to Replication Server Manager	131
Setting up a heterogeneous replication system	133

Introduction to Replication Server Manager

Replication Server Manager (RSM) consists of:

- RSM Client (Sybase Central and the Replication Server plug-in), an Open Client application that lets you monitor, diagnose, and manage a replication system. The RSM Client includes a graphical user interface (GUI) integrated with Sybase Central.
- RSM Server, an Open Server application that translates and communicates requests from an RSM Client. The RSM Server performs the following tasks:
 - Retrieves metadata from data servers
 - Manages and monitors Replication Agents, Replication Servers, Open Servers, and other RSM servers
 - Monitors and responds to replication system events

RSM relationships

RSM, as a client, can log in to one or more RSM Servers. Each server, in turn, can have a relationship with any data server, Replication Server, and Replication Agent that you want RSM Server to manage.

At least one RSM Server is required. The RSM Server must run somewhere on your network. The RSM Server is available on both UNIX and Microsoft Windows platforms.

If RSM is used to access the system log of a Sybase replication system component (Replication Server, Sybase Replication Agent, or DirectConnect gateway), an RSM Server must have access to the log file from the server where the RSM Server is running.

RSM permissions

You must have the following information for each server to be added to the RSM Server domain:

- The server name
- The type of server
- If you are adding a non-Sybase data server, you need to know:
 - The access server you will use to obtain metadata from the non-Sybase data server (a Replication Agent or a DirectConnect gateway)
 - The name of the server or the name of the Replication Agent that you use to manage the server
 - The user ID and password used to administer the server

Note The user ID must have System Administrator privileges on the server.

- The user ID and password of the RSSD primary user (Replication Server only)
- The user ID and password of the RSSD Database Owner, if it is different from the RSSD primary user (Replication Server only)
- The location of the log file to monitor

RSM connectivity

Before an RSM Server can interact with the servers in a replication system, you must define the RSM Server domain. The RSM Server domain is the set of Replication Servers, data servers, Replication Agents, RSM Servers, and Open Servers that you want an RSM Server to manage.

The servers that you add to a RSM Server domain can be anywhere in your network. This allows you to manage replication systems that are distributed worldwide.

An RSM Server must have an entry in its interfaces file (*interfaces* for UNIX or *sql.ini* for Windows) for every server (data server, Replication Server, Replication Agent, and so on) in its domain.

If the RSM Client and the RSM Server run on different machines, the *sql.ini* file on the RSM Client machine needs entries only for the RSM Servers to which the RSM Client will connect.

Note An RSM Client must have an entity in its interfaces file for every RSM Server that it connects to.

RSM limitations

The following limitations apply to Replication Server Manager version 12.6:

- The RSM Client (Sybase Central) is available only on Microsoft Windows platforms.
- Support for Adaptive Server Anywhere (ASA) as a data source is not provided in RSM. RSM supports ASA as a generic Open Server component, meaning it cannot be managed or configured, but it can be tested to determine if it is active.
- RSM does not manage or monitor the Replication Agent for DB2 UDB for OS/390. RSM can manage a replication system where DB2 for OS/390 is either a primary or replicate data server.
- RSM cannot invoke the Sybase Replication Agent `pdb_xlog` command. For more information about the `pdb_xlog` command, see the Sybase Replication Agent *Administration Guide*.
- Online help is the primary documentation for Replication Server Manager. You can also find RSM information in Chapter 3, “Managing Replication Server with Sybase Central” in the Replication Server *Administration Guide*.

Setting up a heterogeneous replication system

Replication Server Manager (RSM) provides the ability to setup, manage, and monitor replication systems that contain heterogeneous or non-Sybase data servers.

Replication system configuration options

Non-Sybase data servers can be either the source or the target of replication data (that is, primary or replicate databases). To move data from a non-Sybase

data server requires a Replication Agent. To move data to a non-Sybase data server requires a DirectConnect database gateway.

To manage a non-Sybase data server in a replication environment, RSM must retrieve the server's metadata, such as table names, column properties, stored procedure parameters, and so forth. Since RSM cannot directly access a non-Sybase data server, it obtains metadata either through a Replication Agent (for primary data servers), or through a DirectConnect database gateway (for replicate data servers).

Primary-only database connections

In a Sybase-only replication system (with only Adaptive Server database servers), Replication Server communicates directly with all of the data servers using Open Client connectivity. Replication Server database connections are configured as either replicate-only, or primary-and-replicate (bidirectional) in a Sybase-only replication system.

In a replication system with non-Sybase data servers, a primary-only connection is used for a non-Sybase data server that is only a source of data (that is, no data flows out to the non-Sybase data server).

In a primary-only configuration, RSM retrieves the metadata for the non-Sybase data server through the Sybase Replication Agent.

The primary-only solution is implemented in RSM. The Replication Server database connection is set up as a normal primary-and-replicate connection, but RSM filters out the status of the DSI thread and reports only the state of the Rep Agent User thread. The DSI thread still exists in the Replication Server in a suspended state.

Note You cannot use a primary-only configuration to replicate data from DB2 for OS/390. RSM does not manage the Replication Agent for DB2 UDB for OS/390, so the Replication Agent cannot retrieve the DB2 metadata for RSM. For instructions on setting up a DB2 for OS/390 database as a replication source (primary database), see “DB2 Universal Database for OS/390 replication environment” on page 136.

Required servers

The following servers are required to set up a primary-only configuration:

- Non-Sybase primary data server
- Sybase Replication Agent designed for the non-Sybase database
- Replication Server
- Adaptive Server replicate data server
- RSM Server

Primary-only limitations	RSM does not support either a ping cycle or a heartbeat process for a primary-only connection because the DSI thread in the Replication Server is in a suspended state. If either the ping cycle or the heartbeat process is needed, you must set up a bidirectional (primary-and-replicate) configuration.
Replicate-only database connection	<p>A replicate-only database connection is used when a non-Sybase data server is only the target of replicated data (that is, no data flows from the non-Sybase data server). The replicate-only connection is a standard Replication Server replicate database connection. For a replicate-only connection, the Replication Server contains a DSI thread, but no Rep Agent User thread. For a non-Sybase data server, replicate data passes through a DirectConnect database gateway server. RSM uses the DirectConnect server to retrieve metadata from the non-Sybase replicate data server.</p> <p>All non-Sybase data servers that are supported by RSM use the same configuration to set up a replicate-only connection.</p>
Required servers	<p>The following servers are required to set up a replicate-only configuration:</p> <ul style="list-style-type: none">• Adaptive Server primary data server• Replication Server• DirectConnect database gateway server designed for the non-Sybase replicate data server (or MainframeConnect in a gatewayless environment)• Non-Sybase replicate data server• RSM Server
Replicate-only limitations	<p>The following limitations apply to replicate-only connections:</p> <ul style="list-style-type: none">• RSM does not display table or stored procedure metadata for either Adaptive Server or non-Sybase replicate databases.• RSM does not validate Maintenance User permissions at the replicate database when it creates a subscription for a non-Sybase replicate data server.• RSM does not validate the table columns or stored procedure parameters by comparing the primary and replicate databases when it creates a subscription for a non-Sybase replicate data server.
Bidirectional database connection	A bidirectional (or primary-and-replicate) connection is used when a non-Sybase data server acts as both the source and target of replication. For a bidirectional connection, the Replication Server database connection contains both a DSI thread and a Rep Agent User thread. RSM reports the status of both of these threads.

The bidirectional configuration uses a Sybase Replication Agent to retrieve the data from a non-Sybase data server and a DirectConnect database gateway server (or MainframeConnect in a gatewayless environment) to send data to the non-Sybase replicate database. As with the primary-only configuration, RSM uses the Sybase Replication Agent to retrieve metadata from the non-Sybase primary data server.

Note You cannot use this form of a bidirectional connection to replicate data to or from a DB2 for OS/390 database. For instructions on setting up a bidirectional DB2 for OS/390 environment, see “DB2 Universal Database for OS/390 replication environment” on page 136.

Required servers

The following servers are required to set up a bidirectional configuration:

- Non-Sybase data server that acts as both a primary database and a replicate database
- Sybase Replication Agent designed for the non-Sybase database
- Replication Server
- DirectConnect database gateway server designed for the non-Sybase data server (or MainframeConnect in a gatewayless environment)
- RSM Server

Bidirectional limitations

The following limitations apply to bidirectional connections:

- RSM does not validate Maintenance User permissions at the replicate database when it creates a subscription for a non-Sybase replicate data server.
- RSM does not validate the table columns or stored procedure parameters by comparing the primary and replicate databases when it creates a subscription for a non-Sybase replicate data server.

DB2 Universal Database for OS/390 replication environment

A replication environment with DB2 Universal Database for OS/390 (DB2 for OS/390) data servers differs from other non-Sybase replication environments because RSM does not manage the Replication Agent for DB2 UDB for OS/390. Additionally, the Replication Agent cannot retrieve metadata from the DB2 for OS/390 database.

As with other non-Sybase data servers, DB2 for OS/390 has three environment configurations:

- Primary-only
- Replicate-only
- Bidirectional

A DB2 for OS/390 replicate-only environment uses the same server configuration as other non-Sybase, replicate-only configurations. The following sections describe the DB2 for OS/390 primary-only and bidirectional configurations.

DB2 for OS/390 primary-only database connection

A primary-only connection is used when the DB2 for OS/390 database is only the source of replicated data (that is, data does not flow out to the DB2 database). RSM uses a DirectConnect for OS/390 gateway server to retrieve metadata from DB2 for OS/390.

As with other non-Sybase primary-only configurations, the DB2 for OS/390 primary-only solution is implemented in RSM.

The Replication Server database connection is set up as a primary-and-replicate connection, but RSM filters out the status of the DSI thread, and reports only the state of the Rep Agent User thread. The DSI thread still exists in the Replication Server in a suspended state.

Required servers

The following servers are required to set up a DB2 for OS/390 primary-only configuration:

- DB2 Universal Database for OS/390 primary data server
- Replication Agent for DB2 UDB for OS/390
- Replication Server
- Replicate data server
- DirectConnect for OS/390
- RSM Server

DB2 for OS/390 primary-only limitations

The following limitations apply to a DB2 for OS/390 primary-only connection:

- RSM does not support either a ping cycle or a heartbeat process for a primary-only connection because the DSI thread in the Replication Server is in a suspended state. If either the ping cycle or the heartbeat process is needed, set up a bidirectional configuration.

	<ul style="list-style-type: none">• RSM cannot manage or monitor the Replication Agent for DB2 UDB for OS/390. You must set all Replication Agent configuration parameters manually.
DB2 for OS/390 bidirectional database connection	A bidirectional (primary-and-replicate) connection is used when the DB2 for OS/390 database is both a source and target of replication. The DB2 for OS/390 bidirectional connection configuration uses a Replication Agent for DB2 UDB for OS/390 to retrieve data from the DB2 database, and a DirectConnect for OS/390 (or MainframeConnect in a gatewayless environment) to send data to the database. RSM can use DirectConnect for OS/390 to retrieve metadata from the DB2 for OS/390 database.
Required servers	<p>The following servers are required to set up a DB2 for OS/390 bidirectional configuration:</p> <ul style="list-style-type: none">• DB2 Universal Database for OS/390 data server• Replication Agent for DB2 UDB for OS/390• Replication Server• DirectConnect for OS/390 (or MainframeConnect in a gatewayless environment)• RSM Server
DB2 for OS/390 bidirectional limitations	<p>The following limitations apply to a DB2 for OS/390 bidirectional connection:</p> <ul style="list-style-type: none">• RSM does not validate Maintenance User permissions at the replicate database when it creates a subscription for a non-Sybase replicate data server.• RSM does not validate the table columns or stored procedure parameters by comparing the primary and replicate databases when it creates a subscription for a non-Sybase replicate data server.• RSM cannot manage or monitor the Replication Agent for DB2 UDB for OS/390. You must set all Replication Agent configuration parameters manually.

Setting up Replication Server

Replicating non-Sybase data requires unique function strings, datatype definitions, and translations to translate data from one datatype to another.

Note For more detailed information, refer to “Replication Server heterogeneous datatype support” on page 115 and “Emulating rs_init activity for a non-Sybase database” on page 122.

Replication Server stores function strings, datatype definitions, and translations in the RSSD. These objects are loaded in the RSSD using scripts provided by Sybase. The scripts are located in `$$SYBASE/$$SYBASE_RS/scripts` and are named using the following conventions:

- Function-string class scripts – *hds_xxx_funcstrings.sql*
Example: *hds_db2_funcstrings.sql*
- UDD (datatype definition) class scripts – *hds_xxx_udds.sql*
Example: *hds_db2_udds.sql*
- Translation scripts – *hds_clt_xxx_to_xxx.sql*
Example: *hds_clt_db2_to_oracle.sql*

Each non-Sybase data server has a function-string class and a UDD class that contains function strings and UDDs that are used to translate the data into its native format. Additionally, each translation that defines how a non-Sybase datatype is translated to another non-Sybase datatype is associated with a function-string class.

All of these objects—function strings, UDDs, and translations—must be loaded in the RSSD of the correct Replication Server to produce the correct datatype transformations.

The following list describes issues that affect where these objects should be loaded:

- Replicate Replication Servers require:
 - The function-string class (and all associated function strings) for all replicate data servers that it manages.
 - The UDD class (and all associated UDDs) for all replicate data servers that it manages, and for all primary data servers that replicate data to the replicate data servers that it manages. The primary data UDDs are required for the translations.
 - Translations for all primary/replicate data server pairs. For example, for replication between Oracle and DB2, the Oracle to DB2 translations are required.

- Primary Replication Servers require:
 - The function-string class for a bidirectional connection to a primary data server
 - The function-string class and the UDD class to perform materialization.
 - The UDD class for all primary data servers that it manages so that replication definitions can use UDDs in declared datatypes.
 - The UDD class for all replicate data servers, if replication definitions use UDDs in published datatypes.

❖ **To set up Replication Server to translate data for a non-Sybase database**

- 1 Modify each script to issue a `use db` command before any other commands, where `db` is the database name of the RSSD for the Replication Server.

Make a copy of the script if it is needed in more than one RSSD. For example, add the following lines to the beginning of the scripts:

```
use emb
go
```

- 2 Load the function-string class scripts first. From the command prompt enter:

```
isql -Uuid -Ppassword -SRSSD
< hds_dbtype_funcstrings.sql
```

- 3 Load the UDD class scripts next. From the command prompt enter:

```
isql -Uuid -Ppassword -SRSSD
< hds_dbtype_udds.sql
```

- 4 Load the translation scripts last. From the command prompt enter:

```
isql -Uuid -Ppassword -SRSSD
< hds_dbtype_to_dbtype.sql
```

- 5 Shut down and restart the Replication Server.

Note This last step is required. Modifying the RSSD does not refresh the Replication Server memory cache.

Setting up a replicate database with HDS scripts

Replication Server requires two tables named `rs_info` and `rs_lastcommit` at the replicate database. Additionally, for Microsoft SQL Server, Replication Server requires a stored procedure named `rs_update_lastcommit`.

Sybase provides scripts to install and load these objects for each supported non-Sybase data server. These scripts are part of the heterogeneous datatype support (HDS) feature in Replication Server. These scripts are located in the `$SYBASE/$SYBASE_RS/scripts` directory and are named using the following convention: *hds_xxx_setup_for_replicate.sql*

For example: *hds_db2_setup_for_replicate.sql*

❖ To set up a non-Sybase replicate database

- 1 Modify the script to identify the database in which the objects are created. The syntax is dependent on the data server.
- 2 Use the DirectConnect gateway server connected to the replicate database to load the script.

Note The user ID used to load the script must have permissions to create the objects.

Issue the following command:

```
isql -Uid -Ppassword -SDCserver  
< hds_db2_setup_for_replicate.sql
```

Setting up a Replication Agent

This section describes how to set up a Replication Agent within an RSM environment. The Replication Agent server must be added before the non-Sybase data server.

Note RSM does not manage the Replication Agent for DB2 UDB for OS/390.

❖ To set up a Replication Agent

- 1 Identify the data server and database name that will be used by the Replication Agent when sending data to the Replication Server. This name must match the Replication Agent `rs_source_ds` and `rs_source_db`

configuration parameters and the data server and database name specified in the Replication Server database connection.

For replicate-only and bidirectional configurations, the data server name should match the DirectConnect server entry in the Sybase interfaces file.

The database name should match the database name in DB2 Universal Database, Informix, or Microsoft SQL Server. For Oracle servers, the database name does not need to match any particular name.

2 Add the Replication Agent server to RSM.

Select the RSM icon from the Sybase Central desktop, and double-click on the Add Server icon.

The Add Server wizard starts.

3 Select the Replication Agent server name from the drop-down list.

If the Replication Agent is not in the list, cancel the wizard and add the server to the RSM server interfaces file.

Choose Next to continue.

4 Select the Replication Agent type from the drop-down list box.

Choose Next to continue.

5 Enter the user ID and password for the Replication Agent System Administrator.

This information is used by RSM to configure and monitor the Replication Agent.

Choose Next to continue.

6 Enter the data server name and database name you identified in step 1.

Choose Next to continue.

7 Select how the Replication Agent log file will be monitored.

Choose Next to continue.

You can now use RSM to configure the Replication Agent.

8 Set the Replication Agent configuration parameters `rs_source_ds` and `rs_source_db` that identify the data server and database name used by the Replication Server.

These names were identified in step 1.

Setting up a non-Sybase data server with RSM

The RSM creates a “pseudo server” on the Sybase Central desktop to represent a non-Sybase data server. RSM does not communicate directly with non-Sybase data servers, but it collects metadata about the database through either the Sybase Replication Agent or a DirectConnect gateway server. This interface server is identified in the Add Server wizard.

❖ To set up a non-Sybase data server

- 1 Add the non-Sybase data server to the RSM.

Select the RSM icon from the Sybase Central desktop, and double-click on the Add Server icon.

The Add Server wizard starts.

- 2 Select the interface server name from the drop-down list.

If the server is not in the list, cancel the wizard and add the server to the RSM Server interfaces file.

Different connection configurations use different interface servers:

- Replicate-only – DirectConnect gateway server
- DB2 for OS/390 – DirectConnect for OS/390 server
- Primary-only/bidirectional – Replication Agent server

- 3 When adding a non-Sybase server for bidirectional replication, select the Replication Agent from the drop-down list.

Choose Next to continue.

- 4 Select the non-Sybase server type from the drop-down list box.

Choose Next to continue.

- 5 Identify the interface server type: Replication Agent or DirectConnect.

If the server is for a DB2 for OS/390 or replicate-only configuration, check the DirectConnect box. If the server is for a primary-only configuration, check the Replication Agent box. If the server is for a bidirectional configuration, check both boxes.

Choose Next to continue.

- 6 Depending on the selection from the previous dialog box, the Add Server wizard prompts you for a database name.

For primary replication configurations, this name must match the database name used in the Replication Agent `rs_source_db` configuration parameter.

Choose Next to continue.

- 7 If you select a DirectConnect server as the interface server type, the Add Server wizard prompts you for a user ID and password.

This information is used by RSM to monitor the data server. Login information for the Replication Agent is collected when the Replication Agent is added to the RSM environment.

Choose Next to continue.

- 8 If you are adding a primary-only DB2 for OS/390 data server, select the Suspend DSI check box.

This switch tells RSM not to report the status of the DSI thread. Primary-only configurations using Sybase Replication Agent automatically set this switch.

- 9 Open the server Properties dialog box.
- 10 Select the non-Sybase server icon on the Sybase Central desktop, right-click and select Properties from the context menu.
- 11 Open the Communication tab.
- 12 Select the Suspend DSI check box, and choose OK.

Note You must install the Informix schema tables when setting up an Informix replicate data server. For more information, see the *Informix Guide to SQL*.

Creating a database connection

The next step in setting up a replication environment is to add a database connection for the non-Sybase data server.

❖ To create a database connection

- 1 Open the database icon under the non-Sybase data server, and double-click the Add Connection icon.

The Add Connection wizard starts.

- 2 From the drop-down list, select the Replication Server that will manage the database connection.
Choose Next to continue.
- 3 Enter the System Administrator user ID and password for the Replication Server.
Choose Next to continue.
- 4 The next dialog box displays the name of the non-Sybase data server and database.
Choose Next to continue.
- 5 Enter the Maintenance User ID and password for the non-Sybase data server.
Choose Next to continue.
- 6 Select the connection type, either replicate-only or primary-and-replicate (bidirectional).
If you are creating a connection for a primary-only configuration, select Primary and replicate.
Choose Next to continue.
- 7 Modify the database connection parameters.
Non-Sybase database connections require different parameter settings. RSM pre-sets most connections, but you should double-check all required parameters. If the parameters are not set correctly, the connection will not stay in a ready state.
 - Open the connection Properties dialog box.
 - Click once on the newly created database connection on the Sybase Central desktop, right-click on the connection icon, and select Properties from the pop-up menu.
 - On the General tab, verify that the function-string class is set to the non-Sybase function-string class that matches the server type.
 - On the Advanced tab, verify that the Enabled check box in the Batch Replication group is cleared.
 - For an Informix data server, set the Command separator drop-down list (on the Advanced tab) to “;”.
- 8 Choose OK.

Resuming a database connection

The last step to setting up a replication environment is to resume the database connection.

❖ To resume a database connection

- 1 Select the non-Sybase database connection from the Sybase Central desktop, right-click it, and select Resume.
- 2 On the Resume Connection dialog box, select the appropriate check boxes for the type of connection:
 - Replicate-only – select the Resume DSI check box.
 - Primary-only – select the Resume Replication Agent check box.
 - Bidirectional – select both the Resume DSI and the Resume Replication Agent check boxes.
- 3 Choose OK.

The status of the database connection should change in a few moments. If the status does not change, review the Replication Server log for error messages. See Chapter 9, “Troubleshooting Heterogeneous Replication Systems” for more details.

Replication System Configuration Examples

This chapter gives examples of several replication system configurations with heterogeneous or non-Sybase data servers and describes the issues involved with each configuration.

Topic	Page
Non-Sybase primary to Adaptive Server replicate	147
Adaptive Server primary to non-Sybase replicate	148
Non-Sybase primary to non-Sybase replicate	150
Bidirectional non-Sybase to non-Sybase replication	152

Non-Sybase primary to Adaptive Server replicate

The simplest heterogeneous replication scenario is replicating one-way from a non-Sybase primary database to an Adaptive Server replicate database. The only unique requirements are the addition of a Replication Agent designed to extract transaction data from the non-Sybase primary database, and the application of the HDS feature of Replication Server to translate primary database native datatypes to Adaptive Server datatypes.

Replication system components

The following components are required for a non-Sybase primary to Adaptive Server replicate configuration:

- Non-Sybase primary data server (for example, Oracle)
- Replication Agent designed for the primary data server (for example, Sybase Replication Agent)
- Replication Server
- Adaptive Server replicate data server

Replication system issues

You must consider the following issues in a non-Sybase primary to Adaptive Server replicate configuration:

- The Replication Server database connection for the primary database must include a valid user ID and password for the primary database, even though this user ID does not apply transactions to the primary database in this configuration.
- If you use a Replication Server heterogeneous datatype support (HDS) datatype definition to define a column's datatype in a replication definition, the *hds_clt_XXX_to_ase.sql* script must be applied to the RSSD of the replicate Replication Server to provide automatic translation of that primary database datatype to an Adaptive Server native datatype.

If you can use Replication Server native datatypes to represent all the primary column datatypes (that is, no datatype definitions are required), then no HDS translation scripts need be applied.

- Sybase recommends that you do not configure the Replication Agent to perform any datatype translation, unless all replicate databases require the data in that translated form. In that case, it may be a more effective use of resources to let the Replication Agent do the translation once, rather than having replicate Replication Servers do it for each replicate database.

Adaptive Server primary to non-Sybase replicate

A simple heterogeneous replication scenario is replicating one-way from an Adaptive Server primary database to a non-Sybase replicate. The only unique requirements are the addition of a DirectConnect database gateway to apply transaction data to the replicate database, and the application of the HDS feature of Replication Server to translate Adaptive Server datatypes to native datatypes of the replicate database.

Replication system components

The following components are required for an Adaptive Server primary to non-Sybase replicate configuration:

- Adaptive Server primary database

- Replication Server
- DirectConnect database gateway designed for the replicate data server (for example, DirectConnect for Informix)
- Non-Sybase replicate data server (for example, Informix)

Replication system issues

You must consider the following issues in an Adaptive Server primary to non-Sybase replicate configuration:

- The Replication Server database connection for the replicate database must include a valid user ID and password (the Maintenance User) for the replicate database. This user ID must have authority to apply replicate transactions in the replicate database.
- The *hds_XXX_setup_for_replicate.sql* script must be executed against the replicate database to create the *rs_info* and *rs_lastcommit* tables in the replicate database.
- The Replication Server replicate database connection must be created referencing the correct function-string class for this particular replicate database (for example, *rs_informix_function_class*.)
- The *hds_XXX_funcstrings.sql* script must be applied to the RSSD of the replicate Replication Server.
- The *hds_clt_ase_to_XXX.sql* script must be applied to the RSSD of the replicate Replication Server to provide class-level translation of Adaptive Server native datatypes to replicate database native datatypes.

Note Class-level translations are also required to perform translations on system variable values used by Replication Server to maintain the *rs_lastcommit* table in the replicate database.

Non-Sybase primary to non-Sybase replicate

This scenario varies in complexity, depending on the mix of non-Sybase data servers. If the primary and replicate databases are the same type (for example, Oracle to Oracle), fewer class-level translations are required (you do not have to translate an Oracle DATE datatype to another Oracle DATE datatype). If the primary and replicate databases are different types, you must apply additional class-level translations.

Replication system components

The following components are required for a non-Sybase primary to non-Sybase replicate configuration:

- Non-Sybase primary data server (for example, Oracle)
- Replication Agent designed for the primary data server (for example, Sybase Replication Agent)
- Replication Server
- DirectConnect database gateway designed for the replicate data server (for example, DirectConnect for Informix)
- Non-Sybase replicate data server (for example, Informix)

Replication system issues

You must consider the following issues in a non-Sybase primary to non-Sybase replicate configuration:

- The Replication Server primary database connection must include a valid user ID and password for the primary database. This user ID must have authority to apply replicate transactions (even if no transactions will be replicated to the primary database).
- The *hds_XXX_setup_for_replicate.sql* script for the replicate database must be executed against the replicate database to create the *rs_info* and *rs_lastcommit* tables in the replicate database.
- The Replication Server replicate database connection must be created referencing the correct function-string class for this particular replicate database (for example, *rs_informix_function_class*).

- The *hds_xxx_funcstrings.sql* script must be applied to the RSSD of the replicate Replication Server.
- The *hds_clt_ase_to_xxx.sql* script must be applied to the RSSD of the replicate Replication Server to provide class-level translation of Adaptive Server native datatypes to the replicate database native datatypes.

Even if you are not replicating from an Adaptive Server primary database, Adaptive Server native datatypes are used in the default Replication Server function string calls to maintain the *rs_lastcommit* table. In addition, replication definitions for a non-Sybase primary database may use Adaptive Server native datatypes.

- If you use a Replication Server HDS datatype definition to define a column's datatype in a replication definition, the *hds_clt_xxx_to_xxx.sql* script must be applied to the RSSD of the replicate Replication Server to provide automatic translation of that primary database datatype to a replicate database native datatype.

For example, with an Oracle primary database and an Informix replicate database, you must apply the *hds_clt_oracle_to_informix.sql* script to the RSSD of the replicate Replication Server.

If you can use existing Replication Server native datatypes to represent all the columns (that is, no datatype definitions are required), then no HDS translation scripts need be applied.

Note No additional translation script is required, if the same type of non-Sybase database is used for both primary and replicate databases (for example, Informix to Informix).

Bidirectional non-Sybase to non-Sybase replication

In this scenario, replication occurs both to and from each database. Each non-Sybase database must have both a Replication Agent *and* a DirectConnect database gateway.

Replication system components

The following components are required for a bidirectional non-Sybase primary to non-Sybase replicate configuration:

- Non-Sybase primary data server (for example, Oracle)
- Replication Agent designed for the primary data server (for example, Sybase Replication Agent)
- DirectConnect database gateway designed for the “primary” data server acting as a replicate database (for example, DirectConnect for Oracle)
- Replication Server
- DirectConnect database gateway designed for the replicate data server (for example, DirectConnect for Informix)
- Replication Agent designed for the “replicate” data server acting as a primary database (for example, Sybase Replication Agent)
- Non-Sybase replicate data server (for example, Informix)

Replication system issues

From a technical standpoint, you can set up a bidirectional replication scenario using only two Replication Server database connections (one “primary-and-replicate” connection for each database). However, this document describes four connections, one for each database acting as a primary database, and one for each database acting as a replicate database.

Note In the following description of bidirectional replication issues, the two databases are referred to as Database #1 and Database #2, since both databases take on both “primary” and “replicate” roles in the replication system.

You must consider the following issues in a bidirectional non-Sybase primary to non-Sybase replicate configuration:

- The Replication Server primary database connection for Database #1 must include a valid user ID and password for the primary database. This user ID must be the same user ID specified in the Replication Server replicate database connection for Database #1 (the Maintenance User). This user ID must have authority to apply transaction operations to replicate tables in Database #1.
- The Replication Agent for Database #1 must be configured to bypass Maintenance User transactions to prevent a transaction from returning from the replicate tables in Database #1. Refer to the appropriate Replication Agent *Administration Guide* for details on configuring the Replication Agent to bypass Maintenance User transactions.
- The Replication Server primary database connection for Database #2 must include a valid user ID and password for the primary database. This user ID must be the same user ID specified in the Replication Server replicate database connection for Database #2 (the Maintenance User). This user ID must have authority to apply transaction operations to replicate tables in Database #2.
- The Replication Agent for Database #2 must be configured to bypass Maintenance User transactions to prevent a transaction from returning from the replicate tables in Database #2. Refer to the appropriate Replication Agent *Administration Guide* for details on configuring the Replication Agent to bypass Maintenance User transactions.
- The *hds_XXX_setup_for_replicate.sql* script for each database must be executed against the database to create the *rs_info* and *rs_lastcommit* tables that Replication Server requires for a replicate database.
- The Replication Server database connection for replicate Database #1 must be created referencing the correct function-string class designed for this particular database as a replicate database (for example, *rs_oracle_function_class*.)
- The Replication Server database connection for replicate Database #2 must be created referencing the correct function-string class designed for this particular database as a replicate database (for example, *rs_informix_function_class*.)
- The *hds_XXX_funcstrings.sql* script for each database must be applied to the RSSD of the replicate Replication Server for each database.
- The *hds_clt_ase_to_XXX.sql* script must be applied to the RSSD of each replicate Replication Server to provide class-level translation of Adaptive Server native datatypes to replicate database native datatypes.

Even if you are not replicating from an Adaptive Server primary database, Adaptive Server native datatypes are used in the default Replication Server function-string calls to maintain the `rs_lastcommit` table. In addition, replication definitions for a non-Sybase primary database may use Adaptive Server native datatypes.

- If you use a Replication Server HDS datatype definition to define a column's datatype in a replication definition (for either Database #1 or Database #2), the *hds_clt_XXX_to_YYY.sql* script must be applied to the RSSD of the replicate Replication Server to provide automatic translation of the primary database datatype to a replicate database native datatype.

For example, you must apply the *hds_clt_informix_to_oracle.sql* script to the RSSD of the replicate Replication Server for Database #1. You must apply the *hds_clt_oracle_to_informix.sql* script to the RSSD of the replicate Replication Server for Database #2.

If you can use existing Replication Server native datatypes to represent all the columns (that is, no datatype definitions are required), then no HDS translation scripts need be applied.

Note No additional translation script is required if the same type of non-Sybase database is used for both primary and replicate databases (for example, Informix to Informix).

Administering Heterogeneous Replication Systems

This chapter describes administration tasks for Sybase replication systems with heterogeneous or non-Sybase data servers.

Topic	Page
Replication system maintenance	155
Replication system monitoring	157

Note This chapter describes only administration tasks that are unique to a Sybase replication system with heterogeneous or non-Sybase data servers. For information about basic replication system administration, see the *Replication Server Administration Guide*.

Replication system maintenance

This section describes the tasks required to maintain these components of a Sybase replication system:

- Primary data servers and Replication Agents
- Replication Server
- Replicate data servers and DirectConnect database gateways

Primary data servers and Replication Agents

In a Sybase replication system, the typical ongoing maintenance requirement at a primary database is monitoring and adjusting transaction log resources. All Replication Agents use some kind of log resource, either the native transaction log of the database, or database objects

created and administered by the Replication Agent. In either case, an increase in transaction volume can increase the demand for log resources.

You should review unusual events for their impact on the transaction log. A one-time or periodic maintenance operation, which can impact a large number of database rows, can have a severe impact on transaction log and Replication Agent performance, if only temporarily.

Replication Server

Replication Server may suffer a minor performance impact when it has to process class-level translations, which provide automatic datatype translation for a replicate database.

If you remove a non-Sybase database from a Replication Server (that is, *all* connections associated with that database are removed), you can remove the class-level translations supporting that database if there are no other databases of the same type that are managed by the Replication Server (for example, you remove an Oracle replicate database and there are no other Oracle databases managed by that Replication Server).

At a primary Replication Server, you can remove the translations for a specific database type only if no primary database connection exists for that database type and no replication definitions use a map to clause referencing a user-defined datatype for that database type.

Note Use caution when removing class-level translations. You must verify that for a particular Replication Server, *no* database connection for that database type exists at that Replication Server.

Replicate data servers and DirectConnect gateways

DirectConnect database gateways do not require any special maintenance; however, your resource or capacity planning at the replicate database site should include an estimate of the quantity of replicate transactions expected to be received by the replicate database from all primary databases to which the replicate subscribes.

Replication system monitoring

Replication Server Manager (RSM) provides heartbeat and latency functionality to support a replication system with heterogeneous or non-Sybase data servers. The non-Sybase data servers supported by RSM monitoring are:

- DB2 Universal Database
- Informix Dynamic Server
- Microsoft SQL Server
- Oracle

Replication system monitoring overview

The term *latency* refers to the time difference between when a transaction is completed at the primary database and when it is completed at the replicate database. A *heartbeat* is used to establish a data flow between the primary and replicate databases to measure system latency. In a replication system that already has data movement between primary and replicate databases, a heartbeat is not required to measure latency.

You can start heartbeats and latency graphs from the Sybase Central desktop (RSM Client) by right-clicking a connection listed under a Replication Server, or by right-clicking a database entry listed under a data server and selecting Heartbeats or Latency.

RSM Server heartbeat feature

The heartbeat functionality in RSM Server consists of creating a table in the primary database, defining a replication definition with customized function strings on the table, and creating a subscription to the replication definition. Once the subscription is marked as valid, RSM Server populates the table and updates it according to the *heartbeat interval* specified (every 60 seconds, by default).

Each time the table is updated, data is replicated to the replicate database. The function strings defined on the replication definition change all inserts, updates, and deletes into commits, which become null operations on the replicate database. Since no data is actually inserted on the replicate there is no need for a replicate heartbeat table.

RSM Server latency calculations

Latency is calculated by reading the `rs_lastcommit` table on the replicate database. Although a heartbeat does not insert data into a replicate database, it does cause the `rs_lastcommit` table to be updated, allowing latency from the primary database to the replicate database to be determined.

Changing the RSM Server heartbeat

You can modify, stop, or remove an RSM Server heartbeat process after it is started. To modify the heartbeat, change the Heartbeat interval and select Apply.

The Stop button instructs RSM Server to delete the appropriate row in the heartbeat table, and stop the heartbeat thread that updates the table. Remove performs the same action as Stop, but it also removes the replication definition and subscription that were created for the heartbeat table. The table itself is not deleted.

Monitoring non-Sybase heartbeat and latency

When you start a heartbeat, you must enter a user ID and password. The user ID must have authority to create the heartbeat table in the primary database. The user ID and password must also be added to the Replication Server and given create object permission.

The heterogeneous database specific function strings, user-defined datatypes (UDDs), and any necessary class-level translations should be loaded into the Replication Server RSSD.

For example, before setting up a heartbeat between Adaptive Server and Informix, run the following Replication Server scripts against the Replication Server RSSD:

- *hds_informix_funcstrings.sql*
- *hds_informix_udds.sql*
- *hds_clt_ase_to_informix.sql*
- *hds_clt_informix_to_ase.sql*

Note Because function-string values are cached, you must restart the Replication Server for the new settings to take effect.

Database connections for non-Sybase data servers are composed of two parts:

- The data server name, which will be the name of a Sybase DirectConnect service
- The database name

For more information, see “Database connections” on page 105.

For Oracle and Informix, the database name is arbitrary, although it must match the value of the Replication Agent `rs_source_db` configuration parameter.

For Microsoft SQL Server, the database name refers to an actual database, similar to Adaptive Server.

When starting a heartbeat with a Microsoft SQL Server primary database, you should not use the `sa` user ID. The owner of tables created by `sa` is `dbo`. The heartbeat table is created, but RSM Server looks for a table owned by `sa` and does not find it. You can add `dbo` as a Replication Server user ID with `create object` permission, however, and use that user ID to start the heartbeat.

For DB2 Universal Database, the database name is also the name of an actual database, but the concept of a database in DB2 is different than it is for Adaptive Server. The heartbeat table is created in the database that is named in the database connection for DB2, so the user ID used to start the heartbeat must have `CREATE TABLE` permission in that database.

For example, if the database connection is `DB2_directconnect.my_database`, the `CREATE TABLE` command is sent to DB2 through the DirectConnect gateway in the form:

```
CREATE TABLE RSM_HEARTBEAT (...) IN DATABASE MY_DATABASE
```

The user ID used to start the heartbeat is the owner of the heartbeat table. The Maintenance User ID defined on that connection, however, owns the `RS_LASTCOMMIT` table in DB2. This table is read by RSM Server to determine latency when DB2 is the replicate database. The Maintenance User ID is retrieved automatically from the Replication Server, so there is no need to specify it when you want to check latency. Checking latency does not require a user ID because no tables are created or modified.

You must specify the time difference between the RSM Server and a non-Sybase data server to get meaningful latency values. You can enter this value under the Monitoring tab of the data server's Properties dialog box. If the non-Sybase data server's clock time is ahead of the RSM Server time clock, enter the difference in seconds as a positive number. Otherwise, enter the difference as a negative number. Remember to take into account time zone differences, if any.

Troubleshooting non-Sybase heartbeat and latency

To trace the heartbeat processes in RSM Server, activate the trace option RSML HBEATS. Use the trace options EVE GENERAL_1 and EVE GENERAL_2 to trace the latency process.

In Replication Server, you can use the DSI_BUF_DUMP trace flag to display the commands sent to the DSI connection, which is the DirectConnect database gateway for a non-Sybase data server. You do not need the diagnostic version of Replication Server to use this trace flag. Refer to the *Replication Server Troubleshooting Guide* for more information about Replication Server traces.

For latency calculations to work correctly, the DirectConnect gateways used to access non-Sybase databases must have the value of the DateTimeResults configuration parameter set to datetime. In addition, any Sybase Replication Agent must have the value of the ltl_origin_time_required configuration parameter set to true.

Removing heartbeat objects when the heartbeat is not running

The first time a heartbeat is started, a replication definition and a subscription are created. They remain defined unless you click Remove in the Heartbeat dialog box while the heartbeat is running.

❖ **To remove the replication definition and subscription when the heartbeat is not running**

- 1 From the Sybase Central desktop, right-click a Replication Server icon and choose Properties.
- 2 Select the Catalog tab of the Replication Server Properties dialog box.
- 3 Select Subscriptions as the object type and find the subscription named rsm_heartbeat.
- 4 Delete the subscription by right-clicking the name and choosing Delete.
- 5 In the Drop Subscription Information dialog box, select Without Purge in the Method box.
- 6 Choose OK.
- 7 Once the subscription is removed, select Replication Definitions from the Object type box and find the heartbeat replication definition named heartbeat_XXX, where XXX is the 3-digit database ID of the primary database.
- 8 Right-click the replication definition name and choose Delete.

- 9 In the Database Connection Properties dialog box, the function-string class should be set to the appropriate type (for example, `rs_informix_function_class` for an Informix database).
- 10 Under the Advanced tab of the Database Connection Properties dialog box:
 - Verify that Command separator is set to “;” for Informix database connections.
 - Under Batch Replication, verify that the Enabled check box is unselected for all non-Sybase data servers.

Troubleshooting Heterogeneous Replication Systems

This chapter describes common problems and troubleshooting procedures for Sybase replication systems with heterogeneous or non-Sybase data servers.

Topic	Page
Troubleshooting overview	163
Inbound and outbound queue problems	164
Replication failure problems	168
Troubleshooting specific errors	171

Troubleshooting overview

Common 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, are described in the *Replication Server Troubleshooting Guide*.

For non-Sybase primary and replicate databases, the Replication Agent and DirectConnect gateway documentation provide troubleshooting information for each specific database.

This chapter describes some basic troubleshooting you can do if replication fails in a Sybase replication system with heterogeneous or non-Sybase data servers.

Inbound and outbound queue problems

The *inbound queue* is where Replication Server stores the data it receives from a primary database (through a Replication Agent or another Replication Server). The *outbound queue* is where Replication Server stores the data it needs to send to a replicate site (either a replicate database or another Replication Server).

This section describes how to troubleshoot problems with these Replication Server queues.

Inbound queue problems

You can tell that the Replication Server inbound queue for a primary database is not being updated if you issue the Replication Server `admin who,sqm` command at the primary Replication Server and the results indicate that:

- The number of blocks being written in the Replication Server inbound queue for the connection in question is not changing.
- The number of duplicate messages being detected is not increasing.

❖ To determine the reason the inbound queue is not being updated

- 1 Verify the Replication Server connection Rep Agent User thread status.

You can issue an `admin who` command in the primary Replication Server to review the status of the Rep Agent User thread for the Replication Server database connection in question.

- If there is no Rep Agent User thread for the connection, the connection was not created with the `with log transfer on` clause. You can alter the Replication Server database connection to turn log transfer processing on, if needed.
- If the Rep Agent User thread status is down, the Replication Agent is not actively connected to the Replication Server. A down status is typical for Replication Agents that only connect to Replication Server when there is work to be sent, and then disconnect after a period of inactivity.

- 2 Verify that the expected Replication Agent is executing.

Verify that the expected Replication Agent is active, and that the values of the Replication Agent `rs_source_ds` and `rs_source_db` configuration parameters match the desired Replication Server connection name.

Refer to the appropriate Replication Agent documentation for other tests to validate that the Replication Agent is executing.

- 3 Verify that the expected table or procedure is marked for replication.

In RSM, the properties of the primary object should have the Replicate check box selected.

Replication Agent documentation describes the Replication Agent commands you can use to check replication status.

Sybase Replication Agent provides for separate *enabling* of replication, in addition to marking. In this case, make sure the marked object is also enabled for replication.

- 4 Verify that the Replication Agent is scanning new records.

If the database object is marked for replication, the log scanning process of the Replication Agent should record that additional information is being scanned.

To verify that new records are being scanned:

- 1 Start tracing in the Replication Agent.
- 2 Update or execute a primary database object that has been marked for replication.
- 3 Verify that scanning occurs.

Refer to the appropriate Replication Agent documentation to determine the trace flags you can use to validate the scanning process.

- 5 Verify that logging occurs for the primary object.

For Replication Agents that use the primary database transaction or recovery log, this check is not required.

For Replication Agents that use a trigger-based transaction log (Sybase Replication Agent for Informix, Microsoft SQL Server, and Oracle), validate that the log is being populated.

You can use the following procedure to validate that the Sybase Replication Agent trigger-based log is being updated:

- 1 Shut down (suspend) the Replication Agent log scanning process.
- 2 Verify that there is no other activity against marked objects.
- 3 To get a count of the current records in the transaction log, issue the statement

```
select count(*) from ra_tran_log_
```

where *ra_* is the value of the Replication Agent *pdb_xlog_prefix* configuration parameter.

- 4 Update or execute a primary database object that has been marked for replication.

- 5 Issue the same `select count(*)` statement.

If the row count does not change, the transaction was not logged. If the row count increases, logging is occurring.

If the row count decreases, a log truncation process is executing and you must retry the test after the truncation process is complete.

- 6 Verify that the transactions to be replicated are not being executed by the Replication Server Maintenance User.

Transactions being processed by the Maintenance User defined in the Replication Server connection for this database are typically bypassed by the Replication Agent.

For some Replication Agents, bypassing Maintenance User transactions is optional.

Refer to the appropriate Replication Agent *Administration Guide* for details.

Outbound queue problems

You can tell that the Replication Server outbound queue for a replicate database is not being updated if you issue the Replication Server `admin who,sqm` command at the replicate Replication Server and the results indicate that:

- The number of blocks being written in the Replication Server outbound queue for the connection in question is not changing.
- The number of duplicate messages being detected is not increasing.

Problems between inbound and outbound queues are often naming problems.

❖ To determine the reason the outbound queue is not being updated

- 1 Verify that any Replication Server routes are active.

Refer to the Replication Server *Troubleshooting Guide* for route validation techniques between primary and replicate Replication Servers.

- 2 Verify that the Replication Server connection DSI thread is not down.

Issue an admin who command in the replicate Replication Server to review the status of the DSI thread for the Replication Server connection.

If the DSI thread status is down, the Replication Server is not connected to the replicate database (or DirectConnect gateway). Review the Replication Server log for errors and attempt to resume the connection.

- 3 Verify that the DSI thread connection is not in “Loss Detected” mode by viewing the replicate Replication Server log for “Loss Detected” messages for the DSI thread in question.

When Replication Server detects a loss, no further messages are accepted on the DSI thread connection.

Refer to the Replication Server *Administration Guide* for information about recovering from this error.

- 4 Verify the primary replication definition.

The primary Replication Server inbound queue can receive data, but not apply that data to any replication definition. This error typically occurs when the name of the replication definition does not match the name presented in the Log Transfer Language (LTL) that was created by the Replication Agent. This becomes more likely when you are using different non-Sybase database types with different default character cases.

Replication Server processing of replication commands is case-sensitive. In a replication system with non-Sybase data servers, you must ensure that the LTL generated by Replication Agents matches the Replication Server connection names and replication definition object names.

Some Replication Agents always use lowercase names when they communicate with Replication Server (for example, Adaptive Server and DB2 Universal Database).

Sybase Replication Agent provides a configurable option to specify the character case. The best option is to pick one character case (uppercase or lowercase) and use it consistently with all Replication Server connections, replication definitions, and subscription names.

The character case validation process is manual. You can use the `rs_helprep` command to verify the name of a replication definition. You can then turn on LTL tracing in the Replication Agent and verify that the name provided in the LTL trace matches the spelling and character case of the name specified in the replication definition.

If the character case appears to be incorrect, review the Replication Agent documentation to verify the default character case settings and any possible configuration changes. If a name is misspelled, you must delete and then re-create the replication definition.

Replication failure problems

The most common replication failure problems in a replication system with heterogeneous or non-Sybase data servers are:

- Trigger-based logging failures in the Replication Agent
- Replicate database fails to be updated

This section describes how to troubleshoot these problems.

Trigger-based logging failures

Sybase Replication Agent uses a trigger-based log solution for Informix, Microsoft SQL Server, and Oracle databases. Use the following procedure to isolate the problem if an object marked for replication fails to update the log.

Note Replication Agents for DB2 Universal Database (on either OS/390 platforms, or UNIX, Windows NT, or Windows 2000 platforms) are not trigger-based.

❖ To find the problem when a trigger-based log fails to update

- 1 Verify that no exceptions have been logged.

The Replication Agent can log exceptions to an exception table (the default name is `ra_exception`) in the primary database.

Select the contents of the exception table. The table should be empty or contain a single empty entry. If the table is not empty, use the information in the exception to determine the problem.

- 2 Check the primary database log for error messages.
- 3 Verify that error messages are returned to the client application.

It may be that an error message is being generated, but the client application is suppressing it.

- 4 Attempt the command as a different user.

If logging occurs using a different user ID, the problem may be related to database permissions.

- 5 Check any related primary database system tables used to record error information.

For example, in Oracle, the `all_errors` system table can contain information about triggers and procedures that have been invalidated.

Check the database error table(s) for entries with object names that start with the same prefix as Replication Agent transaction log objects, that is, the value of the `pdb_xlog_prefix` configuration parameter.)

- 6 Verify that triggers or procedures are enabled.

Some primary databases allow triggers or procedures to be enabled or disabled at the system level.

Replicate database is not updated

If the Replication Server outbound queue is being updated, but transaction data is not being applied at the replicate database, use the following procedure to determine the reason.

❖ To determine why replicate transactions are not applied at the replicate database

- 1 Determine if the subscription contains a where clause.

Verify that the transaction data expected will pass any where clause in the subscription definition. Use the `rs_helpsub` stored procedure to list the text of the subscription.

- 2 Verify HDS installation.

If you are using Replication Server HDS to support replication to or from a non-Sybase data server, verify that the HDS scripts have been properly applied.

See “Expected datatype translations do not occur” on page 172 for more information.

- 3 Verify that the `rs_lastcommit` table is set up correctly.

If you are using Replication Server HDS to support replication to or from a non-Sybase data server, verify that the HDS scripts have been properly applied.

Refer to “Updates to rs_lastcommit fail” on page 171 for more detail.

- 4 Review the replicate Replication Server log for errors.
- 5 Review the replicate database log for errors.
- 6 Verify manual access to replicate objects.

Log in to the replicate database (or DirectConnect gateway) using the Replication Server connection Maintenance User ID, and verify that you have update authority to the replicate table or procedure.

- 7 Validate commands sent to the replicate database:
 - Turn on the DSI_BUF_DUMP trace in the replicate Replication Server and record to the Replication Server log the commands being sent to the replicate database.
 - Verify that these commands, when manually applied, produce the expected results.

Note The DSI_BUF_DUMP trace flag can be used with any Replication Server. (By contrast, the similar DSI_CMD_DUMP trace flag is available only with the diagnostic version of Replication Server.) Refer to the Replication Server *Troubleshooting Guide* for more information about the Replication Server traces.

- 8 Turn on tracing at the DirectConnect gateway to see what commands are being received.

For example, these parameters in a DirectConnect for Oracle configuration file cause DirectConnect to write additional information to the *DCO.log* file:

- network_tracing = 1
- traces = 1,2,3,4,5,6,10

Refer to the appropriate DirectConnect documentation for specific trace availability and syntax.

Troubleshooting specific errors

This section describes troubleshooting for specific errors you may encounter in a Sybase replication system with heterogeneous or non-Sybase data servers.

Date information does not include time values

This problem can occur with non-Sybase primary databases. Replication Agents have configuration parameters that must be set to request that the native date or time datatypes in the primary database be converted to Sybase datetime formats for replication. If this conversion is not requested, the default date format returned from the primary database is used, which may not provide a time value. For example, the Oracle default date display format is *DD-MON-YY*. If the Replication Agent is not configured to convert this display to a format that includes time (for example, *CCYY-MM-DD hh.mm.ss.mmm*), then the time value will be all zeroes whenever the date is replicated.

See the Replication Agent for DB2 UDB for OS/390 *User's and Troubleshooting Guide* for more information about the `Date_in_char` parameter. See the Sybase Replication Agent *Administration Guide* for more information about the `pdb_convert_datetime` parameter.

When using a Replication Server datatype for a non-Sybase data server, you must ensure that the date datatype specified in a replication definition map to clause includes a time value.

Note The Sybase Replication Agent `pdb_convert_datetime` configuration parameter must be set *before* a table or procedure is marked for replication in an Informix, Microsoft SQL Server, or Oracle database. The trigger code performs the translation, and the Replication Agent uses that parameter value to determine how to create the trigger.

Updates to `rs_lastcommit` fail

When replicating into a non-Sybase replicate database, the replicate Replication Server updates the `rs_lastcommit` table as soon as the connection is resumed. If the replicate Replication Server error log displays a syntax error while updating the `rs_lastcommit` table, the following procedure may help identify the problem.

❖ **To troubleshoot rs_lastcommit update failure**

- 1 Verify that the table exists in the replicate database, and the schema matches the script in the *hds_xxx_setup_for_replicate.sql* file.

Note For a non-Sybase database, you should create the rs_lastcommit table by executing the *hds_xxx_setup_for_replicate.sql* script supplied with Replication Server.

- 2 Verify access authority.

Log in to the replicate database using the Replication Server Maintenance User ID and password specified in the create connection command for that database connection.

Verify that this user ID can update the rs_lastcommit table. You should be able to insert and delete a dummy entry without error.

- 3 Trace the actual command.

Turn on tracing in the replicate Replication Server (DSI_BUF_DUMP trace) or in the DirectConnect gateway and resume the Replication Server connection.

Identify the failing statement and correct as necessary.

Expected datatype translations do not occur

The most common reason for a datatype translation failure is an incomplete installation of the necessary user-defined datatypes (UDDs) and translations. You can use the following procedure to validate your UDD and translation installation.

❖ **To validate UDD and translation installation**

- 1 Reboot the Replication Servers. Replication Server caches all function-string information at start-up.

Subsequent changes to the function strings stored in the RSSD do not take effect until the Replication Server is restarted.

- 2 Verify that UDDs have been defined.

The Replication Server script *hds_xxx_udds.sql* provides the SQL statements necessary to apply UDDs to the RSSD for specific non-Sybase replicate databases.

For each datatype, the script issues a delete followed by an insert. You can re-run these scripts without failure.

These scripts must be executed at each Replication Server where:

- The non-Sybase data server is a replicate database, and
- The non-Sybase data server primary replication definition refers to a UDD using the map_to clause.

Verify that your copy of the scripts has been updated with the correct use statement for the database name of the RSSD.

- 3 Verify that class-level translations have been applied to the replicate Replication Server.

The Replication Server script *hds_clt_from_to_xxx.sql* provides the SQL statements necessary to apply class-level translations to the RSSD of the replicate Replication Server for a specific combination of non-Sybase primary databases to non-Sybase replicate databases.

Note The *hds_clt_ase_to_xxx.sql* is required for any non-Sybase replicate database. For example, if you are replicating from Informix to Oracle, the *hds_clt_ase_to_oracle.sql* translations must be applied to ensure Replication Server updates to the rs_lastcommit table are properly translated and applied to the replicate database.

You can re-run these scripts without failure. Verify that your copy of the scripts has been updated with the correct use statement for the database name of the RSSD.

Note Class-level translations adversely affect performance. Do not install translations that are not required for your replication system.

- 4 Verify that your replicate database Replication Server connection is associated with the appropriate function-string class.

To take advantage of class-level translations, the replicate Replication Server connection must use the correct non-Sybase function-string class.

You can use the Replication Server rs_helpdb command to determine which function-string class is defined for a database connection.

Function-string classes for replicate databases are:

- Adaptive Server Enterprise – *rs_sqlserver_function_class*

- DB2 Universal Database on OS/390 platforms – *rs_db2_function_class*
- DB2 Universal Database on UNIX, Windows NT, and Windows 2000 platforms – *rs_udb_function_class*
- Informix – *rs_informix_function_class*
- Microsoft SQL Server – *rs_msss_function_class*
- Oracle – *rs_oracle_function_class*

You can use the Replication Server `admin show_function_classes` command to display a list of active function-string classes.

You can use the Replication Server `alter connection` command to change the function-string class of an existing database connection.

- 5 Verify that the non-Sybase function-string classes have been updated with appropriate function strings.

Replication Server script *hds_XXX_funcstrings.sql* provides the SQL statements necessary to apply function strings to the RSSD of the replicate Replication Server for a specific non-Sybase replicate database.

For each function string, the scripts issue a delete followed by an insert. You can re-run these scripts without failure.

Verify that your copy of the scripts has been updated with the correct use statement for the database name of the RSSD.

- 6 Use the Replication Server `admin translate` command.

The `admin translate` command allows you to verify the results of a specific translation.

Use this command to verify that the translation engine is providing the translation results you expect.

Refer to the Replication Server *Administration Guide* for more information on heterogeneous datatype support (HDS) and the `admin translate` command.

Replication Agent triggers do not fire

The Sybase Replication Agent uses a trigger-based approach to logging transactions for Informix, Microsoft SQL Server, and Oracle. If you determine the triggers are not firing (in other words, no entries are being made to the

ra_tran_log_ table), the following procedure can help you determine where the problem is.

❖ **To troubleshoot Replication Agent trigger failure**

- 1 Verify marked and enabled status of primary objects.

Use the Replication Agent `pdb_setreptable, table_name` command (or `pdb_setrepproc, proc_name` for a stored procedure) to display the marked and enabled status for the object in the primary database.

The status for an individual table or procedure must be true for both marked and enabled attributes.

Refer to the Sybase Replication Agent *Administration Guide* for more information on these commands and how to change the marked and enabled status.

- 2 Verify access to primary database objects.

Normally, if a permission violation occurs, you should see an access error message when attempting to update the primary table or execute the primary procedure.

Use a standard client access tool (for example, `isql`, `sqlplus`, or `dbaccess`) to issue a simple insert or execute to the primary object and check for any error messages. Make sure your client access tool is configured to display all messages.

- 3 Verify that triggers or procedures are not disabled.

Most data servers provide some means to disable a trigger or stored procedure.

- 4 Check the Replication Agent exceptions table.

The Sybase Replication Agent can be configured to record errors to an exception table and bypass logging until the error is corrected.

In the primary database, query the `ra_exception_` table for entries. If you find any, resolve the error(s) and clear the exception table.

- 5 Check the primary database system log for any error messages.

Check any vendor specific system error tables (for example, the Oracle `ALL_ERRORS` table may contain errors on triggers and stored procedures).

- 6 Manually execute the trigger or procedure calls.

You can find the code executed by the trigger or stored procedure by examining the marking script that was created when the primary object was marked for replication.

The object-marking scripts are located in the *mark.sql* file in the *\$SYBASE/rax-12_6/instance_name/scripts/object_name/installed/* subdirectory.

Refer to the Sybase Replication Agent *Administration Guide* for more detailed information about the marking scripts and their location.

Determining log scan position

The Sybase Replication Agent transaction log table (*ra_tran_log_*) contains a column that identifies when a particular operation has occurred.

The following sample queries can help identify how these tables can be used to count the number of operations not yet processed by the Replication Agent:

- For an Informix database, use:

```
select count(*) from ra_tran_log_ where process_seq  
< 1
```

- For a Microsoft SQL Server database, use:

```
select count(*) from ra_tran_log_ (NOLOCK) WHERE  
tran_id NOT IN (SELECT ra_id FROM ra_proc_seq_  
(NOLOCK))
```

- For an Oracle database, use:

```
SELECT COUNT(*) from RA_TRAN_LOG_ where PROCESS_SEQ  
< 1
```

LTL generation and tracing

This section describes how to trace the Log Transfer Language (LTL) commands sent to a primary Replication Server, as well as other significant Replication Agent traces.

Replication Agent for DB2 UDB for OS/390

You can use the configuration parameters described in this section to obtain additional information that is not normally presented by Replication Agent for DB2 UDB.

Note There is usually some performance impact when you use these parameters. You should review the full description of a parameter in the Replication Agent for DB2 UDB *Installation Guide* for OS/390 before using it.

- If you need additional tracing to help debug the information passed to a Replication Agent user exit, you can set the value of the `API_com_test` configuration parameter to Y. You can also use this trace when no exit is being used.
- The `LTL_test_only` configuration parameter controls whether LTM for MVS connects to Replication Server and sends transaction operations for replication. When the value of the `LTL_test_only` parameter is Y, LTL that would normally be sent to Replication Server is written to the *LTLOUT* file, instead.

Note No replication takes place when the value of the `LTL_test_only` parameter is Y.

- The `Suppress_col_names` configuration parameter determines whether LTM for MVS suppresses column names from the LTL that is sent to Replication Server.

If you are tracing LTL output, setting the value of this parameter to N ensures column names are present in the generated LTL.

- The `trace=1,4` configuration parameter traces Sybase Log Extract calls to the Replication API, and sends the output to *SYSPRINT*. Set the value of the `trace=1,4` parameter to Y to turn on this trace.
- The `trace=1,11` configuration parameter writes LTL that is passed to Replication Server to *LTLOUT*. Set the value of the `trace=1,11` parameter to Y to turn on this trace.

Note The `trace=1,11` trace option continues to send LTL to the primary Replication Server, whereas the `LTL_test_only` option does not.

- The Use_repdef configuration parameter allows LTM for MVS to send LTL to Replication Server that contains only the columns specified in the replication definition.

Setting the value of the Use_repdef parameter to N may increase the amount of information provided in an LTL trace.

Sybase Replication Agent

You can use the trace flags and configuration parameters described in this section to obtain additional information that is not normally presented by the Sybase Replication Agent.

Note There is usually some performance impact when you use these trace flags and parameters. You should review the full description of a trace flag or parameter in the Sybase Replication Agent *Administration Guide* before using it.

Trace flags

Normal trace output is sent to the Replication Agent trace log file. You can use the log_trace_name command to find the name and path of the Replication Agent trace log file.

Output from the LTITRACELTL trace point is sent to a separate LTL output log file (*LTITRACELTL.log*).

The following trace flags are particularly useful for troubleshooting Sybase Replication Agent problems:

- LRTRACE – traces general execution of the Log Reader component
- LTITRACE – traces general execution of the Log Transfer Interface component
- LTITRACELTL – enables LTL statement tracing in the *LTITRACELTL.log* file
- RACONTRC – traces connection and query execution
- RACONTRCSQL – traces SQL statements sent to the primary database

Configuration parameters

The settings of the following Sybase Replication Agent configuration parameters affect the trace information:

- compress_ltl_syntax – provides more verbose description of LTL commands when set to false

- `connect_to_rs` – allows LTL to be generated without actual connection or sending information to Replication Server when set to false
- `log_trace_verbose` – provides more verbose description of traced components when set to true
- `use_rssd` – provides a complete generation of LTL commands without modification for replication definition information when set to false
- `column_compression` – sends complete column information (all columns in after images) in the generated LTL for update operations when set to false

For a complete description of Sybase Replication Agent trace flags and configuration settings, refer to the Sybase Replication Agent *Administration Guide*.

Appendixes

Appendixes in this part provide supplemental information that can help you design and implement a replication system with heterogeneous or non-Sybase data servers.

- Appendix A, “Datatype Translation and Mapping,” lists the class-level datatype translations for non-Sybase data servers supported by Replication Server, and the Replication Server datatype names for non-Sybase datatypes.
- Appendix B, “Materialization Issues,” describes the issues related materializing subscriptions in a replication system with heterogeneous or non-Sybase data servers.
- Appendix C, “Heterogeneous Database Reconciliation,” describes the issues related to synchronizing and reconciling databases in a replication system with heterogeneous or non-Sybase data servers.
- Appendix D, “Replication with Adaptive Server Anywhere,” describes primary and replicate data server issues for Adaptive Server Anywhere in a Sybase replication system.

Datatype Translation and Mapping

This appendix lists the class-level datatype translations for non-Sybase data servers supported by Replication Server, and the Replication Server datatype names for non-Sybase datatypes.

Name	Page
DB2 datatypes	184
Informix datatypes	187
Microsoft SQL Server datatypes	190
Oracle datatypes	192

For each supported non-Sybase data server, Replication Server provides class-level translations that define the default mapping from one datatype to another. Translations are provided for:

- Non-Sybase datatypes that do not correspond directly to Adaptive Server datatypes
- Adaptive Server datatypes that do not correspond directly to the non-Sybase datatypes
- Non-Sybase datatypes that do not correspond directly to the datatypes of another supported non-Sybase data server

Note Class-level translations are *not* provided for any datatype that corresponds directly to a datatype in another data server.

DB2 datatypes

This section lists class-level translations (default datatype mapping) for DB2 datatypes and Replication Server datatype names for DB2 datatypes.

Note This information about datatype translation applies to DB2 Universal Database (UDB) in either mainframe environments (such as OS/390 or Z/OS), or UNIX or Microsoft Windows environments.

DB2 class-level translations

The following sections list the class-level translations for Adaptive Server datatypes to DB2 datatypes, DB2 datatypes to Adaptive Server datatypes, and DB2 datatypes to datatypes of other supported non-Sybase data servers:

- “Adaptive Server to DB2 datatypes”
- “DB2 to Adaptive Server datatypes”
- “DB2 to Informix datatypes”
- “DB2 to Microsoft SQL Server datatypes”
- “DB2 to Oracle datatypes”

Adaptive Server to
DB2 datatypes

Table A-1 lists class-level translations from Adaptive Server datatypes to DB2 datatypes.

Table A-1: Class-level translation from Adaptive Server to DB2 datatypes

Adaptive Server datatype	DB2 datatype
binary	CHAR FOR BIT DATA
bit	TINYINT
date	DATE (UNIX and Windows only)
datetime	TIMESTAMP
decimal	DECIMAL
money	NUMERIC
numeric	NUMERIC
real	REAL (UNIX and Windows only)
smalldatetime	TIMESTAMP
smallmoney	NUMERIC
time	TIME (UNIX and Windows only)

Adaptive Server datatype	DB2 datatype
varbinary	VARCHAR FOR BIT DATA

DB2 to Adaptive
Server datatypes

Table A-2 lists class-level translations from DB2 datatypes to Adaptive Server datatypes.

Table A-2: Class-level translation from DB2 to Adaptive Server datatypes

DB2 datatype	Adaptive Server datatype
CHAR FOR BIT DATA	binary
DATE	datetime
DOUBLE (UNIX and Windows only)	float
REAL (UNIX and Windows only)	real
TIME	datetime
TIMESTAMP	datetime
VARCHAR FOR BIT DATA	varbinary

DB2 to Informix
datatypes

Table A-3 lists class-level translations from DB2 datatypes to Informix datatypes.

Table A-3: Class-level translation from DB2 to Informix datatypes

DB2 datatype	Informix datatype
CHAR FOR BIT DATA	binary
DATE	date
DOUBLE (UNIX and Windows only)	float
REAL (UNIX and Windows only)	real
TIME	datetime (time only)
TIMESTAMP	datetime fraction(5)
VARCHAR FOR BIT DATA	binary

DB2 to Microsoft SQL
Server datatypes

Table A-4 lists class-level translations from DB2 datatypes to Microsoft SQL Server datatypes.

Table A-4: Class-level translation from DB2 to Microsoft SQL Server datatypes

DB2 datatype	Microsoft SQL Server datatype
CHAR FOR BIT DATA	binary
DATE	datetime
DOUBLE (UNIX and Windows only)	float
REAL (UNIX and Windows only)	real
TIME	datetime

DB2 to Oracle
datatypes

DB2 datatype	Microsoft SQL Server datatype
TIMESTAMP	datetime
VARCHAR FOR BIT DATA	varbinary

Table A-5 lists class-level translations from DB2 datatypes to Oracle datatypes.

Table A-5: Class-level translation from DB2 to Oracle datatypes

DB2 datatype	Oracle datatype
CHAR FOR BIT DATA	RAW
DATE	DATE
DOUBLE (UNIX and Windows only)	FLOAT
REAL (UNIX and Windows only)	REAL
TIME	DATE (with time)
TIMESTAMP	DATE (with time)
VARCHAR FOR BIT DATA	RAW

Replication Server datatype names for DB2

Table A-6 lists the Replication Server user-defined datatype (UDD) names that identify DB2 datatypes for DB2 data servers on OS/390 platforms.

Table A-6: Replication Server names for DB2 OS/390 datatypes

DB2 OS/390 datatype	Replication Server name
CHAR FOR BIT DATA	rs_db2_char_for_bit
DATE	rs_db2_date
DECIMAL	rs_db2_decimal, rs_db2_numeric
TIME	rs_db2_time
TIMESTAMP	rs_db2_timestamp
VARCHAR FOR BIT DATA	rs_db2_varchar_for_bit

Table A-7 lists the Replication Server user-defined datatype (UDD) names that identify DB2 datatypes for DB2 data servers on UNIX and Microsoft Windows platforms.

Table A-7: Replication Server names for DB2 UNIX and Windows datatypes

DB2 UNIX and Windows datatype	Replication Server name
CHAR FOR BIT DATA	rs_udb_char_for_bit
DATE	rs_udb_date

DB2 UNIX and Windows datatype	Replication Server name
DOUBLE	rs_udb_double
INTEGER	rs_udb_bigint
REAL	rs_udb_real
TIME	rs_udb_time
TIMESTAMP	rs_udb_timestamp
VARCHAR FOR BIT DATA	rs_udb_varchar_for_bit

Informix datatypes

This section lists class-level translations (default datatype mapping) for Informix datatypes and Replication Server datatype names for Informix datatypes.

Informix class-level translations

The following sections list class-level translations for Adaptive Server datatypes to Informix datatypes, Informix datatypes to Adaptive Server datatypes, and Informix datatypes to datatypes of other supported non-Sybase data servers:

- “Adaptive Server to Informix datatypes”
- “Informix to Adaptive Server datatypes”
- “Informix to DB2 Datatypes”
- “Informix to Microsoft SQL Server datatypes”
- “Informix to Oracle datatypes”

Adaptive Server to
Informix datatypes

Table A-8 lists class-level translations from Adaptive Server datatypes to Informix datatypes.

Table A-8: Class-level translation from Adaptive Server to Informix datatypes

Adaptive Server datatype	Informix datatype
binary	binary
date	date
datetime	datetime fraction(3)

Adaptive Server datatype	Informix datatype
smalldatetime	datetime fraction(0)
time	datetime (time only)
varbinary	binary

Informix to Adaptive
Server datatypes

Table A-9 lists class-level translations from Informix datatypes to Adaptive Server datatypes.

Table A-9: Class-level translation from Informix to Adaptive Server datatypes

Informix datatype	Adaptive Server datatype
binary	binary
date	datetime
datetime fraction(0)	datetime
datetime fraction(1)	datetime
datetime fraction(2)	datetime
datetime fraction(3)	datetime
datetime fraction(4)	datetime
datetime fraction(5)	datetime
datetime (time only)	datetime

Informix to DB2
Datatypes

Table A-10 lists class-level translations from Informix datatypes to DB2 datatypes.

Table A-10: Class-level translation from Informix to DB2 datatypes

Informix datatype	DB2 datatype
binary	CHAR FOR BIT DATA
date	DATE
datetime fraction(0)	TIMESTAMP
datetime fraction(1)	TIMESTAMP
datetime fraction(2)	TIMESTAMP
datetime fraction(3)	TIMESTAMP
datetime fraction(4)	TIMESTAMP
datetime fraction(5)	TIMESTAMP
datetime (time only)	TIME

Informix to Microsoft
SQL Server datatypes

Table A-11 lists class-level translations from Informix datatypes to Microsoft SQL Server datatypes.

Table A-11: Class-level translation from Informix to Microsoft SQL Server datatypes

Informix datatype	Microsoft SQL Server datatype
binary	binary
date	datetime
datetime fraction(0)	datetime
datetime fraction(1)	datetime
datetime fraction(2)	datetime
datetime fraction(3)	datetime
datetime fraction(4)	datetime
datetime fraction(5)	datetime
datetime (time only)	datetime

Informix to Oracle
datatypes

Table A-12 lists class-level translations from Informix datatypes to Oracle datatypes.

Table A-12: Class-level translation from Informix to Oracle datatypes

Informix datatype	Oracle datatype
binary	RAW
date	DATE
datetime fraction(0)	DATE (with time)
datetime fraction(1)	DATE (with time)
datetime fraction(2)	DATE (with time)
datetime fraction(3)	DATE (with time)
datetime fraction(4)	DATE (with time)
datetime fraction(5)	DATE (with time)
datetime (time only)	DATE (with time)

Replication Server datatype names for Informix

Table A-13 lists the Replication Server user-defined datatype (UDD) names that identify Informix datatypes.

Table A-13: Replication Server names for Informix datatypes

Informix datatype	Replication Server name
binary	rs_informix_binary
date	rs_informix_date
datetime fraction(0)	rs_informix_datetime0
datetime fraction(1)	rs_informix_datetime1

Informix datatype	Replication Server name
datetime fraction(2)	rs_informix_datetime2
datetime fraction(3)	rs_informix_datetime3
datetime fraction(4)	rs_informix_datetime4
datetime fraction(5)	rs_informix_datetime5
datetime (time only)	rs_informix_datetime_time
integer	rs_informix_int8
integer	rs_informix_serial8

Microsoft SQL Server datatypes

This section lists class-level translations (default datatype mapping) for Microsoft SQL Server datatypes and Replication Server datatype names for Microsoft SQL Server datatypes.

Microsoft SQL Server class-level translations

The following sections list class-level translations for Microsoft SQL Server datatypes to datatypes of other supported non-Sybase data servers:

- “Microsoft SQL Server to DB2 datatypes”
- “Microsoft SQL Server to Informix datatypes”
- “Microsoft SQL Server to Oracle datatypes”

Note Class-level translations are not supplied for Adaptive Server datatypes to Microsoft SQL Server datatypes (or Microsoft SQL Server datatypes to Adaptive Server datatypes) because Microsoft SQL Server datatypes are directly compatible with Adaptive Server datatypes and they require no translation.

Microsoft SQL Server
to DB2 datatypes

Table A-14 lists class-level translations from Microsoft SQL Server datatypes to DB2 datatypes.

Table A-14: Class-level translation from Microsoft SQL Server to DB2 datatypes

Microsoft SQL Server datatype	DB2 datatype
binary	CHAR FOR BIT DATA
bit	TINYINT
datetime	TIMESTAMP
decimal	DECIMAL
money	NUMERIC
numeric	NUMERIC
smalldatetime	TIMESTAMP
smallmoney	NUMERIC
varbinary	VARCHAR FOR BIT DATA

Microsoft SQL Server
to Informix datatypes

Table A-15 lists class-level translations from Microsoft SQL Server datatypes to Informix datatypes.

Table A-15: Class-level translation from Microsoft SQL Server to Informix datatypes

Microsoft SQL Server datatype	Informix datatype
binary	binary
datetime	datetime fraction(3)
smalldatetime	datetime fraction(0)
varbinary	binary

Microsoft SQL Server
to Oracle datatypes

Table A-16 lists class-level translations from Microsoft SQL Server datatypes to Oracle datatypes.

Table A-16: Class-level translation from Microsoft SQL Server to Oracle datatypes

Microsoft SQL Server datatype	Oracle datatype
binary	RAW
datetime	DATE (with time)
money	DECIMAL
smalldatetime	DATE
smallmoney	DECIMAL
varbinary	RAW

Replication Server datatype names for Microsoft SQL Server

All Microsoft SQL Server datatypes are compatible with the corresponding Adaptive Server datatypes. Only one Microsoft SQL Server datatype has a user-defined datatype definition.

Table A-17 lists the Replication Server user-defined datatype (UDD) name that identifies a Microsoft SQL Server datatype.

Table A-17: Replication Server names for Microsoft SQL Server datatypes

Microsoft SQL Server datatype	Replication Server name
integer	rs_mssql_bigint

Oracle datatypes

This section lists class-level translations (default datatype mapping) for Oracle datatypes and Replication Server datatype names for Oracle datatypes.

Oracle class-level translations

The following sections list class-level translations for Adaptive Server datatypes to Oracle datatypes, Oracle datatypes to Adaptive Server datatypes, and Oracle datatypes to datatypes of other supported non-Sybase data servers:

- “Adaptive Server to Oracle datatypes”
- “Oracle to Adaptive Server Datatypes”
- “Oracle to DB2 datatypes”
- “Oracle to Informix datatypes”
- “Oracle to Microsoft SQL Server datatypes”

Adaptive Server to
Oracle datatypes

Table A-18 lists class-level translations from Adaptive Server datatypes to Oracle datatypes.

Table A-18: Class-level translation from Adaptive Server to Oracle datatypes

Adaptive Server datatype	Oracle datatype
binary	RAW
date	DATE

Adaptive Server datatype	Oracle datatype
datetime	DATE (with time)
money	DECIMAL
smalldatetime	DATE
smallmoney	DECIMAL
time	DATE (with time)
varbinary	RAW

Oracle to Adaptive
Server Datatypes

Table A-19 lists class-level translations from Oracle datatypes to Adaptive Server datatypes.

Table A-19: Class-level translation from Oracle to Adaptive Server datatypes

Oracle datatype	Adaptive Server datatype
RAW	varbinary
DATE	datetime
TIMESTAMP (0)	datetime
TIMESTAMP (1)	datetime
TIMESTAMP (2)	datetime
TIMESTAMP (3)	datetime
TIMESTAMP (4)	datetime
TIMESTAMP (5)	datetime
TIMESTAMP (6)	datetime
TIMESTAMP (7)	datetime
TIMESTAMP (8)	datetime
TIMESTAMP (9)	datetime

Oracle to DB2
datatypes

Table A-20 lists class-level translations from Oracle datatypes to DB2 datatypes.

Table A-20: Class-level translation from Oracle to DB2 datatypes

Oracle datatype	DB2 datatype
RAW	CHAR FOR BIT DATA
DATE	DATE
DATE (with time)	TIMESTAMP
FLOAT	DOUBLE (UNIX and Windows only)
INTEGER	INTEGER (UNIX and Windows only)
TIMESTAMP (0)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (1)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (2)	TIMESTAMP (UNIX and Windows only)

Oracle datatype	DB2 datatype
TIMESTAMP (3)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (4)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (5)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (6)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (7)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (8)	TIMESTAMP (UNIX and Windows only)
TIMESTAMP (9)	TIMESTAMP (UNIX and Windows only)

Oracle to Informix datatypes

Table A-21 lists class-level translations from Oracle datatypes to Informix datatypes.

Table A-21: Class-level translation from Oracle to Informix datatypes

Oracle datatype	Informix datatype
RAW	binary
DATE	date
DATE (with time)	datetime fraction(0)
TIMESTAMP (0)	datetime fraction(0)
TIMESTAMP (1)	datetime fraction(1)
TIMESTAMP (2)	datetime fraction(2)
TIMESTAMP (3)	datetime fraction(3)
TIMESTAMP (4)	datetime fraction(4)
TIMESTAMP (5)	datetime fraction(5)
TIMESTAMP (6)	datetime fraction(5)
TIMESTAMP (7)	datetime fraction(5)
TIMESTAMP (8)	datetime fraction(5)
TIMESTAMP (9)	datetime fraction(5)

Oracle to Microsoft SQL Server datatypes

Table A-22 lists class-level translations from Oracle datatypes to Microsoft SQL Server datatypes.

Table A-22: Class-level translation from Oracle to Microsoft SQL Server datatypes

Oracle datatype	Microsoft SQL Server datatype
RAW	varbinary
DATE	datetime
TIMESTAMP (0)	datetime
TIMESTAMP (1)	datetime
TIMESTAMP (2)	datetime
TIMESTAMP (3)	datetime

Oracle datatype	Microsoft SQL Server datatype
TIMESTAMP (4)	datetime
TIMESTAMP (5)	datetime
TIMESTAMP (6)	datetime
TIMESTAMP (7)	datetime
TIMESTAMP (8)	datetime
TIMESTAMP (9)	datetime

Replication Server datatype names for Oracle

Table A-23 lists the Replication Server user-defined datatype (UDD) names that identify Oracle datatypes.

Table A-23: Replication Server names for Oracle datatypes

Oracle datatype	Replication Server name
RAW	rs_oracle_binary
DATE	rs_oracle_date
INTEGER	rs_oracle_int
NUMBER	rs_oracle_float
NUMBER	rs_oracle_decimal
TIMESTAMP	rs_oracle_datetime

Materialization Issues

This appendix describes the subscription materialization issues that you must consider when implementing a replication system with heterogeneous or non-Sybase data servers. It also describes how to materialize subscriptions to primary tables in a non-Sybase database.

Topic	Page
Materialization overview	197
Unloading data from a primary database	199
Datatype translation issues	199
Loading data into replicate databases	201
Atomic bulk materialization	202
Nonatomic bulk materialization	205

Materialization overview

Materialization is a process of creating and activating subscriptions and copying data from a primary database to a replicate database, thereby initializing the replicate database.

Before you can replicate data from a primary database, you must set up and populate each replicate database so that the replicate objects (tables and so forth) are in a state consistent with those in the primary database.

Types of materialization

Two types of subscription materialization are supported by Replication Server:

- *Bulk materialization* – manually creating and activating a subscription and populating a replicate database using data unload and load utilities outside the control of the replication system.
- *Automatic materialization* – creating a subscription and populating a replicate database using Replication Server commands.

For more information about subscription materialization methods, see the *Replication Server Administration Guide*.

Heterogeneous materialization issues

To materialize subscriptions to primary data in a non-Sybase data server, you must use a bulk materialization method.

With bulk materialization methods, you must coordinate and manually perform the following activities:

- Define, activate, and validate the subscription (or create the subscription without materialization).
- Unload the subscription data at the primary database.
- Move the unloaded data to the replicate database site.
- Load the primary data into the replicate database tables.
- Resume the database connection from the replicate Replication Server to the replicate data server so that the replicate database can receive replicated transactions.
- Resume replication at the Replication Agent instance.

Bulk materialization options

There are two bulk materialization options for subscriptions to primary data in a non-Sybase database:

- *Atomic bulk materialization*
 - Stop updates to the primary table and dump the subscription data from the primary database.
 - In the replicate Replication Server, define the subscription.
 - In the primary database, use the `rs_marker` stored procedure to activate the subscription using the `with suspension` option.
 - Load the subscription data into the replicate table.
 - Resume the database connection from the replicate Replication Server to the replicate database.
 - In the replicate Replication Server, validate the subscription.
- *Nonatomic bulk materialization*
 - In the replicate Replication Server, use the `set autocorrection` command.
 - In the replicate Replication Server, define the subscription.
 - In the primary database, use the `rs_marker` stored procedure to activate the subscription using the `with suspension` option.
 - Dump the subscription data from the primary database.

-
- In the primary database, use the `rs_marker` stored procedure to validate the subscription.
 - Load the subscription data into the replicate table.
 - Resume the database connection from the replicate Replication Server to the replicate database.
 - When the subscription becomes valid at all Replication Servers, turn off autocorrection.

Unloading data from a primary database

Part of the subscription materialization process involves unloading the subscription data from the primary table so it can be loaded into the replicate table. *Subscription data* is the data in the primary table that is requested by the subscription.

Data unloading utilities are usually provided with data server software. You can use one of the OEM-supplied data unloading utilities, or a database unload utility of your choice.

Note Once subscription data is unloaded from a primary database, you may need to perform datatype translation on the unloaded data before loading the data into the replicate database. For more information, see “Datatype translation issues” on page 199.

Datatype translation issues

If you use the heterogeneous datatype support (HDS) feature of Replication Server to perform either column- or class-level translations on replicated data, you must perform datatype translations on the subscription data you unload from the primary database for materialization:

- If any column-level translation is specified in a replication definition for the data, you must perform the datatype translation specified in each replication definition for the primary data.

- If class-level translations are specified for the database connection to a replicate database (in the function-string class for the database connection), you must perform the datatype translations specified for each replicate database connection.

Note When Replication Server processes data for replication, column-level translations are performed before class-level translations. Therefore, if you must perform datatype translations on subscription data based on both column-level translation and class-level translations, you must perform the column-level translation first.

Datatype definitions are specified in scripts provided with Replication Server. These scripts define the attributes of the non-Sybase database native datatypes in terms of Replication Server native datatypes. You can use the datatype definitions in these scripts to determine how to translate the datatypes in the subscription data that you unload from the primary database.

These are the scripts that support non-Sybase data servers:

- *hds_db2_udds.sql* – DB2 Universal Database on OS/390 platforms
- *hds_udb_udds.sql* – DB2 Universal Database on UNIX, Windows NT, and Windows 2000 platforms
- *hds_informix_udds.sql* – Informix
- *hds_mssql_udds.sql* – Microsoft SQL Server
- *hds_oracle_udds.sql* – Oracle

Class-level translations identify primary database native datatypes and the replicate database native datatypes that data should be translated into (for example, Oracle DATE should be translated to DB2 TIMESTAMP).

A set of scripts is provided with Replication Server for each combination of supported primary and replicate data servers. For example, the following scripts are provided for DB2 for OS/390 as a replicate database:

- *hds_clt_ase_to_db2.sql* – translations for Adaptive Server datatypes to DB2 datatypes
- *hds_clt_informix_to_db2.sql* – translations for Informix datatypes to DB2 datatypes
- *hds_clt_mssql_to_db2.sql* – translations for Microsoft SQL Server datatypes to DB2 datatypes

-
- *hds_clt_oracle_to_db2.sql* – translations for Oracle datatypes to DB2 datatypes

You can use the datatype translations specified in these scripts to determine how to translate the datatypes in the subscription data that you unload from the primary database.

Note Datatype translations based on class-level translation scripts must be performed *after* column-level translations.

For more information about Replication Server HDS datatype translations, see “Replication Server heterogeneous datatype support” on page 115.

Loading data into replicate databases

Part of the subscription materialization process involves loading the subscription data from the primary table into the replicate table.

Note After subscription data is unloaded from a primary database, you may need to perform datatype translation on the unloaded data before loading the data into the replicate database.

If you are using Adaptive Server Enterprise as the data server for the replicate database, you can use the Sybase bcp utility to load subscription data into the replicate database.

If you are using a non-Sybase data server as the data server for the replicate database, you can use the load utility of your choice to load subscription data into the replicate database.

For more information about using bcp with Adaptive Server, see Sybase Adaptive Server Enterprise *Utility Guide*.

Atomic bulk materialization

Atomic bulk materialization assumes that all applications updating the primary table can be suspended while a copy of the table is made. The copy is then loaded at the replicate site.

You can use this method to retrieve data from the primary database if you can (at least temporarily) suspend updates to the primary data.

Preparing for materialization

Before you start an atomic bulk materialization procedure, verify the following:

- The primary table exists and contains data.
- You have access to a user ID with ownership or select privilege on the primary table (or a column to be replicated in the primary table).
- The replicate table exists and contains the appropriate column(s) with the appropriate datatype(s).
- You have successfully configured all Replication Servers in your replication system.
- You have created the replication definition correctly at the primary Replication Server.
- If you are using Sybase Replication Agent for a DB2 Universal Database, Informix, Microsoft SQL Server, or Oracle primary database:
 - You have successfully created the Replication Agent transaction log objects in the primary database.
 - You have marked and enabled replication for the primary table in the primary database.
 - You have started the Replication Agent instance and put it in the *Replicating* state.

Performing atomic bulk materialization

❖ To perform atomic bulk materialization

- 1 Use `isql` to log in to the replicate Replication Server as the System Administrator (`sa`):

```
isql -Usa -Psa_password -SRRS_servername
```

where *sa* is the System Administrator user ID, *sa_password* is the password for the System Administrator user ID, and *RRS_servername* is the server name of the replicate Replication Server.

- 2 Define the subscription at the replicate Replication Server:

```
1> define subscription subscription_name
2> for replication_definition
3> with replicate at dataserver.database
4> [where search_conditions]
5> go
```

The `dataserver.database` must match the Replication Server connection name you use for the replicate database.

- 3 Check the subscription at both the primary and replicate Replication Servers. Use the following command to verify that the subscription status is `DEFINED`:

```
1> check subscription subscription_name
2> for replication_definition
3> with replicate at dataserver.database
4> go
```

- 4 Lock the primary table to prevent primary transaction activity. This prevents updates to the primary table during materialization.
- 5 Unload the subscription data at the primary site using your site's preferred database unload method to select or dump the data from the primary table.

Note When unloading subscription data from the primary table, make sure you select only the columns specified in the replication definition and the rows specified in the subscription.

- 6 Perform any datatype translations necessary for the subscription data.
If any column-level translation is specified in the replication definition for this data, you must perform the datatype translation specified in the replication definition.

If class-level translations are specified for the subscription, you must perform the datatype translations specified for the subscription.

- 7 Activate the subscription using the `with suspension` option at the replicate Replication Server:

```
1> activate subscription subscription_name
2> for replication_definition
3> with replicate at dataserver.database
4> with suspension
5> go
```

- 8 Wait for the subscription to become active at both the primary and replicate Replication Servers. Execute the `check subscription` command at both the primary and replicate Replication Servers to verify that the subscription status is `ACTIVE`.

When the subscription status is `ACTIVE` at the replicate Replication Server, the database connection for the replicate database is suspended.

- 9 Restore the primary table to read-write access (`unlock`).
- 10 Load the subscription data into the replicate database using the `bcp` utility or your site's preferred database load utility.
- 11 From the replicate Replication Server, resume the database connection for the replicate database:

```
1> resume connection
2> to dataserver.database
3> go
```

- 12 Validate the subscription at the replicate Replication Server:

```
1> validate subscription subscription_name
2> for replication_definition
3> with replicate at dataserver.database
4> go
```

- 13 Wait for the subscription to become valid at both the primary and replicate Replication Servers, then execute the `check subscription` command at both the primary and replicate Replication Servers to verify that the status is `VALID`.

When you complete this procedure, the subscription is created, the replicate data is consistent with the primary data, and replication is in progress.

If replication is not in progress when you complete this procedure, see Chapter 9, "Troubleshooting Heterogeneous Replication Systems."

See also

- Replication Server *Reference Manual* for information on Replication Command Language (RCL) commands.
- Replication Server *Administration Guide* for information on configuring Replication Servers and materialization methods.

Nonatomic bulk materialization

Nonatomic bulk materialization assumes applications updating the primary table cannot be suspended while a copy of the table is made. Therefore, nonatomic materialization requires the use of the Replication Server *autocorrection* feature to get the replicate database synchronized with the primary database.

Note You cannot use nonatomic materialization if the replicate minimal columns feature is set for the replication definition for the primary table.

Preparing for materialization

Before you start a nonatomic bulk materialization procedure, verify the following:

- The primary table exists and contains data.
- You have access to a user ID with ownership or select privilege on the primary table (or a column to be replicated in the primary table).
- The replicate table exists and contains the appropriate column(s).
- You have successfully configured all Replication Servers in your replication system.
- You created the replication definition correctly at the primary Replication Server.
- If you are using Sybase Replication Agent for a DB2 Universal Database, Informix, Microsoft SQL Server, or Oracle primary database:
 - You have successfully created the Replication Agent transaction log objects in the primary database.

- You have marked and enabled replication for the primary table in the primary database.
- You have started the Replication Agent instance and put it in the *Replicating* state.

Performing nonatomic bulk materialization

❖ To perform nonatomic bulk materialization

- 1 Use `isql` to log in to the replicate Replication Server as the System Administrator (`sa`):

```
isql -Usa -Psa_password -SRRS_servername
```

where *sa* is the System Administrator user ID, *sa_password* is the password for the System Administrator user ID, and *RRS_servername* is the server name of the replicate Replication Server.

- 2 Turn on the autocorrection feature at the replicate Replication Server:

```
1> set autocorrection on
2> for replication_definition
3> with replicate at dataserver.database
4> go
```

- 3 Define the subscription using the `with suspension` option at the replicate Replication Server:

```
1> define subscription subscription_name
2> for replication_definition
3> with replicate at dataserver.database
4> with suspension
5> go
```

The `dataserver.database` must match the Replication Server connection name you use for the replicate database.

- 4 In the primary database, invoke the `rs_marker` stored procedure to activate the subscription.
- 5 Check the subscription at both the primary and replicate Replication Servers. Use the following command to verify that the subscription status is `ACTIVE`:

```
1> check subscription subscription_name
2> for replication_definition
3> with replicate at dataserver.database
4> go
```

When the subscription status is ACTIVE at the replicate Replication Server, the database connection for the replicate database is suspended.

- 6 Unload the subscription data at the primary site using your site's preferred database unload method to select or dump the data from the primary tables.

Note When unloading subscription data from the primary table, make sure you select only the columns specified in the replication definition and the rows specified in the subscription.

- 7 Perform any datatype translations necessary for the subscription data.

If any column-level translation is specified in the replication definition for this data, you must perform the datatype translation specified in the replication definition.

If class-level translations are specified for the subscription, you must perform the datatype translations specified for the subscription.
- 8 In the primary database, invoke the `rs_marker` stored procedure to validate the subscription.
- 9 Wait for the subscription to become valid at both the primary and replicate Replication Servers, then execute the `check subscription` command at both the primary and replicate Replication Servers to verify that the status is VALID.
- 10 Load the subscription data into the replicate database using the `bcp` utility or your site's preferred database load utility.
- 11 From the replicate Replication Server, resume the database connection for the replicate database:

```
1> resume connection
2> to dataserver.database
3> go
```
- 12 Wait for the subscription to become valid at both the primary and replicate Replication Servers, then execute the `check subscription` command at both the primary and replicate Replication Servers to verify that the status is VALID.

When the subscription's status is VALID at the replicate Replication Server, the replicate database is synchronized with the primary database and you can turn off autocorrection.

13 Turn off the autocorrection feature at the replicate Replication Server:

```
1> set autocorrection off
2> for replication_definition
3> with replicate at dataserver.database
4> go
```

When you complete this procedure, the subscription is created, the replicate data is consistent with the primary data, and replication is in progress.

If replication is not in progress when you complete this procedure, see Chapter 9, “Troubleshooting Heterogeneous Replication Systems.”

See also

- Replication Server *Reference Manual* for information on Replication Command Language (RCL) commands.
- Replication Server *Administration Guide* for information on configuring Replication Servers and materialization methods.

Heterogeneous Database Reconciliation

This appendix describes the issues involved with comparing and reconciling data from different databases in a heterogeneous replication system.

Topic	Page
Sybase rs_subcmp utility	209
Database comparison application	210

Sybase rs_subcmp utility

The Sybase rs_subcmp utility allows you to compare primary and replicate tables in Adaptive Server databases, and reconcile any differences found. Sybase provides the rs_subcmp executable program with Replication Server.

Some other database vendors may provide a similar “compare” utility that can perform the same function for their own databases, but there is no equivalent utility to support different types of non-Sybase data servers (for example, comparing tables in an Oracle database to tables in a Microsoft SQL Server database).

For non-Sybase database support, the only options are:

- Acquire third-party tools that provide such functionality, or
- Build your own application.

Database comparison application

You can develop a custom application to perform the same functions as the `rs_subcmp` utility. The application's complexity depends on the number of different data server types, the complexity of the tables to be compared, the amount of data translation involved, and so forth.

The following list describes the major issues that a database comparison application must accommodate to be successful in a heterogeneous replication environment:

- **Connectivity** – the application must be able to communicate with both the primary and replicate databases. If multiple database vendors are involved, ODBC and JDBC protocols can provide a common interface and functionality.
- **Sort order** – the default sort order may be different for different databases. The application may need to force the sort order to improve comparison performance.
- **Character sets** – some primary and replicate databases may store character data in different character sets. Your custom application may need to support these translations.
- **Object identification** – primary and replicate tables may not be named the same, or have exactly the same schema or column names. The comparison application may need to accept very explicit instructions for location, database, and table and column names to be referenced.
- **Subset comparison** – the application may need to compare only a portion of a table. The ability to specify a where clause type of select for both primary and replicate tables may be important.
- **Latency** – in a replication system, there is always some latency (a measure of the time it takes a primary transaction to appear in a replicate table). A comparison application must include some tolerance to distinguish between rows that are “not there” and “not there yet.”
- **Data transformation** – the application must be able to handle difference in precision and format between different databases, the same way Replication Server supports class-level translations. An option to simplify the application may be to allow certain columns to be excluded from the comparison process, based on datatype (for example, do not compare the DATE datatypes of different database vendors).
- **Large object (LOB) data** – Large object (for example, LOB, CLOB, TEXT, or IMAGE) datatypes cause additional processing issues because of their

size. To improve performance, you can limit the number of bytes used for comparison, if the likelihood of a “non-match” can still be relied on.

For more information about the `rs_subcmp` utility, see the *Replication Server Administration Guide* and the *Replication Server Reference Manual*.

Replication with Adaptive Server Anywhere

This appendix describes the issues involved with using Adaptive Server Anywhere (ASA) in a Replication Server system.

Topic	Page
ASA primary data servers	213
ASA replicate data servers	216

ASA primary data servers

This section describes the primary database issues and considerations specific to the Adaptive Server Anywhere (ASA) data server in a Sybase replication system.

ASA is a different relational database than Sybase Adaptive Server Enterprise, which is the original primary data server supported by Replication Server.

Adaptive Server Enterprise is designed specifically for high-performance OLTP (online transaction processing) and mixed workload enterprise computing. By contrast, ASA is designed for the following applications:

- Embedded database

Many applications (personal information managers, document management systems—nearly any application that stores information) require a database “behind the scenes.” ASA is intended to be the database for these applications. The UltraLite deployment option is intended for embedded environments that have limited resources.

- Mobile computing

With its SQL Remote replication, ASA extends transaction-based computing throughout the enterprise. The UltraLite deployment

option and MobiLink synchronization technology provide full database functionality on devices with limited resources.

- Workgroup server

Workgroups ranging in size from a few people to a few hundred people in an organization may require a data server that they can share. ASA is a multiuser server that can provide a high-performance database for workgroups, well-suited for (but not limited to) environments where administration and hardware resources are limited.

ASA Replication Agent

As a primary data server in a replication system, ASA interacts with the Adaptive Server Anywhere Replication Agent. The ASA Replication Agent is responsible for identifying and transferring transactions from the ASA primary database to a primary Replication Server.

ASA primary database permissions

A script file is provided with ASA to set up an ASA database to function as a primary database. The *rssetup.sql* script file performs the following operations to set up ASA as a primary database for replication:

- Creates a user ID named dbmaint, with password dbmaint, with DBA permissions. This is the Replication Server Maintenance User ID and password required to connect to the primary database.
- Creates a user named sa, with password sysadmin, with DBA permissions. This is the user ID used by Replication Server when materializing data.
- Adds sa and dbmaint to a group named rs_systabgroup.

Replication intrusions and impacts in ASA

Setting REPLICATE ON for an ASA table places additional information in the database transaction log whenever an insert, update, or delete operation occurs on the table. The additional information significantly increases the log resources used by the server. The ASA Replication Agent uses this information to submit the full *before image* of the row, where required, to Replication Server for replication.

Primary database connectivity for ASA

By default, the ASA Replication Agent uses the sa user ID and password created by the *rssetup.sql* command file to log in to the ASA primary database.

You must create the user ID and password that the ASA Replication Agent uses to log in to the primary Replication Server. Record these values in the rs_user and rs_pw configuration parameters in the ASA Replication Agent configuration file. This user ID must have been defined and granted connect source permission in the primary Replication Server.

A Replication Server database connection name is made up of two parts: a data server name (server_name) and a database name (db_name). Record these

	<p>values in the RS_source_ds and RS_source_db configuration parameters in the ASA Replication Agent configuration file.</p> <p>The primary Replication Server does not use the source_db portion of the Replication Server database connection. Instead, Replication Server obtains the database name from the command line of the data server identified in the source_ds portion of the Replication Server database connection. You must, however, include a database name in the Replication Server create connection statement to conform to the syntax.</p>
Primary database limitations in ASA	<p>One ASA Replication Agent is required for each ASA primary database from which transactions are replicated.</p> <p>The ASA Replication Agent does not “pre-read” data from the RSSD of the primary Replication Server, as other Replication Agents do.</p> <p>You cannot substitute an ASA Replication Agent for an Adaptive Server Enterprise Replication Agent (RepAgent), since the ASA and Adaptive Server Enterprise transaction logs have different formats.</p>
ASA primary database configuration issues	<p>In ASA, all database object identifiers are case insensitive (that is, uppercase and lowercase are treated as the same). In Adaptive Server Enterprise, database object identifiers are case sensitive by default. With ASA, you must ensure that the case of database object identifiers match in all parts of the SQL statements to ensure compatibility with Adaptive Server Enterprise.</p>
Replication definitions for primary tables in ASA	<p>Because the ASA Replication Agent does not “pre-read” data from the RSSD of the primary Replication Server, the entire row for all replicated operations is sent to the primary Replication Server. By contrast, other Replication Agents can selectively send only those columns identified by the replication definition.</p>
ASA primary datatype translation issues	<p>ASA supports data of zero length that is not NULL. However, non-null long, varchar, and long binary data of zero length are replicated to a replicate site as NULL.</p> <p>If a primary table has columns with unsupported datatypes, you can replicate the data if you create a replication definition using a compatible supported datatype. For example, to replicate a double column, you could define the column as float in the replication definition.</p>
ASA system management issues	<p>Replication Server Manager support for ASA is not implemented in version 12.6.</p>
Other primary database issues for ASA	<p>One of the differences between the Adaptive Server Enterprise Replication Agent (RepAgent thread) and the ASA Replication Agent is that while the Adaptive Server Enterprise RepAgent thread depends on a temporary recovery database for access to old transactions, the ASA Replication Agent depends on</p>

access to old transaction logs. No temporary recovery database exists for the ASA Replication Agent.

ASA replicate data servers

This section describes the replicate database issues and considerations specific to the Adaptive Server Anywhere (ASA) data server in a Sybase replication system.

ASA is a different relational database from Sybase Adaptive Server Enterprise, which is the original replicate data server supported by Replication Server.

As a replicate data server in a replication system, ASA interacts directly with the replicate Replication Server. The replicate Replication Server logs in to the ASA replicate database and applies replicated transactions.

ASA replicate database permissions

A script file is provided with ASA to set up an ASA database to function as a replicate database. The *rssetup.sql* script file performs the following operations to set up ASA as a replicate database for replication:

- Creates a user ID named dbmaint, with password dbmaint, with DBA permissions. This is the Replication Server Maintenance User ID and password required to connect to the primary database.
- Creates a user named sa, with password sysadmin, with DBA permissions. This is the user ID used by Replication Server when materializing data.
- Adds sa and dbmaint to a group named rs_systabgroup.
- Creates the system tables and stored procedures necessary in the replicate database to support Replication Server.

Replication intrusions and impacts in ASA

The only significant intrusions or impacts to the ASA replicate database are the database objects created by the *rssetup.sql* script to support Replication Server replicate database operations.

The *rssetup.sql* script creates two tables in the replicate database to support Replication Server operations:

- *rs_info* – contains information about the sort order and character set used by the replicate database. You should confirm that the insert statements for this table set the proper character set and sort order for your replicate database.

Replicate database connectivity for ASA	<ul style="list-style-type: none"> rs_lastcommit – holds information about the last successfully committed replicate transaction applied to the replicate database. <p>You do not need to use a database gateway when you use ASA as a replicate data server. The replicate Replication Server connects directly to the ASA replicate data server.</p> <p>A Replication Server database connection name is made up of two parts: a data server name (server_name) and a database name (db_name). The replicate Replication Server looks for an interfaces file entry for the replicate database server_name specified in the database connection.</p> <p>Replication Server logs in to the replicate data server using the user_name and password specified in the database connection. For ASA replicate databases, the user_name and password should be the dbmaint user ID and password that were created when the <i>rssetup.sql</i> script was executed.</p> <p>The replicate Replication Server does not use the source_db portion of the Replication Server database connection. Instead, Replication Server obtains the database name from the command line of the data server identified in the source_ds portion of the Replication Server database connection. You must, however, include a database name in the Replication Server create connection statement to conform to the syntax.</p> <p>You must make an entry in the Replication Server interfaces file to identify the host and port where the ASA replicate data server is listening. The interfaces file entry name must match the server_name portion of the Replication Server database connection.</p>
Replicate database limitations in ASA	<p>Support for ASA as a data source is not provided in Replication Server Manager (RSM) version 12.6. RSM does support ASA as a generic Open Server component, which means that RSM cannot manage or configure ASA data servers, but it can test them to see if they are active.</p>
ASA replicate datatype translation issues	<p>The Heterogeneous Datatype Support (HDS) feature of Replication Server does not support ASA databases. No user-defined datatypes, translations, or function strings are provided.</p>

Glossary

active database	In a warm standby application, a database that is replicated to a standby database. See also warm standby application .
Adaptive Server	The Sybase version 11.5 and later relational database server.
applied function	A replicated function, associated with a function replication definition, that Replication Server delivers from a primary database to a subscribing replicate database. The function passes parameter values to a stored procedure that is executed at the replicate database. See also replicated function delivery , request function , and function replication definition .
article	A replication definition extension for tables or stored procedures that can be an element of a publication. Articles may or may not contain where clauses, which specify a subset of rows that the replicate database receives.
asynchronous procedure delivery	A method of replicating, from a source to a destination database, a stored procedure that is associated with a table replication definition.
asynchronous command	A command that a client submits where the client is not prevented from proceeding with other operations before the completion status is received. Many Replication Server commands function as asynchronous commands within the replication system.
atomic materialization	A materialization method that copies subscription data from a primary to a replicate database through the network in a single atomic operation, using a select operation with a holdlock. No changes to primary data are allowed until data transfer is complete. Replicate data may be applied either as a single transaction or in increments of ten rows per transaction, which ensures that the replicate database transaction log does not fill. Atomic materialization is the default method for the create subscription command. See also nonatomic materialization , bulk materialization and no materialization .
autocorrection	A setting applied to replication definitions, using the set autocorrection command, to prevent failures caused by missing or duplicate rows in a copy of a replicated table. When autocorrection is enabled, Replication Server converts each update or insert operation into a delete followed by

	an insert. Autocorrection should <i>only</i> be enabled for replication definitions whose subscriptions use nonatomic materialization.
base class	A function-string class that does not inherit function strings from a parent class. See also function-string class .
bitmap subscription	A type of subscription that replicates rows based on bitmap comparisons. Create columns using the int datatype, and identify them as the rs_address datatype when you create a replication definition. When you create a subscription, compare each rs_address column to a bitmask using a bitmap comparison operator (&) in the where clause. Rows matching the subscription's bitmap are replicated.
bulk materialization	A materialization method whereby subscription data in a replicate database is initialized outside of the replication system. For example, data may be transferred from a primary database using media such as magnetic tape, diskette, CD-ROM, or optical storage disk. Bulk materialization involves a series of commands, starting with define subscription. You can use bulk materialization for subscriptions to table replication definitions or function replication definitions. See also atomic materialization , nonatomic materialization , and no materialization .
centralized database system	A database system where data is managed by a single database management system at a centralized location.
class	See error class and function-string class .
class tree	A set of function-string classes, consisting of two or more levels of derived and parent classes, that derive from the same base class. See also function-string class .
client	A program connected to a server in a client/server architecture. It may be a front-end application program executed by a user or a utility program that executes as an extension of the system.
Client/Server Interfaces (C/SI)	The Sybase interface standard for programs executing in a client/server architecture.
concurrency	The ability of multiple clients to share data or resources. Concurrency in a database management system depends upon the system protecting clients from conflicts that arise when data in use by one client is modified by another client.
connection	A connection from a Replication Server to a database. See also Data Server Interface (DSI) and logical connection .

coordinated dump	A set of database dumps or transaction dumps that is synchronized across multiple sites by distributing an <code>rs_dumpdb</code> or <code>rs_dumptran</code> function through the replication system.
database	A set of related data tables and other objects that is organized and presented to serve a specific purpose.
database generation number	Stored in both the database and the RSSD of the Replication Server that manages the database, the database generation number is the first part of the origin queue ID (QID) of each log record. The origin queue ID ensures that the Replication Server does not process duplicate records. During recovery operations, you may need to increment the database generation number so that Replication Server does not ignore records submitted after the database is reloaded.
database replication definition	<p>A description of a set of database objects—tables, transactions, functions, system stored procedures, and DDL—for which a subscription can be created.</p> <p>You can also create table replication definitions and function replication definitions. See also table replication definition and function replication definition.</p>
database server	A server program, such as Sybase Adaptive Server, that provides database management services to clients.
data definition language (DDL)	The set of commands in a query language, such as Transact-SQL, that describes data and their relationships in a database. DDL commands in Transact-SQL include those using the create, drop, and alter keywords.
data manipulation language (DML)	The set of commands in a query language, such as Transact-SQL, that operates on data. DML commands in Transact-SQL include select, insert, update, and delete.
data server	A server whose client interface conforms to the Sybase Client/Server Interfaces and provides the functionality necessary to maintain the physical representation of a replicated table in a database. Data servers are usually database servers, but they can also be any data repository with the interface and functionality Replication Server requires.
Data Server Interface (DSI)	Replication Server threads corresponding to a connection between a Replication Server and a database. DSI threads submit transactions from the DSI outbound queue to a replicate data server. They consist of a scheduler thread and one or more executor threads. The scheduler thread groups the transactions by commit order and dispatches them to the executor threads. The executor threads map functions to function strings and execute the transactions.

	in the replicate database. DSI threads use an Open Client connection to a database. See also outbound queue and connection .
data source	A specific combination of a database management system (DBMS) product such as a relational or non-relational data server, a database residing in that DBMS, and the communications method used to access that DBMS from other parts of a replication system. See also database and data server .
decision support application	A database client application characterized by ad hoc queries, reports, and calculations and few data update transactions.
declared datatype	<p>The datatype of the value delivered to the Replication Server from the Replication Agent:</p> <ul style="list-style-type: none">• If the Replication Agent delivers a base Replication Server datatype, such as datetime, to the Replication Server, the declared datatype is the base datatype.• Otherwise, the declared datatype must be the UDD for the original datatype at the primary database.
default function string	The function string that is provided by default for the system-provided classes rs_sqlserver_function_class and rs_default_function_class and classes that inherit function strings from these classes, either directly or indirectly. See also function string .
dematerialization	The optional process, when a subscription is dropped, whereby specific rows that are not used by other subscriptions are removed from the replicate database.
derived class	A function-string class that inherits function strings from a parent class. See also function-string class and parent class .
direct route	A route used to send messages directly from a source to a destination Replication Server, with no intermediate Replication Servers. See also indirect route and route .
disk partition	See partition .
distributed database system	A database system where data is stored in multiple databases on a network. The databases may be managed by data servers of the same type (for example, Adaptive Server) or by heterogeneous data servers.
Distributor	A Replication Server thread (DIST) that helps to determine the destination of each transaction in the inbound queue.

dump marker	A message written by Adaptive Server in a database transaction log when a dump is performed. In a warm standby application, when you are initializing the standby database with data from the active database, you can specify that Replication Server use the dump marker to determine where in the transaction stream to begin applying transactions in the standby database. See also warm standby application .
Embedded Replication Server System Database (ERSSD)	The Adaptive Server Anywhere (ASA) database that stores Replication Server system tables. You can choose whether to store Replication Server system tables on the ERSSD or the Adaptive Server RSSD. See also Replication Server System Database (RSSD) .
error action	A Replication Server response to a data server error. Possible Replication Server error actions are ignore, warn, retry_log, log, retry_stop, and stop_replication. Error actions are assigned to specific data server errors.
error class	A name for a collection of data server error actions that are used with a specified database.
exceptions log	A set of three Replication Server system tables that holds information about transactions that failed on a data server. The transactions in the log must be resolved by a user or by an intelligent application. You can use the rs_helpexception stored procedure to query the exceptions log.
Failover	<p>Sybase Failover allows you to configure two version 12.0 and later Adaptive Servers as companions. If the primary companion fails, that server's devices, databases, and connections can be taken over by the secondary companion.</p> <p>For more detailed information about how Sybase Failover works in Adaptive Server, refer to <i>Using Sybase Failover in a High Availability System</i>, which is part of the Adaptive Server Enterprise documentation set.</p> <p>For instructions on how to enable Failover support for non-RSSD Replication Server connections to Adaptive Server, see "Configuring the Replication System to Support Sybase Failover" in Chapter 16, "Replication System Recovery," of the <i>Replication Server Administration Guide</i>.</p>
fault tolerance	The ability of a system to continue to operate correctly even though one or more of its component parts is malfunctioning.
function	A Replication Server object that represents a data server operation such as insert, delete, select, or begin transaction. Replication Server distributes such operations to other Replication Servers as functions. Each function consists of a function name and a set of data parameters. In order to execute the function in a destination database, Replication Server uses function strings to convert a

	function to a command or set of commands for a type of database. See also user-defined function , and replicated function delivery .
function replication definition	A description of a replicated function used in replicated function delivery. The function replication definition, maintained by Replication Server, includes information about the parameters to be replicated and the location of the primary version of the affected data. See also replicated function delivery .
function scope	The range of a function's effect. Functions have replication definition scope or function-string class scope. A function with replication definition scope is defined for a specific replication definition, and cannot be applied to other replication definitions. A function with function-string class scope is defined once for a function-string class and is available only within that class.
function string	A string that Replication Server uses to map a database command to a data server API. For the <code>rs_select</code> and <code>rs_select_with_lock</code> functions only, the string contains an input template, used to match function strings with the database command. For all functions, the string also contains an output template, used to format the database command for the destination data server.
function-string class	A named collection of function strings used with a specified database connection. Function-string classes include those provided with Replication Server and those you have created. Function-string classes can share function string definitions through function-string inheritance. The three system-provided function-string classes are <code>rs_sqlserver_function_class</code> , <code>rs_default_function_class</code> , and <code>rs_db2_function_class</code> . See also base class , class tree , derived class , function-string inheritance , and parent class .
function-string inheritance	The ability to share function string definitions between classes, whereby a derived class inherits function strings from a parent class. See also derived class , function-string class , and parent class .
function-string variable	An identifier used in a function string to represent a value that is to be substituted at run time. Variables in function strings are enclosed in question marks (?). They represent column values, function parameters, system-defined variables, or user-defined variables.
function subscription	A subscription to a function replication definition (used in applied function delivery).
generation number	See database generation number .
heterogeneous data servers	Data servers that are supplied by more than one vendor, used together in a distributed database system.

high availability (HA)	Very low downtime. Computer systems that provide HA usually provide 99.999% availability, or roughly five minutes unscheduled downtime per year.
hibernation mode	A Replication Server state in which all DDL commands, except admin and sysadmin commands, are rejected; all routes and connections are suspended; most service threads, such as DSI and RSI, are suspended; and RSI and RepAgent (or LTM) users are logged off and not allowed to log on. Used during route upgrades, and may be turned on for a Replication Server to debug problems.
hot standby application	A database application in which the standby database can be placed into service without interrupting client applications and without losing any transactions. See also warm standby application .
ID Server	One Replication Server in a replication system is the ID Server. In addition to performing the usual Replication Server tasks, the ID Server assigns unique ID numbers to every Replication Server and database in the replication system, and maintains version information for the replication system.
inbound queue	A stable queue used to spool messages from a Replication Agent to a Replication Server.
indirect route	A route used to send messages from a source to a destination Replication Server, through one or more intermediate Replication Servers. See also direct route and route .
interfaces file	A file containing entries that define network access information for server programs in a Sybase client/server architecture. Server programs may include Adaptive Servers, SQL Servers, gateways, Replication Servers, and Replication Agents such as LTM for SQL Server. The interfaces file entries enable clients and servers to connect to each other in a network.
latency	The measure of the time it takes to distribute to a replicate database a data modification operation first applied in a primary database. The time includes Replication Agent processing, Replication Server processing, and network overhead.
local-area network (LAN)	A system of computers and devices, such as printers and terminals, connected by cabling for the purpose of sharing data and devices.
locator value	The value stored in the rs_locator table of the Replication Server's RSSD that identifies the latest log transaction record received and acknowledged by the Replication Server from each previous site during replication.

logical connection	A database connection that Replication Server maps to the connections for the active and standby databases in a warm standby application. See also connection and warm standby application .
login name	The name that a user or a system component such as Replication Server uses to log in to a data server, Replication Server, or Replication Agent.
Log Transfer Language (LTL)	A subset of the Replication Command Language (RCL). A Replication Agent such as RepAgent or LTM for SQL Server uses LTL commands to submit to Replication Server the information it retrieves from primary database transaction logs.
Log Transfer Manager (LTM)	The Replication Agent program for Sybase SQL Server. See also Replication Agent and RepAgent thread .
Maintenance User	A data server login name that Replication Server uses to maintain replicate data. In most applications, Maintenance User transactions are not replicated.
materialization	The process of copying data specified by a subscription from a primary database to a replicate database, thereby initializing the replicate table. Replicate data can be transferred over a network, or, for subscriptions involving large amounts of data, loaded initially from media. See also atomic materialization , bulk materialization , no materialization , and nonatomic materialization .
materialization queue	A stable queue used to spool messages related to a subscription being materialized or dematerialized.
missing row	A row missing from a replicated copy of a table but present in the primary table.
mixed-version system	<p>A replication system containing Replication Servers of different software versions that have different capabilities based on their different software versions and site versions. Mixed-version support is available only if the system version is 11.0.2 or greater.</p> <p>For example, a replication system containing Replication Servers version 11.5 or later and version 11.0.2 is a mixed-version system. A replication system containing Replication Servers of releases earlier than release 11.0.2 is not a mixed-version system, because any newer Replication Servers are restricted by the system version from using certain new features. See also site version and system version.</p>
more columns	Columns in a replication definition exceeding 250, but limited to 1024. More columns are supported by Replication Server version 12.5 and later.

multi-site availability (MSA)	Methodology for replicating database objects—tables, functions, transactions, system stored procedures, and DDL from the primary to the replicate database. See also database replication definition .
name space	The scope within which an object name must be unique.
nonatomic materialization	A materialization method that copies subscription data from a primary to a replicate database through the network in a single operation, without a holdlock. Changes to the primary table are allowed during data transfer, which may cause temporary inconsistencies between replicate and primary databases. Data is applied in increments of ten rows per transaction, which ensures that the replicate database transaction log does not fill. Nonatomic materialization is an optional method for the create subscription command. See also autocorrection , atomic materialization , no materialization , and bulk materialization .
network-based security	Secure transmission of data across a network. Replication Server supports third-party security mechanisms that provide user authentication, unified login, and secure message transmission between Replication Servers.
no materialization	A materialization method that lets you create a subscription when the subscription data already exists at the replicate site. Use the create subscription command with the without materialization clause. You can use this method to create subscriptions to table replication definitions and function replication definitions. See also atomic materialization and bulk materialization .
online transaction processing (OLTP) application	A database client application characterized by frequent transactions involving data modification (inserts, deletes, and updates).
origin queue ID (QID)	Formed by the Replication Agent, the qid uniquely identifies each log record passed to the Replication Server. It includes the date and time from the primary data server, and the database generation number. See also database generation number .
orphaned row	A row in a replicated copy of a table that does not match an active subscription.
outbound queue	A stable queue used to spool messages. The DSI outbound queue spools messages to a replicate database. The RSI outbound queue spools messages to a replicate Replication Server.
parallel DSI	A method of configuring a database connection so that transactions are applied to a replicate data server using multiple DSI threads operating in parallel, rather than a single DSI thread. See also connection and Data Server Interface (DSI) .

parameter	An identifier representing a value that is provided when a procedure executes. Parameter names are prefixed with an @ character in function strings. When a procedure is called from a function string, Replication Server passes the parameter values, unaltered, to the data server. See also searchable parameter .
parent class	A function-string class from which a derived class inherits function strings. See also function-string class and derived class .
partition	A raw disk partition or operating system file that Replication Server uses for stable queue storage. Only use operating system files in a test environment.
physical connection	See connection .
primary data	The definitive version of a set of data in a replication system. The primary data is maintained on a data server that is known to all of the Replication Servers with subscriptions for the data.
primary database	Any database that contains data that is replicated to another database through the replication system.
primary fragment	A horizontal segment of a table that holds the primary version of a set of rows.
primary key	A set of table columns that uniquely identifies each row.
primary site	A Replication Server where a function-string class or error class is defined. See error class and function-string class .
principal user	The user who starts an application. When using network-based security, Replication Server logs in to remote servers as the principal user.
projection	A vertical slice of a table, representing a subset of the table's columns.
publication	A group of articles from the same primary database. A publication lets you collect replication definitions for related tables and/or stored procedures and then subscribe to them as a group. You collect replication definitions as articles in a publication at the source Replication Server and subscribe to them with a publication subscription at the destination Replication Server. See also article and publication subscription .
publication subscription	A subscription to a publication. See also article and publication .
published datatype	The datatype of the column after the column-level translation (and before a class-level translation, if any) at the replicate data server. The published datatype must be either a Replication Server base datatype or a UDD for the datatype in the target data server. If the published datatype is omitted from the replication definition, it defaults to the declared datatype

query	In a database management system, a request to retrieve data that meets a given set of criteria. The SQL database language includes the <code>select</code> command for queries.
quiescent	A replication system in which all data-changing operations (or transactions) have been propagated to their destinations. Some Replication Server commands or procedures require that you first quiesce the replication system.
remote procedure call (RPC)	A request to execute a procedure that resides in a remote server. The server that executes the procedure could be a Adaptive Server, a Replication Server, or a server created using Open Server. The request can originate from any of these servers or from a client application. The RPC request format is a part of the Sybase Client/Server Interfaces.
RepAgent thread	The Replication Agent for Adaptive Server Enterprise. The RepAgent thread is an Adaptive Server thread. It sends transaction log information from the primary database to the primary Replication Server.
Rep Agent User thread	The thread on a Replication Server database connection that a Replication Agent connects with, on behalf of a primary database. See also Data Server Interface (DSI) .
replicate database	Any database that contains data that is replicated from another database through the replication system.
replicated function delivery	A method of replicating, from a source to a destination database, a stored procedure that is associated with a function replication definition. See also applied function , request function , and function replication definition .
replicated stored procedure	An Adaptive Server stored procedure that is marked as replicated using the <code>sp_setrepproc</code> or the <code>sp_setreplicate</code> system procedure. Replicated stored procedures can be associated with function replication definitions or table replication definitions. See also replicated function delivery and asynchronous procedure delivery .
replicated table	A table that is maintained by Replication Server, in part or in whole, in databases at multiple locations. There is one primary version of the table, which is marked as replicated using the <code>sp_setreptable</code> or the <code>sp_setreplicate</code> system procedure; all other versions are replicated copies.
Replication Agent	A program or module that sends transaction log information from a primary data server to a primary Replication Server. The Replication Agent for Adaptive Server Enterprise is the RepAgent thread. Sybase provides separate Replication Agent software products to support non-Sybase data servers in a replication system.

Replication Command Language (RCL)	The commands used to manage information in Replication Server.
replication definition	<p>Usually, a description of a table for which subscriptions can be created. The replication definition, maintained by Replication Server, includes information about the columns to be replicated and the location of the primary version of the table.</p> <p>You can also create function replication definitions; sometimes the term “table replication definition” is used to distinguish between table and function replication definitions. See also function replication definition.</p>
Replication Server	The Sybase server program that maintains replicated data, typically on a LAN, and processes data transactions received from other Replication Servers on the same LAN or on a WAN.
Replication Server Interface (RSI)	A thread that logs in to a destination Replication Server and transfers commands from the RSI outbound stable queue to the destination Replication Server. There is one RSI thread for each destination Replication Server that is a recipient of commands from a primary or intermediate Replication Server. See also outbound queue and route .
Replication Server Manager (RSM)	An application for managing and monitoring a replication system and its components. Includes status and monitoring functions, diagnostics and troubleshooting functions, asynchronous notification of user-defined events, and operational control over many aspects of the system.
Replication System Administrator	The System Administrator who manages routine operations in the Replication Server.
Replication Server System Database (RSSD)	The Adaptive Server database containing a Replication Server system tables. You can choose whether to store Replication Server system tables on the RSSD or the Adaptive Server Anywhere (ASA) ERSSD. See also Embedded Replication Server System Database (ERSSD) .
Replication Server system Adaptive Server	The Adaptive Server with the database containing a Replication Server’s system tables (the RSSD).
replication system	A data processing system where data is replicated in multiple databases to provide remote users with the benefits of local data access. Specifically, a replication system that is based upon Replication Server and includes other components such as Replication Agents and data servers.
replication system domain	All replication system components that use the same ID Server.

request function	A replicated function, associated with a function replication definition, that Replication Server delivers from a replicate database to a primary database. The function passes parameter values to a stored procedure that is executed at the primary database. See also replicated function delivery , request function , and function replication definition .
route	A 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. See also direct route and indirect route .
route version	The lower of the site version numbers of the route's source and destination Replication Servers. Replication Server version 11.5 and later use the route version number to determine which data to send to the replicate site. See also site version .
row migration	The process whereby column value changes in rows in a primary version of a table cause corresponding rows in a replicate version of the table to be inserted or deleted, based on comparison with values in a subscription's where clause.
SQL Server	The Sybase relational database pre-11.5 server.
schema	The structure of the database. DDL commands and system procedures change system tables stored in the database. Supported DDL commands and system procedures can be replicated to standby databases when you use Replication Server version 11.5 or later and Adaptive Server version 11.5 or later.
searchable column	A column in a replicated table that can be specified in the where clause of a subscription or article to restrict the rows replicated at a site.
searchable parameter	A parameter in a replicated stored procedure that can be specified in the where clause of a subscription to help determine whether or not the stored procedure should be replicated. See also parameter .
secondary truncation point	See truncation point .
site	An installation consisting of, at minimum, a Replication Server, data server, and database, and possibly a Replication Agent, usually at a discrete geographic location. The components at each site are connected over a WAN to those at other sites in a replication system. See also primary site .
site version	The version number for an individual Replication Server. Once the site version has been set to a particular level, the Replication Server enables features specific to that level, and downgrades are not allowed. See also software version , route version , and system version .

software version	The version number of the software release for an individual Replication Server. See also site version and system version .
Stable Queue Manager (SQM)	A thread that manages the stable queues. There is one Stable Queue Manager (SQM) thread for each stable queue accessed by the Replication Server, whether inbound or outbound.
Stable Queue Transaction (SQT) interface	A thread that reassembles transaction commands in commit order. A Stable Queue Transaction (SQT) interface thread reads from inbound stable queues, puts transactions in commit order, then sends them to the Distributor (DIST) thread or a DSI thread, depending on which thread required the SQT ordering of the transaction.
stable queues	Store-and-forward queues where Replication Server stores messages destined for a route or database connection. Messages written into a stable queue remain there until they can be delivered to the destination Replication Server or database. Replication Server builds stable queues using its disk partitions. See also inbound queue , outbound queue , and materialization queue .
standalone mode	A special Replication Server mode used for initiating recovery operations.
standby database	In a warm standby application, a database that receives data modifications from the active database and serves as a backup of that database. See also warm standby application .
stored procedure	A collection of SQL statements and optional control-of-flow statements stored under a name in a Adaptive Server database. Stored procedures supplied with Adaptive Server are called system procedures. Some stored procedures for querying the RSSD are included with the Replication Server software.
subscription	A request for Replication Server to maintain a replicated copy of a table, or a set of rows from a table, in a replicate database at a specified location. You can also subscribe to a function replication definition, for replicating stored procedures.
subscription dematerialization	See dematerialization .
subscription materialization	See materialization .
subscription migration	See row migration .
Sybase Central	A graphical tool that provides a common interface for managing Sybase and Powersoft products. Replication Server uses Replication Server Manager as a Sybase Central plug-in. See also Replication Server Manager (RSM) .

symmetric multiprocessing (SMP)	On a multiprocessor platform, the ability of an application's threads to run in parallel. Replication Server supports SMP, which can improve server performance and efficiency.
synchronous command	A command that a client considers complete only after the completion status is received.
system function	A function that is predefined and part of the Replication Server product. Different system functions coordinate replication activities, such as <code>rs_begin</code> , or perform data manipulation operations, such as <code>rs_insert</code> , <code>rs_delete</code> , and <code>rs_update</code> .
system-provided classes	The error class <code>rs_sqlserver_error_class</code> and the function-string classes <code>rs_sqlserver_function_class</code> , <code>rs_default_function_class</code> , and <code>rs_db2_function_class</code> that Replication Server provides. Function strings are generated automatically for the system-provided function-string classes and for any derived classes that inherit from these classes, directly or indirectly. See also error class and function-string class .
system version	The version number for a replication system that represents the version for which new features are enabled, for Replication Servers of release 11.0.2 or earlier, and below which no Replication Server can be downgraded or installed. For a Replication Server version 11.5, your use of certain new features requires a site version of 1150 and a system version of at least 1102. See also mixed-version system , site version , and software version .
table replication definition	See replication definition .
table subscription	A subscription to a table replication definition.
thread	A process running within Replication Server. Built upon Sybase Open Server, Replication Server has a multi-threaded architecture. Each thread performs a certain function such as managing a user session, receiving messages from a Replication Agent or another Replication Server, or applying messages to a database. See also Data Server Interface (DSI) , Distributor , and Replication Server Interface (RSI) .
transaction	A mechanism for grouping statements so that they are treated as a unit: either all statements in the group are executed or no statements in the group are executed.
Transact-SQL	The relational database language used with Adaptive Server. It is based on standard SQL (Structured Query Language), with Sybase extensions.

truncation point	<p>An Adaptive Server database that holds primary data has an active truncation point, marking the transaction log location where Adaptive Server has completed processing. This is the primary truncation point.</p> <p>The RepAgent for an Adaptive Server database maintains a secondary truncation point, marking the transaction log location separating the portion of the log successfully submitted to the Replication Server from the portion not yet submitted. The secondary truncation point ensures that each operation enters the replication system before its portion of the log is truncated.</p>
user-defined function	<p>A function that allows you to create custom applications that use Replication Server to distribute replicated functions or asynchronous stored procedures between sites in a replication system. In replicated function delivery, a user-defined function is automatically created by Replication Server when you create a function replication definition.</p>
variable	<p>See function-string variable.</p>
version	<p>See mixed-version system, site version, software version, and system version.</p>
warm standby application	<p>An application that employs Replication Server to maintain a standby database for a database known as the active database. If the active database fails, Replication Server and client applications can switch to the standby database.</p>
wide-area network (WAN)	<p>A system of local-area networks (LANs) connected together with data communication lines.</p>
wide columns	<p>Columns in a replication definition containing char, varchar, binary, varbinary, unichar, univarchar, or Java inrow data that are wider than 255 bytes. Wide columns are supported by Replication Server version 12.5 and later.</p>
wide data	<p>Wide data rows, limited to the size of the data page on the data server. Adaptive Server supports page sizes of 2K, 4K, 8K, and 16K. Wide data is supported by Replication Server version 12.5 and later.</p>
wide messages	<p>Messages larger than 16K that span blocks. Wide messages are supported by Replication Server version 12.5 and later.</p>

Index

A

Adaptive Server Anywhere
 as primary database 213–216
 as replicate database 216–217
 configuration parameters 214
 dbmaint user ID 214, 216
 long binary datatype 215
 long datatype 215
 primary database configuration 215
 primary database permissions 214
 replicate database setup 216–217
 replicate on command 214
 Replication Agent for 213–216
 transaction log 215
 varchar datatype 215
 zero-length data 215
Adaptive Server Enterprise
 as primary database 148–149
 as replicate database 147–148
 binary datatype 75
 bit datatype 85, 120
 char datatype 41, 85, 120
 CIS feature 99–100, 106
 datetime datatype 40, 41, 98
 numeric datatype 90
 Replication Agent for 7
 varbinary datatype 75
 varchar datatype 41, 85, 120
administering replication systems 155–161
atomic bulk materialization 198, 202–205

B

bcp utility 204, 207
bidirectional replication 113–114
 database connections 135, 138
 Replication Agent filtering transactions 40, 46,
 54, 61, 68

with non-Sybase data servers 152–154

C

case sensitivity. *See* character case; identifiers
character case
 in replication definitions and subscriptions 125–
 127
 of configuration parameters 46, 53, 61, 68
 of object names in Adaptive Server Anywhere 215
 of object names in DB2 39, 46–47
 of object names in Informix 53, 54
 of object names in Microsoft SQL Server 61–62
 of object names in Oracle 69, 70
class-level translations
 See also heterogeneous datatype support (HDS)
 DB2 for OS/390 datatypes 77, 184–187
 DB2 for UNIX and Windows datatypes 81–82,
 184–187
 Informix datatypes 86, 187–190
 installing 124
 Microsoft SQL Server datatypes 92, 190–192
 Oracle datatypes 96, 192–195
CLASSPATH system variable 50, 65
columns
 maximum number for replication 53, 60, 68
 with large-object (LOB) datatype 25–27
commands
 create connection 40, 46, 53, 61, 68, 102, 108,
 110, 112, 124
 pdb_xlog 45, 53, 60, 68, 133
 replicate on, ASA 214
 resume connection 108, 111
 rs_dump 28
 rs_dumptran 29
 rs_init 27
 rs_marker 29
 rs_subcmp 30, 209–211
communications

- JDBC protocol 50, 58, 65
- RSM 132–133
- TCP/IP 37, 45, 52, 59, 67
- comparing databases 209–211
- Component Integration Services (CIS) 99–100, 106
 - See also* Adaptive Server Enterprise
- configuration parameters
 - LTM for MVS 37–41
 - Replication Agent for Informix 50–52
 - Replication Agent for Microsoft SQL Server 58–60
 - Replication Agent for Oracle 65–67
 - Replication Agent for UDB 43–45
- connect source** permission 102, 108
- conventions, document style xiii
- conventions, syntax xiv
- cs_convert** utility 115

D

- database gateways 8, 18–19, 106, 109–112
 - for DB2 for OS/390 replicate database 73–78
 - for DB2 for UNIX and Windows 78
 - for Informix replicate database 83–88
 - for Microsoft SQL Server replicate database 88–91, 93
 - for Oracle replicate database 94–95, 97–99
 - troubleshooting 169–170
- databases
 - loading data into replicate 201
 - materialization 28–30, 197–208
 - owner-qualified object names 24–25
 - primary database 3, 6, 33–71
 - reconciling 209–211
 - replicate database 3, 8, 72–99
 - Replication Server connection 105–115, 124, 134–138
 - setting up with RSM 143–145
 - unloading data from primary 199
- data-sharing environment, DB2 for OS/390 36, 39
- datatype translations 41, 85–87, 90, 98, 115–116
 - class-level 77, 86–87, 92, 96, 124, 200
 - column-level 199
 - during materialization 199–201
- datatypes
 - binary**, Sybase 75
 - bit**, Sybase 85, 120
 - BLOB**, DB2 75, 80
 - BLOB**, Informix 84
 - BLOB**, Oracle 95
 - boundary of 117–120
 - BYTE**, Informix 85
 - CHAR**, DB2 40, 41
 - char**, Sybase 41, 48, 55, 70, 85, 120
 - class-level translations for DB2 184–187
 - class-level translations for Informix 187–190
 - class-level translations for Microsoft SQL Server 190–192
 - class-level translations for Oracle 192–195
 - CLOB**, DB2 75, 80
 - CLOB**, Informix 84
 - CLOB**, Oracle 95
 - date and time** troubleshooting 171
 - DATE**, DB2 40, 41, 47–48
 - DATE**, Informix 55
 - DATE**, Oracle 70, 98
 - datetime**, Microsoft SQL Server 62
 - datetime**, Sybase 40, 41, 47–48, 55, 62, 70, 98, 171
 - DBCLOB**, DB2 75
 - DECIMAL**, Informix 85
 - decimal**, Microsoft SQL Server 119
 - default HDS translation 183–195
 - default translations for DB2 184–187
 - default translations for Informix 187–190
 - default translations for Microsoft SQL Server 190–192
 - default translations for Oracle 192–195
 - image**, Microsoft SQL Server 88
 - INT**, Informix 84, 120
 - large objects (LOB) 25–27, 75, 84, 88, 95, 121
 - long binary**, ASA 215
 - LONG RAW**, Oracle 95
 - long**, ASA 215
 - LONG**, Oracle 95
 - LVARCHAR**, DB2 80
 - LVARCHAR**, Informix 84
 - materialization 199–201
 - NCLOB**, Oracle 95
 - ntext**, Microsoft SQL Server 88
 - NUMERIC**, Informix 85
 - numeric**, Microsoft SQL Server 121
 - numeric**, Sybase 90

- rs_db2_char_for_bit**, HDS 75
 - rs_db2_varchar_for_bit**, HDS 75
 - rs_mss_numeric**, HDS 121
 - rs_msds_numeric**, HDS 90
 - rs_oracle_date**, HDS 98
 - rs_oracle_datetime**, HDS 98
 - SMALLINT**, Informix 84, 120
 - text**, Microsoft SQL Server 88
 - TIME**, DB2 47–48
 - TIMESTAMP**, DB2 47–48
 - translated by Replication Agent 41, 47–48, 55, 62, 70, 171
 - troubleshooting translations 171, 172–174
 - varbinary**, Sybase 75
 - varchar**, ASA 215
 - varchar**, Sybase 41, 48, 55, 70, 85, 120
- DB2 for OS/390
- as primary database 35–41
 - as replicate database 73–78
 - BLOB** datatype 75
 - CHAR** datatype 40, 41
 - class-level translation scripts 77
 - class-level translations 184–187
 - CLOB** datatype 75
 - data-sharing environment 36, 39
 - DATE** datatype 40, 41
 - DBCLOB** datatype 75
 - default datatype translations 184–187
 - function-string class for 124
 - gatewayless connection 8, 73, 75
 - LTM for MVS 37–41
 - LTMADMIN** user 35
 - LTMLASTCOMMIT** table 36, 74
 - LTMOBJECTS** table 36
 - primary database configuration 39–40
 - primary database connectivity 37–39
 - primary database permissions 35–36
 - replicate database setup 75–77
 - Replication Agent 17, 35–41
 - Replication Server connections 136–138
 - RS_INFO** table 73–74
 - RS_LASTCOMMIT** table 74
 - Sybase Log Extract 39
 - SYGWHOST** macro 38
 - transaction log 35, 36–37
 - translating primary datatypes 41
- DB2 for UNIX and Windows
- as primary database 41–49
 - as replicate database 78–82
 - BLOB** datatype 80
 - class-level translation scripts 81–82
 - class-level translations 184–187
 - CLOB** datatype 80
 - default datatype translations 184–187
 - JDBC driver 46
 - LVARCHAR** datatype 80
 - primary database configuration 46–47
 - primary database connectivity 42–45
 - primary database limitations 45
 - primary database permissions 42
 - replicate database configuration 80–82
 - replicate database connectivity 80
 - replicate database gateway 78
 - replicate database permissions 79
 - Replication Agent for 17, 41–49
 - RS_INFO** table 79
 - RS_LASTCOMMIT** table 79
 - RSM limitations 45, 49
 - transaction log 42
 - translating primary datatypes 47–48
- DB2 Universal Database. *See* DB2 for OS/390; DB2 for UNIX and Windows
- DirectConnect database gateways 18–19, 106, 109–112
- DB2 metadata 48
 - DirectConnect for DRDA 78
 - DirectConnect for Informix 83–88
 - DirectConnect for Microsoft SQL Server 88–91, 93
 - DirectConnect for Oracle 94–95, 97–99
 - DirectConnect for OS/390 73–78
 - Informix metadata 55
 - interfaces* file 111, 142
 - Microsoft SQL Server metadata 62–63
 - Oracle metadata 70
 - troubleshooting 169–170
- drivers, JDBC 50, 58, 65
- required for Oracle 68
- DSI thread 106, 107, 108, 110–111
- command batching 116
- dynamic performance views. *See* Oracle system views

E

error messages

- binary datatype 120
- datatype boundary 117–118, 119–120
- numeric identity 121

F

function strings

- for **rs_dump** command 28
- for **rs_dumptran** command 29
- for **rs_marker** command 29
- installing for HDS 123
- loading HDS 122
- modifying for LOB replication 25–27

function-string classes

- for DB2 for OS/390 124
- HDS feature 122
- modifying for LOB replication 25–27

G

gateway

See database gateways

gatewayless connection environment 8, 73, 75

H

heartbeat, RSM 157–161

heterogeneous datatype support (HDS)

- See also* Replication Server, HDS feature
- compared with CIS 99–100
- DB2 for OS/390 75–77
- DB2 for UNIX and Windows 80–82
- default datatype translation 183–195
- function-string classes 122, 139–140
- hds_db2_connection_sample.sql* script 77
- hds_db2_funcstrings.sql* script 76
- hds_db2_setup_for_replicate.sql* script 73, 75
- hds_db2_udds.sql* script 76
- hds_informix_connection_sample.sql* script 87
- hds_informix_funcstrings.sql* script 86
- hds_informix_setup_for_replicate.sql* script 83, 85

- hds_informix_udds.sql* script 86
- hds_msss_connection_sample.sql* script 92
- hds_msss_funcstrings.sql* script 91
- hds_msss_setup_for_replicate.sql* script 88, 91
- hds_msss_udds.sql* script 91
- hds_oracle_connection_sample.sql* script 97
- hds_oracle_funcstrings.sql* script 96
- hds_oracle_setup_for_replicate.sql* script 94, 95–96
- hds_oracle_udds.sql* script 96
- hds_udb_connection_sample.sql* script 82
- hds_udb_funcstrings.sql* script 80
- hds_udb_setup_for_replicate.sql* script 80
- hds_udb_udds.sql* script 80
- Informix database 85–87
- installing function strings 123
- limitations of 116–122
- Microsoft SQL Server database 91–93
- Oracle database 95–97
- scripts to emulate **rs_init** 123–124
- using 115–124

I

identifiers

- case sensitivity 39
- character case of 125–127
- object names in Adaptive Server Anywhere 215
- object names in DB2 46–47
- object names in Informix 53, 54
- object names in Microsoft SQL Server 61–62
- object names in Oracle 69, 70

@@IDENTITY global variable 57, 60

inbound queue, troubleshooting 164–166

Informix data server

- as primary database 49–56
- as replicate database 83–88
- BLOB** datatype 84
- BYTE** datatype 85
- class-level translation scripts 86
- class-level translations 187–190
- CLOB** datatype 84
- DATE** datatype 55
- DECIMAL** datatype 85
- default datatype translations 187–190

Informix schema 52
INT datatype 84, 120
 JDBC driver 50, 53
 log scan position 176
LVARCHAR datatype 84
 maximum number of columns 53
NUMERIC datatype 85
 numeric precision 120
 primary database configuration 53–54
 primary database limitations 52–53
 primary database permissions 49
 replicate database setup 85–87
 Replication Agent 17, 49–56
 Replication Agent configuration parameters 50–52
rs_info table 83
rs_lastcommit table 83
 RSM limitations 53, 56
 selective update triggers 52
SMALLINT datatype 84, 120
 transaction isolation in 52
 transaction log 50
 translating primary datatypes 55
 triggers 52
 interfaces file 108–109, 217
 interfaces file 74, 84, 89, 95, 111, 132, 142

J

Java Runtime Environment (JRE) 46, 53, 64, 68
 JDBC communications protocol 50, 58, 65
 drivers 68

L

large object (LOB) datatypes
 in DB2 for OS/390 database 26, 75
 in DB2 for UNIX and Windows 80
 in Informix database 84
 in Microsoft SQL Server database 26, 88
 in Oracle database 26, 95
 replication limitations 25–27
 translation limitations 121
 latency of replication 105

 monitoring 157–161
 Log Transfer Language (LTL) 15, 47, 54, 62, 69, 103, 104–105
 problems with 176–179
 LTM for MVS. *See* Replication Agent for DB2 UDB for OS/390
 LTM locator 104
 See also origin queue ID
LTMLASTCOMMIT table, in DB2 for OS/390 database 36, 74
LTMOBJECTS table, in DB2 for OS/390 database 36

M

MainframeConnect for DB2 UDB 8, 73, 75
 maintaining replication systems 155–156
 Maintenance User, Replication Server 112–115, 123
 Replication Agent filtering transactions 113–114
 transactions 40, 46, 53, 61, 68
 user ID 103, 112–114
 materialization 28–30, 197–208
 atomic bulk 198, 202–205
 datatype translation 199–201
 loading data into replicate database 201
 nonatomic bulk 198, 205–208
 unloading data from primary database 199
 metadata
 of DB2 database 48
 of Informix database 55
 of Microsoft SQL Server database 62–63
 of Oracle database 70
 Microsoft SQL Server data server
 as primary database 56–63
 as replicate database 88–93
 class-level translation scripts 92
 class-level translations 190–192
 decimal datatype 119
 default datatype translations 190–192
 identity columns 121
 @@IDENTITY global variable 57, 60
 image datatype 88
 JDBC driver 58
 log scan position 176
 maximum number of columns 60
 ntext datatype 88

- numeric** datatype 121
- numeric precision 118–120
- primary database configuration 61–62
- primary database limitations 60–61
- primary database permissions 56
- replicate database setup 91–93
- Replication Agent 17, 56–63
- Replication Agent configuration parameters 58–60
- rs_info** table 89
- rs_lastcommit** table 89
- RSM limitations 60, 63
- text** datatype 88
- transaction log 57
- translating primary datatypes 62
- monitoring replication systems 157–161
- MVS operating system, *See* OS/390 operating system
- Replication Agent configuration parameters 65–67
- RS_INFO** table 94
- RS_LASTCOMMIT** table 94
- RSM limitations 68
- selective update triggers 67
- system views 64
- TNS Listener process 65
- transaction log 65
- translating primary datatypes 70
- triggers 67
- Oracle system views 64
- origin queue ID 16, 105
- LTM Locator 104
- OS/390 operating system
 - data-sharing environment 36, 39
- SYGWHOST** macro 38
- outbound queue, troubleshooting 166–168
- owner-qualified object names 24–25

N

- nonatomic bulk materialization 198, 205–208
- NULL, zero-length data replicated as 215

O

- Oracle data server
 - as primary database 63–71
 - as replicate database 94–99
 - BLOB** datatype 95
 - class-level translation scripts 96
 - class-level translations 192–195
 - CLOB** datatype 95
 - DATE** datatype 98
 - default datatype translations 192–195
 - JDBC driver 65, 68
 - log scan position 176
 - LONG** datatype 95
 - LONG RAW** datatype 95
 - maximum number of columns 68
 - NCLOB** datatype 95
 - primary database configuration 68–69
 - primary database limitations 67–68
 - primary database permissions 64
 - replicate database setup 95–97
 - Replication Agent 17, 63–71
 - permissions
 - DB2 for OS/390 primary database 35–36
 - DB2 for UNIX and Windows primary database 42
 - Informix primary database 49
 - Microsoft SQL Server primary database 56
 - Oracle primary database 64
 - Sybase ASA primary database 214
- primary databases 3, 6, 33–71
 - Adaptive Server Anywhere 213–216
 - DB2 for OS/390 35–41, 136–138
 - DB2 for UNIX and Windows 41–49
 - heterogeneous replication issues 21–22
 - Informix 49–56
 - maintenance of 156
 - Microsoft SQL Server 56–63
 - Oracle 63–71
 - origin queue ID 16, 105
 - Replication Server connection 134–135
 - unloading data from 199
- problems with inbound queue 164–166
- problems with outbound queue 166–168
- publications (replication definition) 14

P

Q

- QID (origin queue ID) 16, 105
- queues, Replication Server 164–168
 - inbound 164–166
 - outbound 166–168

R

- reconciling databases 209–211
- Rep Agent User thread 106–107, 108
- replicate databases 3, 8, 72–99
 - DB2 for OS/390 73–78
 - DB2 for UNIX and Windows 78–82
 - heterogeneous replication issues 22–23
 - Informix 83–88
 - loading data into 201
 - maintenance of 156
 - Microsoft SQL Server 88–93
 - Oracle 94–99
 - Replication Server connection 135
 - setting up 27, 83, 88, 94, 141, 216
 - Sybase ASA 216–217
 - troubleshooting 169–170
- Replication Agent 6–7, 15–18, 41, 103–105
 - See also* Replication Agent for DB2 UDB for OS/390; Sybase Replication Agent
 - connect source** permission 102, 108
 - filtering Maintenance User transactions 40, 46, 54, 61, 68
 - for ASA 213–216
 - for DB2 for UNIX and Windows 41–49
 - for Informix 49–56
 - for Microsoft SQL Server 56–63
 - for Oracle 63–71
 - LTL batch mode 105
 - LTM locator 104
 - origin queue ID 16, 105
 - Replication Server connection 107–108
 - RSSD parameters for 39, 40, 47, 54, 62, 69
 - setting up with RSM 141–142
 - transaction log 16
 - using the RSSD 35, 39, 40, 41, 45, 47, 49, 52, 54, 56, 59, 62, 63, 67, 69, 215
- Replication Agent for DB2 UDB for OS/390 17, 35–41

See also Replication Agent; Sybase Replication Agent

- Communications_Protocol** parameter 38
- Date_in_char** parameter 40, 41
- DB2 configuration issues 39–40
- interfaces file 108–109
- LTL problems 177–178
- LTM for MVS 37–41
- LTM_process_maint_uid_trans** parameter 40
- LTMADMIN** user 35
- RS_pw** parameter 38
- RS_source_db** parameter 38, 39
- RS_source_ds** parameter 38, 39
- RS_user** parameter 38
- RSSD_database** parameter 39
- RSSD_pw** parameter 39
- RSSD_server** parameter 39
- RSSD_user** parameter 39
- Sybase Log Extract 39
- TCP/IP communications 37
- Use_repdef** parameter 39, 40
- Replication Command Language (RCL) 13
- replication definitions 14, 47, 54, 62, 69
 - map to** clause 40, 116, 171
- Replication Server 7, 12–15
 - ASA replicate database 217
 - behavior as client 101–102
 - behavior as server 101
 - bidirectional database connection 135, 138
 - communication protocols 102
 - connect source** permission 38, 102, 108
 - connection from Sybase Replication Agent 44–45, 51–52, 59, 66–67
 - create connection** command 40, 46, 53, 61, 68, 102, 108, 110, 112, 124
 - database connections 105–115, 124, 134–138, 144–145
 - database connections with DB2 for OS/390 136–138
 - DSI thread 106, 107, 108, 110–111, 116
 - HDS feature 75–77, 80–82, 85–87, 91–93, 95–97, 99–100, 115–124, 139–141, 172–174
 - heterogeneous replication issues 101–127
 - inbound queue 164–166
 - interfaces file 108–109, 217
 - interfaces* file 74, 84, 89, 95, 111

- LTM locator 104
- maintenance of 156
- Maintenance User 40, 46, 53, 61, 68, 103, 112–115, 123
- materializing subscriptions 28–30, 197–208
- outbound queue 166–168
- primary database connection 134–135
- Rep Agent User thread 106–107, 108
- replicate database connection 135
- Replication Agent connection 107–108
- resume connection** command 108, 111
- rs_db2_char_for_bit** datatype 75
- rs_db2_varchar_for_bit** datatype 75
- rs_dump** command 28
- rs_dumptran** command 29
- rs_get_lastcommit** function 74
- rs_init** command 27
- rs_marker** command 29
- rs_mss_numeric** datatype 121
- rs_msss_numeric** datatype 90
- rs_oracle_date** datatype 98
- rs_oracle_datetime** datatype 98
- rs_subcmp** utility 30, 209–211
- RSI user 103
- RSSD 13, 103
- setting up with RSM 138–140
- Sybase ASA replicate database 217
- SysAdmin user 103
- TCP/IP communications 37, 45, 52, 59, 67
- user IDs 102–103
- Replication Server Manager. *See* RSM
- Replication Server System Database. *See* RSSD
- replication system 3–9
 - administration of 155–161
 - components of 6–9
 - configuration of 147–154
 - database connections 134–138
 - database gateway 8
 - diagram of 4, 5
 - latency 105
 - maintenance of 155–156
 - monitoring of 157–161
 - primary database 3, 6
 - problems with 163–179
 - replicate database 3, 8
 - Replication Agent 6–7
 - Replication Server 7
 - Replication Server Manager 9
 - setting up with RSM 133–146
 - troubleshooting 163–179
- resuming database connections 146
- rs_dump** command 28
- rs_dumptran** command 29
- RS_INFO** table, in DB2 for OS/390 database 73–74
- rs_info** table, in Informix database 83
- rs_info** table, in Microsoft SQL Server database 89
- RS_INFO** table, in Oracle database 94
- rs_init** command 27
- rs_init** utility 122–124
- RS_LASTCOMMIT** table, in DB2 for OS/390 database 74
- rs_lastcommit** table, in Informix database 83
- rs_lastcommit** table, in Microsoft SQL Server database 89
- RS_LASTCOMMIT** table, in Oracle database 94
- rs_lastcommit** table, problems with 171–172
- rs_marker** command 29
- rs_subcmp** utility 30, 209–211
- RSM 9, 19, 48, 55, 62–63, 70, 131–146
 - connectivity 132–133
 - getting metadata with 48, 55, 62–63, 70, 121
 - heartbeat 157–161
 - interfaces file 109
 - interfaces* file 132
 - latency calculations 105
 - limitations of 133
 - limitations with Sybase Replication Agent 45, 49, 53, 56, 60, 63, 68
 - monitoring latency 157–161
 - permissions 132
 - relationship with other servers 131–132
 - setting up databases with 143–145
 - setting up Replication Agents 141–142
 - setting up Replication Server 138–140
 - setting up replication system 133–146
- RSSD 13
 - datatype definitions stored in 121
 - installing UDDs in 123, 139–140
 - Replication Agent using the 35, 39, 40, 41, 45, 47, 49, 52, 54, 56, 59, 62, 63, 67, 69, 215
 - Replication Server user ID 103

S

scripts

- DB2 for OS/390 class-level translations 77
 - hds_db2_connection_sample.sql* 77
 - hds_db2_funcstrings.sql* 76
 - hds_db2_setup_for_replicate.sql* 73, 75
 - hds_db2_udds.sql* 76
 - hds_informix_connection_sample.sql* 87
 - hds_informix_funcstrings.sql* 86
 - hds_informix_setup_for_replicate.sql* 83, 85
 - hds_informix_udds.sql* 86
 - hds_mssql_connection_sample.sql* 92
 - hds_mssql_funcstrings.sql* 91
 - hds_mssql_setup_for_replicate.sql* 88, 91
 - hds_mssql_udds.sql* 91
 - hds_oracle_connection_sample.sql* 97
 - hds_oracle_funcstrings.sql* 96
 - hds_oracle_setup_for_replicate.sql* 94, 95–96
 - hds_oracle_udds.sql* 96
 - hds_udb_connection_sample.sql* 82
 - hds_udb_funcstrings.sql* 80
 - hds_udb_setup_for_replicate.sql* 80
 - hds_udb_udds.sql* 80
 - Informix class-level translations 86
 - Microsoft SQL Server class-level translations 92
 - Oracle class-level translations 96
 - Replication Agent sample setup 49, 56, 63, 71
 - rssetup.sql* 214, 216–217
- stored procedures
- and Oracle system views 64
 - for transaction log 50, 57, 65
 - replication limitations 24
 - replication of 14
- subscriptions 14
- materializing 28–30, 197–208
- Sybase Central 19
- Sybase Log Extract, *See* Replication Agent for DB2 UDB for OS/390
- Sybase Replication Agent 17
- See also* Replication Agent; Replication Agent for DB2 UDB for OS/390
 - DB2 configuration issues 46–47
 - DB2 for UNIX and Windows limitations 45
 - filter_maint_userid** parameter 46, 53, 61, 68
 - for DB2 Universal Database 41–49
 - for Informix data server 49–56

- for Microsoft SQL Server database 56–63
- for Oracle database 63–71
- Informix configuration issues 53–54
- Informix limitations 52–53
- JDBC communications 50, 58, 65
- logging problems 168–169
- LTl batch mode 105
- LTl problems 178–179
- ltl_character_case** parameter 46–47, 54, 61–62, 69
- ltl_origin_time_required** parameter 49, 56, 63, 71
- metadata commands 48, 55, 62–63, 70
- Microsoft SQL Server configuration issues 61–62
- Microsoft SQL Server limitations 60–61
- Oracle configuration issues 68–69
- Oracle limitations 67–68
- pdb_convert_datetime** parameter 47–48, 55, 62, 70
- pdb_xlog** command 45, 53, 60, 68, 133
- pds_connection_type** parameter 43, 51, 58, 66
- pds_database_name** parameter 44, 51, 59, 66
- pds_datasource_name** parameter 44
- pds_host_name** parameter 43, 51, 58, 66
- pds_password** parameter 43, 50, 58, 66
- pds_port_number** parameter 44, 51, 58, 66
- pds_server_name** parameter 44, 51, 58, 66
- pds_username** parameter 43, 50, 58, 65
- primary Replication Server connection 44–45, 51–52, 59, 66–67
- rs_host_name** parameter 44, 51, 59, 66
- rs_password** parameter 44, 51, 59, 66
- rs_port_number** parameter 44, 51, 59, 66
- rs_source_db** parameter 44, 51, 59, 66
- rs_source_ds** parameter 44, 51, 59, 66
- rs_username** parameter 44, 51, 59, 66
- RSM limitations 45, 49, 53, 56, 60, 63, 68
- RSSD parameters for 45, 52, 59, 67
- rssd_database_name** parameter 45, 52, 60, 67
- rssd_host_name** parameter 45, 52, 59, 67
- rssd_password** parameter 45, 52, 60, 67
- rssd_port_number** parameter 45, 52, 59, 67
- rssd_username** parameter 45, 52, 60, 67
- sample setup scripts 49, 56, 63, 71
- setting up with RSM 141–142
- TCP/IP communications 45, 52, 59, 67
- transaction log 50, 57, 65

trigger errors 174–176
troubleshooting 168–169
use_rssd parameter 47, 54, 62, 69
SYGWHOST macro, OS/390 38
syntax conventions xiv
SysAdmin user, Replication Server 103
System Administrator user IDs 114

T

tables
LTMLASTCOMMIT, in DB2 for OS/390 database 74
maximum number of columns 53, 60, 68
problems with **rs_lastcommit** 171–172
RS_INFO, in DB2 for OS/390 database 73–74
rs_info, in Informix database 83
rs_info, in Microsoft SQL Server database 89
RS_INFO, in Oracle database 94
RS_LASTCOMMIT, in DB2 for OS/390 database 74
rs_lastcommit, in Informix database 83
rs_lastcommit, in Microsoft SQL Server database 89
RS_LASTCOMMIT, in Oracle database 94
TCP/IP communications protocol 37, 45, 52, 59, 67
TNS Listener process, Oracle 65
transaction logs 16
DB2 for OS/390 35, 36–37
DB2 for UNIX and Windows 42
Informix 50
Microsoft SQL Server 57
Oracle 65
problems with trigger-based 168–169
scan position 176
stored procedures 50, 57, 65
Sybase ASA 215
triggers 50, 57, 65
triggers
errors firing 174–176
for transaction log 50, 57, 65
selective update 52, 67
troubleshooting
date and **time** translations 171
inbound queues 164–166
outbound queues 166–168
replicate databases 169–170
replication systems 163–179

RSM heartbeat 160–161
Sybase Replication Agent triggers 174–176

U

Universal Database. *See* DB2 for OS/390; DB2 for UNIX and Windows
user IDs
dbmaint user, ASA 214, 216
LTMADMIN user, DB2 for OS/390 35
Replication Server 102–103
SysAdmin user 103
System Administrator 114
user-defined datatypes
installing in RSSD 123
utilities
bcp 204, 207
cs_convert 115
rs_init 122–124
rs_subcmp 30, 209–211

V

V_\$SESSION Oracle system view 64
V_\$TRANSACTION Oracle system view 64