SYBASE®

What's New

Replication Server®

15.0

DOCUMENT ID: DC36431-01-1500-02

LAST REVISED: April 2006

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

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

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

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

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

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

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

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

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

Contents

About This Book ..... **vii**

CHAPTER 1 **New Features in RMS 15.0 EBF nos. 13416, 13417, 13418, 13419, 13425** **1**
- Introduction 2
- New triggers 2
  - Partition size triggers 3
  - Queue size triggers 3
  - Connection latency triggers 3
  - Queue latency triggers 4
  - add event trigger enhancement 4
  - drop event trigger enhancement 6
  - get triggers enhancement 7
- Replication Agent connection consolidation 7
  - add server enhancement 8
  - Displaying dbltm connection name 8
- get connections enhancements 9
- filter connection enhancements 9

CHAPTER 2 **New Features in Replication Server Version 15.0** **11**
- Support for longer identifiers 11
- New datatype: bigint 12
- New unsigned integer datatypes 13
- New Unicode datatype: unitext 14
- Replicating computed columns 15
- Replicating encrypted columns 16
- Replicating partitioned tables 16
- Larger disk partitions 18
- Larger text and image size 19
- Embedded Replication Server System Database (ERSSD) enhancement 19
- New password encryption algorithm 20
- Mixed-version enhanced support 21

New interface for monitors and counters 22
Support for isolation levels 23
Bidirectional replication support for DDL in MSA 25
Batching of commands for non-ASE servers 25
SySAM license management 26

CHAPTER 3 System Changes in Replication Server 15.0 29
Configuration parameters 29
Replication Server parameters 29
Database connection parameters 30
RCL commands 31
System tables 32
System stored procedures 33
Function strings 33
Datatypes 34

CHAPTER 4 New Features in Replication Manager 15.0 35
New user interface features 35
Two-tier management solution 35
Three-tier management solution 36
Replication Manager plug-in replaces Replication Server plug-in 36
Online help 36
Visual monitoring of status 37
Event Log pane 37
Background processing 38
Replication Manager logging enhancement 38
Script editors 39
Replication Manager features supported 40
Support for new datatypes 40
Support for DirectConnect 40
Replication support 41
Routes 41
Troubleshooting tools 42
Connection status hide options 45
Warm standby wizards 46
Thread management 49

CHAPTER 5 Introducing Replication Monitoring Services 51
Introducing Replication Monitoring Services 51
Monitoring servers in the replication environment 52
Software requirements and compatibilities 53

Installation 53
Starting and stopping RMS 53
Connecting to RMS in Sybase Central 54
Monitoring a replication environment using RMS 55
Adding and dropping servers for monitoring in Sybase Central 55
Viewing monitored objects in Sybase Central 57
Setting configuration parameters for monitored replication objects 57
Monitoring a logical group of servers 58
Suspending or resuming components in the replication environment 59
Shutting down monitored servers 60
Generating rollup status for servers 60
Generating latency and heartbeat information 61
Adding event triggers 61

CHAPTER 6 New Features in Replication Server Version 12.6 63
Multi-Site availability (MSA) 63
Support for symmetric multiprocessors (SMP) 64
The embedded RSSD (ERSSD) 66
Performance enhancements 66
Better management of empty transactions 66
Internal commit control for parallel processing 66
New Replication Server configuration parameters 67
New database configuration parameters 67
Changed database configuration parameters 68
date and time datatypes 69
Support for sending encrypted passwords 69
New bulk materialization method 70
Chinese character set (GB18030) support 71

CHAPTER 7 System Changes in Replication Server 12.6 73
Configuration parameters 73
Replication Server parameters 73
Database connection parameters 74
RCL commands 75
System stored procedures 76
Function strings 76
System tables 77
Keywords 77

Index 79

About This Book

Audience

This manual is intended for customers who are installing and using Replication Server® 15.0 and the subsequent EBFs.

How to use this book

This book describes the new Sybase® Replication Server and Replication Manager version 15.0 features. It also provides information about the Replication Monitoring Services (RMS), a new component in Replication Server version 15.0 that provides monitoring capabilities for the replication environment.

This book also provides information about the features and system changes in Replication Server version 12.6.

- Chapter 1, “New Features in RMS 15.0 EBF nos. 13416, 13417, 13418, 13419, 13425,” describes the features developed after the Replication Server 15.0 release.
- Chapter 2, “New Features in Replication Server Version 15.0,” describes the Replication Server version 15.0 features.
- Chapter 3, “System Changes in Replication Server 15.0,” describes new and changed commands, stored procedures, configuration parameters, system tables, keywords, and so on in Replication Server version 15.0.
- Chapter 4, “New Features in Replication Manager 15.0,” describes the new features introduced with Replication Manager version 15.0.
- Chapter 5, “Introducing Replication Monitoring Services,” introduces the new component in Replication Server 15.0. The Replication Monitoring Services is a middle-management layer for monitoring large and complex replication environments.
- Chapter 6, “New Features in Replication Server Version 12.6,”describes new and changed commands, stored procedures, configuration parameters, system tables, keywords, and so on in Replication Server version 12.6.
- Chapter 7, “System Changes in Replication Server 12.6,” describes the Replication Server 12.6 features.

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 Sybase Product Manuals Web site. To check for critical product or document information that was added after the release of the product CD, use the Getting Started CD.

- *Installation Guide* for your platform – describes installation and upgrade procedures for all Replication Server and related products.
- *What's New in Replication Server?* (this book) – describes the new features in Replication Server version 15.0 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 Manager plug-in help, which contains information about using Sybase Central™ to manage Replication Server.

Other sources of information

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

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

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

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

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

Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

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

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

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

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

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

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

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

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

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

2 Click MySybase and create a MySybase profile.

Sybase EBFs and software maintenance

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

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

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

3 Select a product.

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

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

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

Conventions

The following style conventions are used in this book:

- In a sample screen display, commands 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 in italics:

  /usr/u/sybase

- The names of programs, utilities, procedures, and commands appear in this type:

  bcp

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

Table 1: Syntax conventions

Key	Definition
command	Command names, command option names, utility names, utility flags, and other keywords are in Arial.
variable	Variables, or words that stand for values that you fill in, are in italics.
{ }	Curly braces indicate that you choose at least one of the enclosed options. Do not include braces in your option.
[ ]	Brackets mean 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 to be typed as part of the command.

Accessibility features

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

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

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

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

If you need help

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

CHAPTER 1

New Features in RMS 15.0 EBF nos. 13416, 13417, 13418, 13419, 13425

This chapter describes the new features introduced with Replication Server 15.0 EBF nos. 13416, 13417, 13418, 13419, and 13425.

Topic	Page
Introduction	2
New triggers	2
Replication Agent connection consolidation	7
get connections enhancements	9
filter connection enhancements	9

Introduction

This EBF contains features developed after the 15.0 release. It includes features needed to support non-ASE replication, as well as additional features. This EBF *completely* replaces the GA version of RMS.

Backward compatibility with Replication Manager (RM) plug-in GA

This EBF supports the GA version of the Replication Manager plug-in. The Replication Manager plug-in can connect to the GA versions of both the RMS and this EBF, and it can function without throwing any exceptions (P1 errors).

To support the GA version of the Replication Manager plug-in, the result sets and RMI objects that RMS returns have not been modified.

Note Neither the RMS nor the Replication Manager plug-in are required to be compatible with all versions of each other. Future versions of the RMS may not provide backward compatibility.

New triggers

This EBF introduces four new triggers:

- Partition size
- Queue size
- Connection latency
- Queue latency

Use add *event* trigger to add these triggers to RMS. You can set up the new triggers to execute a script when the specified threshold is exceeded, or when the size or latency falls below the threshold. Optionally, you can set up the triggers to continuously execute at each monitoring cycle until the trigger is deleted, or until the current size or latency meets the threshold requirement.

To drop these triggers from RMS, use drop *event* trigger.

Partition size triggers

RMS monitors the size of an individual partition or the total size of all partitions in a Replication Server. The partition size trigger executes when the size of a partition does not meet the defined threshold. The partition size threshold is measured as a percentage of disk space used.

Queue size triggers

RMS monitors the size of individual queues in a Replication Server but not the total size of all queues. The queue size trigger executes when the size of queue does not meet the defined threshold. The threshold for a queue is measured in megabytes.

Connection latency triggers

Connection latency triggers notify the user when the defined latency threshold between two database connections is not met. Connection latency is the time it takes a transaction to replicate from the primary database to the replicate database. It is the difference between the destination date and time fields and the origin date and time fields stored in rs_lastcommit table in the replicate database. If the information is not available, the RMS writes an error message to the log and cancels the latency calculation.

Queue latency triggers

Queue latency triggers notify the user when the first block of a queue has remained at the beginning of the queue for more than the specified threshold. RMS monitors each inbound and outbound (DSI) queue, but does not track latency for RSI or materialization queues.

Note RMS, not the Replication Server or data servers, tracks the queue latency. It is therefore important that you coordinate the Replication Server monitoring interval and the latency threshold when you are setting up a queue latency trigger. For example, if the monitoring interval is set to 5 minutes and the queue latency threshold is set to 1 minute, RMS checks latency only after 5 minutes. Since RMS must compare the current time with the previous time, it may take 2 cycles, or 10 minutes, to notice that the queue latency threshold has been exceeded.

---

add event trigger enhancement

The enhanced add *event* trigger command supports the four new triggers. For more information on add *event* trigger, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Syntax

The modified add *event* trigger syntax is:

```
add {status | latency | size} trigger
      [{connection | logical connection | route | queue | rep agent |
        partition} [component_name]]
      [with primary primary_connection]
for server_name
{status changes to state |
  size {exceeds | falls below} size_threshold |
  latency {exceeds | falls below} latency_threshold}
[wait wait_interval]
[continuous continuous_flag]
execute command
```

Parameters

The new parameters are:

- *component_name* – in RMS 15.0, the component name is required if a component type is provided. In this EBF, you can create a size trigger for the sum of all partitions. add *event* trigger includes a type partition, but not the partition's name.

- with primary *primary_connection* – identifies the primary connection for a connection latency trigger. The trigger executes the script if the latency threshold between the primary connection and the replicate connection is not satisfied.
- size exceeds, falls below *size_threshold* – indicates whether the trigger should execute when the size exceeds the threshold or when it falls below the threshold.
- latency exceeds, falls below *latency_threshold* – indicates whether the trigger should execute when the latency exceeds the threshold or when it falls below the threshold.

Examples

Here are examples using the new triggers in RMS:

- Adds a trigger to the Replication Server INVENTORY_RS partition "p1" that executes the script *email.sh* when the partition usage exceeds 80 percent:

```
add size trigger partition p1 for INVENTORY_RS
    size exceeds 80
    continuous true
    execute /sybase/RMS/scripts/email.sh
```

- Adds a trigger to the Replication Server INVENTORY_RS that executes the script *email.sh* when the sum of all partition usage exceeds 75 percent:

```
add size trigger partition for INVENTORY_RS
    size exceeds 75
    execute /sybase/RMS/scripts/email.sh
```

- Adds a trigger to the queue "inventory_pds.vendor(Inbound)" of Replication Server INVENTORY_RS that executes the script *email.sh* when the queue size falls below 100 megabytes:

```
add size trigger queue inventory_pds.vendor(Inbound)
    for INVENTORY_RS
    size falls below 100
    continuous true
    execute /sybase/RMS/scripts/email.sh
```

- Adds a trigger to the replicate connection "inventory_rds.vendor" of replicate Replication Server INVENTORY_RS that will execute the script *email.sh* when the latency from the primary connection "inventory_pds.vendor" exceeds 5 minutes (300 seconds):

```
add latency trigger connection inventory_rds.vendor
    with primary inventory_pds.vendor
    for INVENTORY_RS
```

```
    latency exceeds 300
    execute /sybase/RMS/scripts/email.sh
```

Usage

- You must set for *server_name* to the name of the replicate Replication Server when adding a latency connection trigger. In the following example, INVENTORY_RS is the replicate Replication Server:

  ```
  add latency trigger connection inventory_rds.vendor
      with primary invetory_pds.vendor
      for INVENTORY_RS
      latency exceeds 300
       execute /sybase/RMS/scripts/email.sh
  ```

- You must set the configuration parameter ltl_origin_time_require to "true" when setting up a latency connection trigger where the primary connection is from a Replication Agent or MRA. To set the parameter, connect to the Replication Agent or MRA and execute:

  ```
  ra_config ltl_origin_time_required, true
  ```

drop event trigger enhancement

The enhanced drop *event* trigger command supports the four new triggers. For more information on drop *event* trigger, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Syntax

The modified drop *event* trigger syntax is:

```
drop {status | latency | size} trigger
    [{connection | logical connection | route | queue | rep agent |
      partition} [component_name]]
    [with primary primary_connection]
    for server_name
    {status changes to state |
    size {exceeds | falls below} size_threshold |
    latency {exceeds | falls below} latency_threshold}
```

Parameters

The new parameters are:

- with primary *primary_connection* – identifies the primary connection of the latency connection trigger to drop. This parameter is required when dropping a latency connection trigger.
- size exceeds, falls below *size_threshold* – indicates the size trigger to drop.
- latency exceeds, falls below *latency_threshold* – indicates latency trigger to drop.

Examples

- Drop a partition size trigger:

```
drop size trigger partition p1
   for INVENTORY_RS
   size exceeds 80
```

- Drop a latency connection trigger:

```
drop latency trigger
   connection inventory_rds.vendor
   with primary inventory_pds.ventory
   for INVENTORY_RS
   latency exceeds 300
```

get triggers enhancement

The get triggers result set now returns primary connection.

Column	Description
Primary Connection	The name of the primary connection.

Replication Agent connection consolidation

The Replication Agent connection status is derived from the status of the Replication Agent thread in the Replication Server and the status of the Replication Agent process that is extracting data from the primary database. The Replication Agent process is either a RepAgent thread in ASE, a Replication Agent or a Mirror Replication Agent (MRA) server, or SQL Anywhere LTM, also known as dbltm, for an ASA or IQ server.

In RMS 15.0, the status of the Replication Agent thread in the Replication Server is determined only by the Replication Agent connection status. This EBF integrates the Replication Agent process to determine the complete status. When the Replication Agent thread or the Replication Agent process is down, the RMS returns the state “SUSPENDED” and a description.

The format of the connection name is *ds*.*db*, where *ds* is the name of the data server and *db* is the name of the database. The Replication Agent, MRA, and dbltm store the name of the connection in configuration parameters.

add server enhancement

The enhanced add *server* command supports dbltm servers.

Syntax

The modified add *server* syntax is:

```
add {ASA | ASE | DirectConnect | IQ | Replication Agent | MRA |
Replication Server | RMS | Open Server | dbltm} server_name
    set username [to] user
    [set password [to] passwd]
    [set charset [to] charset]
    [set language [to] lang]
    [set rssd_username [to] rssd_user]
    [set rssd_password [to] rssd_passwd]
    [set rssd_charset [to] rssd_charset]
    [set rssd_language [to] rssd_lang]
    [set monitoring [to] {'true' | 'false'}]
    [set interval [to] interval]
    [set connection_ds [to] ds]
    [set connection_db [to] db]
```

Parameters

The new parameters are:

- *ds* – specifies the name of the primary data server. The dbltm sends *ds*.*db* to the Replication Server when replicating transactions. The *ds* must match the server name used in the Replication Server connection.
- *db* – specifies the name of the primary database. The dbltm server sends *ds*.*db* to the Replication Server when replicating transactions. The *db* must match the database name used in the Replication Server connection.

Note The *ds* and *db* parameters are optional and are valid only for a dbltm server.

Displaying dbltm connection name

This EBF provides the ability to retrieve the dbltm connection name that the user provided when the server was added to the environment. RMS returns the DBMS type for a Replication Agent and a DirectConnect server.

The RMS returns the required information as RMS configuration parameters. To retrieve the connection server and database name, you can issue the following command:

```
configure server customer_ltm RMS
```

get connections enhancements

The enhanced get connections API supports retrieving connections that are associated with a data server or a Replication Agent process. It supports servers other than a Replication Server:

- ASE – get connections returns the connection information for each database in the ASE. The RMS searches all of the Replication Servers in the RMS looking for connections named *ASE_name.database*.
- Replication Agent/MRA – get connections returns the information for the primary connection associated with the Replication Agent. The name of the connection associated with the Replication Agent or MRA is stored in the configuration parameters rs_source_ds and rs_source_db. get connections searches all of the Replication Servers in the RMS to find the connection.
- dbltm – get connections returns the information for the primary connection associated with the dbltm. The connection information for the dbltm is optionally provided when the server is added to the environment. If the information is not available, get connections returns an empty result set and writes a warning message to the RMS log indicating the information is missing.
- DirectConnect – get connections returns the information of all of the connections where the data server matches the name of the DirectConnect server.
- ASA/IQ – get connections returns the information where the data server matches the name of the ASA or IQ server. ASA or IQ server does not use database names.

filter connection enhancements

The implementation of filter *connection* in RMS 15.0 is not consistent with the implementation in the Replication Manager plug-in. RMS 15.0 does not return the status of the connection, and the status of the connection is not rolled up into the status of the Replication Server when a connection filter is turned on. To be consistent with the Replication Manager plug-in implementation, this EBF displays the connection status as “Hidden” when the connection filter is turned on. The status of the connection is still not rolled up into the status of the Replication Server.

CHAPTER 2

New Features in Replication Server Version 15.0

This chapter describes the new features introduced with Replication Server 15.0.

Topic	Page
Support for longer identifiers	11
New datatype: bigint	12
New unsigned integer datatypes	13
New Unicode datatype: unitext	14
Replicating computed columns	15
Replicating encrypted columns	16
Replicating partitioned tables	16
Larger disk partitions	18
Larger text and image size	19
Embedded Replication Server System Database (ERSSD) enhancement	19
New password encryption algorithm	20
Mixed-version enhanced support	21
New interface for monitors and counters	22
Support for isolation levels	23
Bidirectional replication support for DDL in MSA	25
Batching of commands for non-ASE servers	25
SySAM license management	26

Support for longer identifiers

The limit of 30 bytes for selected database and replication object names (identifiers) is extended to 255 bytes for Replication Server version 15.0. Longer identifiers are supported for these objects:

- Tables.

- Columns.
- Stored procedures.
- Parameters – for Replication Server functions and Adaptive Server stored procedures.
- Functions – for function replication definitions.
- Function strings.
- Replication definitions – including table replication definitions, function replication definitions, and database replication definitions.
- Publications.
- Articles.

All other database and replication object identifiers, login names and passwords, transaction names, subscription names, database names, and server names retain the existing 30-byte limit.

Parameter names for Replication Server functions and Adaptive Server stored procedures are the only identifiers that can begin with the @ character.

- Replication Server function parameter names can be as many as 256 bytes including the @ character.
- Adaptive Server stored procedure parameter names can be as many as 255 bytes including the @ character.

Note rs_subcmp supports long identifiers for both table and column names.

Limitation

The create function, alter function, and drop function commands do not support long identifiers. The name of the function and the parameters of these commands cannot exceed 30 bytes.

New datatype: *bigint*

Replication Server 15.0 adds support for the bigint datatype. bigint is a fixed-width, 8-byte datatype. Like int, smallint, and tinyint, bigint is a signed integer datatype. bigint can hold whole numbers between -2^{63} and $+(2^{63} - 1)$.

Table 2-1 describes the range and storage size for all integer datatypes.

Table 2-1: Signed integer datatypes

Integer datatypes	Range	Storage (bytes)
bigint	Whole numbers between -2^{63} and $+(2^{63} - 1)$, inclusive (–9,233,372,036,854,775,808 and +9,233,372,036,854,775,807)	8
int	Whole numbers between -2^{31} and $+(2^{31} - 1)$, inclusive (–2,147,483,648, and +2,147,483,647)	4
smallint	Whole numbers between -2^{15} and $+(2^{15} - 1)$, inclusive (–32,768 and +32,767)	2
tinyint	Positive whole numbers between 0 and 255, inclusive	1

Mixed-version issues

To fully support bigint, the primary and replicate Replication Server must have a site version of 15.0, and the LTL version must be 700. If the LTL version is less than 700 at connect-source time, RepAgent converts bigint columns to numeric.

Note rs_subcmp supports the new bigint datatype.

New unsigned integer datatypes

Replication Server version 15.0 supports four new unsigned integer datatypes. You can use these unsigned datatypes in the same way as their signed equivalents:

- unsigned tinyint – tinyint
- unsigned smallint – smallint
- unsigned int – int
- unsigned bigint – bigint

Signed integers are whole positive or negative numbers. Unsigned integers are only whole positive numbers. The storage sizes of the signed and unsigned integers datatypes are the same. Table 2-2 shows the range of positive numbers supported for each unsigned datatype.

Table 2-2: Unsigned integer datatypes

Integer datatypes	Range	Storage (bytes)
unsigned bigint	Whole numbers between 0 and 18,446,744, 073, 709,551,615, inclusive	8
unsigned int	Whole numbers between 0 and 4,294,967,295, inclusive	4

Integer datatypes	Range	Storage (bytes)
unsigned smallint	Whole numbers between 0 and 65535, inclusive	2
unsigned tinyint	Positive whole numbers between 0 and 255, inclusive	1

Mixed-version issues

To fully support the unsigned integer datatypes, the primary and replicate Replication Server must have a site version of 15.0, and the LTL version must be 700.

If the LTL version is less than 700 at connect-source time, RepAgent makes these datatype conversions:

- unsigned bigint –> numeric
- unsigned int –> numeric
- unsigned smallint –> int
- unsigned tinyint –> tinyint

Replication definitions created with unsigned integer datatypes are not sent to Replication Servers version 12.5.x and earlier. If a replication definition is subscribed by Replication Server version 12.5.x, it cannot be altered to add unsigned integer columns.

Note rs_subcmp supports the new unsigned integer datatypes.

New Unicode datatype: *unitext*

Replication Server version 15.0 adds support for the unitext datatype. unitext is a variable-width, nullable Unicode datatype. Although independent of the text datatype, unitext mirrors its behavior, and can be used wherever text is used.

Like the data in the other Unicode datatypes—unichar and univarchar—unitext data is encoded in UTF-16, which is essentially a 2-byte, fixed-width encoding of Unicode. unitext can hold as many as 1,073,741,823 Unicode characters, or the equivalent of 2,147,483,647 bytes. unitext has no connection to the default character set ID or the default sort order.

The main advantage of Unicode datatypes is efficiency. The UTF-16 character types are approximately 33% more space efficient than UTF-8 for Asian characters.

Mixed-version issues

To fully support unitext, the primary and replicate Replication Server must have a site version of 15.0, the route version must be 15.0, and the LTL version must be 700. If the LTL version is less than 700 at connect-source time, RepAgent converts unitext columns to image.

Note rs_subcmp supports the new unitext datatype.

Replicating computed columns

Computed columns allow you to create an expression and place the result of the expression in a table column. A computed column is:

- Materialized – when its value is computed for each insert or update. Materialized computed columns are stored in the same way as regular columns.
- Virtual – when its value is computed only when referenced in a query. Virtual computed columns are not stored in the table or index page.

A computed column expression is:

- Deterministic – when its value is the same each time it is evaluated.
- Nondeterministic – when its value may be different each time it is evaluated (for example, a date stamp).

Replication Server replicates materialized computed columns in DML statements in the same way it replicates other columns; it does not replicate virtual computed columns.

The replication of computed columns is supported by function strings. In Replication Server version 15.0, the class-level function string rs_set_dml_on_computed is applied at the replicate database DSI when a connection is established. It issues set dml_on_computed “on” after the use database statement. If the replicate Adaptive Server is version 12.5.x or earlier, the command is ignored.

Since Replication Server does not distinguish between computed and regular columns, there are no changes to the syntax for creating or altering replication definitions.

When creating or altering replication definitions for tables containing:

- Deterministic columns – you can choose whether to include those columns in the replication definition. Since deterministic columns always realize the same value, you can create the replication definition without them and allow each replicated insert and update to compute values at the replicate database.
- Nondeterministic columns – you must include nondeterministic computed columns in the replication definition to ensure that the primary and replicate databases remain synchronized.

rs_subcmp support for computed columns

rs_subcmp supports comparison and reconciliation of materialized computed columns; it does not support virtual columns.

If the supported Adaptive Server provides the set dml_on_computed "on" command, rs_subcmp inserts and updates deterministic and nondeterministic materialized columns as regular columns.

Replicating encrypted columns

Replication Server 15.0 supports replication of encrypted columns in Adaptive Server 15.0. Similar to Adaptive Server, Replication Server does not support encrypted columns that contain text and image data.

Replication Server replicates encrypted data as well as the encryption keys. For more information about this feature, see the *New Features Adaptive Server Enterprise 15.0 with Encrypted Columns*.

Note rs_subcmp supports replication of encrypted columns in Adaptive Server.

Replicating partitioned tables

In Replication Server 15.0, partitioned tables introduced in Adaptive Server Enterprise15.0 are replicated in a way that is similar to nonpartitioned tables. The rs_truncate system function in LTL, and the rs_truncate function and function string in Replication Server have been extended to support partitioned tables.

rs_truncate changes in LTL

In LTL, the rs_truncate function has been extended to include partition names, as shown in the following syntax:

```
distribute command_tags applied [owner=owner_name]
table.rs_truncate [partition_name,
[partition_name]...] yielding
```

You must assign a *partition_name* for each partition specified in the truncate table partition command.

rs_truncate function changes in Replication Server

In Replication Server, the rs_truncate function has been extended to accept parameters to support a new truncate table paritition command, shown as follows:

```
truncate table table_name partition partition_name
```

When a truncate partition command is issued, the RepAgent sends the following LTL:

```
applied [owner=owner_name] table_name.rs_truncate,
partition_name _yd
```

rs_truncate function string changes in Replication Server

The partition names are passed as parameters to the rs_truncate function. The rs_truncate function string accepts position-based function-string parameters. The following is a position-based variable that specifies the parameter position in the function, in the LTL command:

```
?n!param?
```

A sample function string for rs_truncate with the position-based variable is as follows:

```
truncate table publishers partition ?1!param?,
?2!param?
```

Examples

Example 1: To replicate truncate table partition as a delete command, alter the function string in the following way:

```
alter function string publisher.rs_truncate
for rs_sqlserver_function_class
output language
'begin transaction
if (?1!param? = '''')   /* NO parameter */
delete publisher
if (?1!param? = ''A'')
delete publishers where c1 < 1000
if (?1!param? = ''B'')
delete publishers where c1 >= 1000
commit transaction'
```

Example 2: To not to truncate table partitions at the replicate server, alter the function string to do nothing if there is a parameter, in the following way:

```
alter function string publisher.rs_truncate
for rs_sqlserver_function_class
output language
'if (?1!param? = '''') delete publisher'
```

Mixed-version issues

For the RepAgent to send an rs_truncate applied subcommand with parameters, the site version must be 15.0 and the LTL version must be 700. If the LTL version is below 700, RepAgent skips the rs_truncate portion of the distribute command.

To replicate the rs_truncate function with a parameter to the replicate Replication Server, the route version must be 1500. If the route version is lower than 1500, the rs_truncate command with the parameter is skipped.

Larger disk partitions

Replication Server version 15.0 extends the maximum disk partition size from 2GB to 1TB. The new limit is applicable after the Replication Server site version is updated to 15.0.

Two new RCL commands support larger disk partitions:

- create partition – makes a partition available to Replication Server. This command replaces the existing add partition command.
  add partition is still supported, for backward compatibility with earlier versions. The syntax and usage of the two commands are identical. The command name has been changed to be consistent with other Replication Server command names.
- alter partition – increases the size of a partition. The syntax is:

  alter partition *logical_partition_name* [expand [size = *size*]]

  For example, to increase the size of the logical partition p2 by 50MB, enter:

  ```
  alter partition p2 expand size = 50
  ```

Larger text and image size

In Replication Server 15.0, the rs_subcmp utility extends the maximum text and image size from 32KB to 2GB. You can specify the new value as length in kilobytes, doing either of:

- Using the -L *text_image_length_in_kilobytes* parameter to set the new text and image length in kilobytes. For example, to set the new length to 64KB, enter:

  ```
  rs_subcmp -L 64 -f subcmp.cfg
  ```

- Providing the new value for text and image length to the parameter TXT_IMG_LEN in the configuration file. Start rs_subcmp using the modified configuration file:

  ```
  rs_subcmp -f subcmp.cfg
  ```

  subcmp.cfg file includes the following configuration parameter and its value:

  ```
  TXT_IMG_LEN = 64
  ```

Note If the value is specified both in the command line and the configuration file, the value in the command line overwrites the value in the configuration file.

Embedded Replication Server System Database (ERSSD) enhancement

ERSSDs were introduced in Replication Server version 12.6. For more information, see Chapter 6, “New Features in Replication Server Version 12.6.”

You can create a route from a Replication Server with an ERSSD, as long as both the source and the destination servers are version 15.0 or later.

To create a route from Replication Server with an ERSSD, use the create route command. Verify that the Replication Agent™ name is in the Replication Server interfaces file; an ERSSD Replication Agent is started as an open server during create route. If the Replication Agent name does not appear in the interfaces file, the command fails.

The default ERSSD Replication Agent name is *erssd_name*_ra. To replace the default name with that of your Replication Agent server, enter:

```
configure replication server
set erssd_ra to <value>
```

Note Sybase provides ERSSD in Adaptive Server Anywhere (ASA) as an option, as well as continues to support the traditional RSSD in Adaptive Server® Enterprise.

New password encryption algorithm

Replication Server 15.0 uses the FIPS-certified Advanced Encryption Standard (AES) algorithm to encrypt new Replication Server user passwords. The AES alogrithm uses the 128-bit encryption key and can be obtained from the Certicom Security Builder library.

Migrating existing encrypted passwords

Use the information in Table 2-3 to migrate existing encrypted passwords in the Replication Server configuration file, and the rs_users and rs_maintusers tables.

Table 2-3: Commands to encrypt passwords in new algorithm

To	Command
Migrate existing user passwords to the new alogrithm	`alter user user set password password` where: • *user* is the login name of the existing user. • *password* is the existing password you want to encrypt using the new alogorithm.
Migrate existing database maintenance user passwords to the new algorithm	`alter connection to data_server.database set password to password` where *password* is the existing password you want to encrypt using the new alogorithm.
Migrate existing route user passwords to the new algorithm	`alter route to dest_replication_server set password to passwd` where: • *dest_replication_server* is the name of the kdestination Replication Server. • *passwd* is the existing password you want to encrypt using the new alogorithm.
Migrate existing user passwords in the configuration file to the new algorithm	Use rs_init to encrypt the passwords using the new algorithm.

Mixed-version issues

To fully support the new password encryption algorithm, both the Replication Server and the rs_init utility must have a site version of 15.0. If the site version is lower than 15.0, an error message displays and encryption is disabled.

Mixed-version enhanced support

In mixed-version environments, interaction between Replication Servers of different versions is restricted to the capabilities of the oldest version. Information associated with new features may not be available to Replication Servers of earlier versions.

Features in the new version, when sent to downstream Replication Servers of earlier versions, can break the multi-site availability (MSA) replication and must be filtered out. In Replication Server 15.0, mixed-version support has been enhanced to provide a mechanism in which the later-version feature data is filtered out by the primary Replication Server before sending it out to the downstream Replication Servers of earlier versions.

New configuration parameter

To enable the primary Replication Server to block certain new feature commands not supported by earlier versions of Replication Server, a new Replication Server configuration parameter, dist_stop_unsupported_cmd, is available for all server and connection-level commands.

Replication Server commands that include the new configuration parameter are described in Table 2-4.

Table 2-4: Modified commands for Mixed-version support

Replication Server command	Syntax
configure replication server	`configure replication server set dist_stop_unsupported_cmd to [on | off]`
alter connection	`alter connection srv.db set dist_stop_unsupported_cmd to [on | off]`
alter logical connection	`alter logical connection lsrv.ldb set dist_stop_unsupported_cmd to [on | off]`

By default, dist_stop_unsupported_cmd is set to off. When this parameter is set to on, DIST suspends itself if a command cannot be sent to some destination Replication Servers. You have to resume the DIST by skipping the entire transaction or by resetting this parameter to off.

When the parameter is set to off, the distributor (DIST) thread skips the newer commands to earlier version Replication Servers.

New interface for monitors and counters

Replication Server 15.0 provides a new, simpler interface for monitoring the counters that provide performance information. By default, all counters are inactive until you turn them on using either:

- The admin stats (admin statistics) command, which activates counters for a specific time period, or
- The Replication Server configuration parameter stats_sampling, which is a toggle that activates or deactivates counters.

Using admin stats, you can specify which counter statistics to report, whether to display those statistics on the screen or save them to the RSSD, and how many seconds to collect the statistics. admin stats lets you specify statistics for:

- Individual counters
- Individual modules
- The sysmon counters, which are a set of counters Sybase has identified as those most valuable for monitoring performance
- All counters

Viewing statistics on screen provides a point-in-time benchmark. Saving statistics to the RSSD lets you accumulate data, so you can see changes in statistics over time, and perform averages and other calculated values. You can specify a sampling period and the number of observations during that sampling period.

When you turn on counter activity using stats_sampling, the counters stay active until you turn them off.

New admin stats options are:

- admin stats, backlog – reports the current backlog in the inbound and outbound stable queues.
- admin stats, { tps | cps | bps } – reports throughput in transactions per second, commands per second, or bytes per second.
- admin stats, status – reports configuration information for counters.

Support for isolation levels

Isolation levels let you control the degree to which data can be accessed by other users during a transaction. With version 15.0, Replication Server decouples isolation levels and serialization methods for the replicate data server, and enables all isolation levels for replicate data servers that Replication Server supports. In versions earlier than 15.0, Replication Server supported only isolation level 3.

Through the use of custom function strings, Replication Server supports all isolation levels the replicate data servers may use. Support is not limited to the ANSI standard only.

Each isolation level specifies the types of actions that are not permitted while concurrent transactions are processing. Higher levels include the restrictions imposed by lower levels.

You can set the isolation level with the database configuration parameter dsi_isolation_level. The ANSI standard levels supported by Adaptive Server are:

- 0 – ensures that data written by one transaction represents the actual data.
- 1 – prevents dirty reads and ensures that data written by one transaction represents the actual data.
- 2 – prevents nonrepeatable reads and dirty reads, and ensures that data written by one transaction represents the actual data.
- 3 – prevents phantom rows, nonrepeatable reads, and dirty reads, and ensures that data written by one transaction represents the actual data.

The default value is the current transaction isolation level for the target data server.

For example, to select isolation level 2 for the Replication Server connection to the TOKYO_DS data server and pubs2 database, enter:

```
alter connection to TOKYO_DS.pubs2
set dsi_isolation_level to '2'
```

In versions earlier than 15.0, you set isolation level 3 and the serialization method at the same time:

```
alter connection to TOKYO_DS.pubs2
set dsi_serialization_method to 'isolation_level_3'
```

With version 15.0, the equivalent is:

```
alter connection to TOKYO_DS.pubs2
set dsi_serializaiton_method to 'wait_for_start'

alter connection to TOKYO_DS.pubs2
set dsi_isolation_level to '3'
```

Bidirectional replication support for DDL in MSA

You can configure multi-site availability (MSA) to set up a two-way replication of data definition language (DDL) transactions between two Adaptive Server databases.

In version 15.0, bidirectional DDL replication support in MSA environments is supported only for non-warm standby databases.

Replication Server 15.0 supports this bidirectional replication using a new configuration parameter dsi_replication_ddl. When dsi_replication_ddl is set to on, DSI sends set replication off to the replicate database, which instructs it to mark the succedent DDL transactions available in the system log not to be replicated. Therefore, these DDL transactions are not replicated back to the original database, which enables the DDL transactions replication in bidirectional MSA replication environment.

❖ **Setting up bidirectional replication**

1 Create a bidirectional MSA replication environment. For steps, see Chapter 12, "Managing Replicated Objects Using Multi-Site Availability," in the *Replication Server Administration Guide Volume 1*.

2 Grant "set session authorization" privilege to a maintenance user on the destination database, as shown in the following example:

```
grant set session authorization to maint_user
```

3 In the destination database, set dsi_replication_ddl to on to enable bidirectional DDL replication, as shown in the following example:

```
alter connection to dataserver.database set dsi_replication on
```

4 Replicate DDL transactions.

Batching of commands for non-ASE servers

Replication Server 15.0 allows you to batch commands for non-ASE database servers. By batching commands, you may be able to achieve improved performance in Replication Server.

New function strings

Support for command batching to non-ASE servers is achieved through the use of two function strings, rs_batch_start and rs_batch_end. These function strings store the SQL translation needed for marking the beginning and end of command batches. Use of these function strings is not necessary for Adaptive Server Enterprise or any other data server where the function strings rs_begin and rs_commit already support the needed functionality.

New DSI connection parameter

A DSI connection parameter, use_batch_markers, is used to control the processing of the two function strings, rs_batch_start and rs_batch_end. You can set use_batch_markers with alter connection and configure connection commands. If use_batch_markers is set to on the function strings rs_batch_start and rs_batch_end are executed. The default is off.

New DSI configuration parameters

DSI must be correctly configured for batching of commands to the replicate data server. There are three DSI configuration parameters to consider for each connection that will be batching commands:

- batch
- batch_begin
- use_batch_markers

For more information about configuring DSI for non-ASE data servers, see Chapter 5, “Replication Server Issues,” in the *Replication Server Heterogeneous Replication Guide*.

SySAM license management

The Sybase Software Asset Management (SySAM) implementation has changed for this version of Replication Server. The changes include:

- Asset management and reporting tools are provided with SySAM version 2.0. These tools allow you to monitor license usage and compliance.
- A single installation method supports all editions of Replication Server.
- SySAM configuration is no longer optional.
- Flexible SySAM configuration options.
- SySAM licenses are no longer shipped along with order fulfillment. You must obtain license certificates from the Sybase Product Download Center (SPDC).

- SySAM license keys include information about the support plan you purchased. You must update these licenses whenever you renew your support plan.
- Licensing policies are strictly and consistently enforced.
- Replication Server can function under grace periods if cannot obtain a license. These grace periods allow you reasonable time to respond to the issues causing license failure. During the grace period, the Replication Server continues to function normally. However, the Replication Server or any of its features, will shut down at the end of the grace period if licensing issues remain unresolved.
- You can receive real-time e-mail notifications about licensing events.
- Licenses issued from SPDC include information about the host machine where the licenses will be deployed. These licenses cannot be used on another machine without being reissued from SPDC.

Note Replication Server 15.0 does not require the REP_SSL license, as SSL now comes as a part of the basic REP_SERVER license.

These changes affect the Replication Server installation and configuration process. For pre-installation and SySAM installation information, see the *Replication Server Installation Guide* for your platform.

Plan your SySAM deployment before installing Replication Server.

CHAPTER 3

System Changes in Replication Server 15.0

This chapter describes system changes introduced with Sybase Replication Server version 15.0.

Topic	Page
Configuration parameters	29
RCL commands	31
System tables	32
System stored procedures	33
Function strings	33
Datatypes	34

Configuration parameters

This section summarizes new and changed configuration parameters for Replication Server 15.0.

Replication Server parameters

Table 3-1 describes new Replication Server configuration parameters that are stored in the rs_config system table and can be set using configure replication server.

Table 3-1: New Replication Server parameters for rs_config table

Parameter	Description
errsd_backup_interval	Interval between backups of database and log. Specified as “nn hours” or “nn minutes” or “nn seconds.” Default: 24 hours.

Parameter	Description
errsd_backup_path	Location of stored backup files. Should be a full directory path. Configuring this path causes immediate backup. Default: Same directory as the transaction log mirror; initial value specified in rs_init.
errsd_backup_start_time	Time the backup starts. Specified as: “hh:mm AM” or “hh:mm “M” using a 12-hour clock, or “hh:mm” using a 24-hour clock. Default: 01:00 AM.
erssd_backup_start_date	Date the backup begins. Specified as “MM/DD/YYYY.” Default: current date.
erssd_ra	Configures Replication Agent name, to create a route from the current site to another Replication Server. This server name must exist in the interfaces name. *erssd_name*_ra

Table 3-2 describes changed Replication Server configuration parameters that are stored in the rs_config system table and can be set using configure replication server.

Table 3-2: Changed Replication Server parameters for rs_config table

Parameter	Description
stats_daemon_sleep_time	Parameter not supported in this version.
send_enc_pw	Renamed to send_enc_pw. The description remains the same, which is as follows: Ensures that Replication Server makes client connections to the RSSD with an encrypted password. Values are on and off (the default).
stats_flush_rssd	Parameter not supported in this version.
stats_reset_afterflush	Parameter not supported in this version.

Database connection parameters

Table 3-3 describes the new database connection configuration parameters.

Table 3-3: New configuration parameters

Parameter	Description
dsi_ignore_commit_order	Removes restrictions when a transaction is submitted or committed to the replicate database.

Parameter	Description
dsi_isolation_level	Specifies the isolation level for transactions. ANSI standard and Adaptive Server supported values are: • 0 – ensures that data written by one transaction represents the actual data. • 1 – prevents dirty reads and ensures that data written by one transaction represents the actual data. • 2 – prevents nonrepeatable reads and dirty reads, and ensures that data written by one transaction represents the actual data. • 3 – prevents phantom rows, nonrepeatable reads, and dirty reads, and ensures that data written by one transaction represents the actual data. **Note** Replication Server support is not limited to ANSI standard values only.
dsi_replication_ddl	Specifies whether or not to replicate the DDL transactions in the transaction log.
dsi_rs_ticket_report	Specifies whether or not to invoke function string rs_ticket_report.
stats_show_zero_counters	Specifies whether or not to display counters with zero observations since the last reset.

Table 3-4 describes changed database configuration parameter.

Table 3-4: Changed database configuration parameter

Parameter	Description of enhancements
dsi_serialization_method	Option isolation_level_3 is not supported in this version.

RCL commands

Table 3-5 summarizes new RCL commands.

Table 3-5: New RCL commands

Command	Description
admin stats	Displays information and statistics about Replication Server counters. Replaces admin statistics.
admin stats, backlog	Displays the current backlog in the Replication Server stable queues.
admin stats, reset	Resets all counters that can be reset. Replaces admin statistics, reset.
admin stats, status	Displays the flushing status for all counters.
admin stats, {tps | cps | bps}	Displays throughput in terms of transactions per second, counts per second, or bytes per second.
alter partition	Changes the size of a partition.
create partition	Makes a partition available to Replication Server. Replaces add partition.

Command	Description
sysadmin erssd	Displays current values for the ERSSD backup time, backup interval, and backup location.

Table 3-6 describes changed RCL commands.

Table 3-6: Changed RCL commands

Command	Description of changes
admin stats_config_connection	Command not supported in this version.
admin stats_config_module	Command not supported in this version.
admin stats_config_route	Command not supported in this version.
admin stats, flush status	Command not supported in this version.
admin statistics, flush_status	Command not supported in this version.
admin stats_intrusive_counter	Command not supported in this version.
alter connection	• New options: dsi_isolation_level – specifies the isolation level for transactions dist_stop_unsupported_cmd – used for mixed-version environments. • Changed option: dsi_serialization method – isolation_level_3 option deleted.
alter logical connection	Enhanced to include a new parameter, dist_stop_unsupported_cmd, for mixed-version environments.
configure replication server	

System tables

Table 3-7 describes system table changes.

Table 3-7: Altered system tables

System table	Description
rs_articles	New minvers column describes minimum Replication Server version required to support this article.
rs_funcstrings	New minvers column describes the minimum Replication Server version required to support this function.
rs_publications	New minvers column describes minimum Replication Server version required to support this publication.

System table	Description
rs_statdetail	Provides new options: • counter_obs – specifies observed counter value. • counter_total – specifies the total counter value for the run or observation period. • counter_last – specifies the last observed counter value for the run or observation period. • counter_max – specifies the maximum counter value for the run or observation period.

System stored procedures

Table 3-8 describes new system stored procedures.

Table 3-8: New system stored procedures

System procedure	Description
rs_dump_stats	Extracts Replication Server statistics collected in the RSSD by admin stats to a comma-delimited format.
rs_ticket	Is executed at the primary database. rs_ticket sends a ticket to the primary Replication Server through rs_marker. If the replicate Replication Server is configured to process rs_ticket, Replication Server calls rs_ticket_report at the replicate database.
rs_ticket_report	Processes a ticket sent by the Replication Server. If at the Replication Server dsi_rs_ticket_report is on and a rs_ticket is received from the primary database, the Replication Server executes rs_ticket_report stored procedure at the replicate database.

Function strings

Table 3-9 describes new function strings.

Table 3-9: New function strings

Function string	Description
rs_batch_end	Specifies the SQL statements required to mark the end of a batch of commands. This function string is used with rs_batch_start.
rs_batch_start	Specifies the SQL statements required in addition to the rs_begin statements to mark the beginning of a batch of commands.
rs_repl_on	Sets replication on in Adaptive Server for a database connection.

Function string	Description
rs_set_ciphertext	Enables set cipher string on, which enables replication of encrypted columns for database connections.
rs_set_dml_on_computed	Enables the replication of materialized computed columns.
rs_set_isolation_level	Passes the isolation level for transactions to Adaptive Server.
rs_ticket_report	Invokes the replicate database stored procedure rs_ticket_report.

Table 3-10 describes changed function strings.

Table 3-10: Changed function strings

Function string	Description of changes
rs_set_isolation_level_3	Function string not supported in this version.
rs_truncate	Includes support for the truncate table partition command.
rs_writetext	Includes support to modify unitext data in a replicate database.

Datatypes

Table 3-11 describes new datatypes.

Table 3-11: New datatypes

Datatype	Description and range
bigint	Fixed-width, signed, 8-byte datatype. Holds whole numbers between –263 and +(263 – 1), inclusive.
unsigned bigint	Unsigned integer datatype. Holds whole numbers between 0 and 18,446,744, 073, 709,551,615, inclusive.
unsigned int	Unsigned integer datatype. Holds whole numbers between 0 and 4,294,967,295, inclusive.
unsigned smallint	Unsigned integer datatype. Holds whole numbers between 0 and 65535, inclusive.
unsigned tinyint	Unsigned integer datatype. Holds whole numbers between 0 and 255, inclusive.
unitext	Variable-width, nullable Unicode datatype coded in UTF-16. Used as text datatype.

CHAPTER 4

New Features in Replication Manager 15.0

This chapter describes new features in the 15.0 version of the Replication Manager, a plug-in to Sybase Central.

Topic	Page
New user interface features	35
Replication Manager features supported	40

New user interface features

The Replication Manager supports several new interface features that promote ease-of-use and productivity.

Two-tier management solution

Replication Manager can manage replication environments by connecting directly to servers in the environment without communicating through a management server layer. This two-tier management solution lets you manage small, simple replication environments with fewer than ten servers.

To set up a replication environment in Sybase Central, you select the servers contained in the environment from a drop-down list, which is derived from the local interfaces file, then provide a user name and password that the Replication Manager uses to connect to these servers. These passwords are encrypted when stored in the Sybase Central repository.

Three-tier management solution

If you are managing a large or complex replication environment, you might want to install the Replication Monitoring Services (RMS) server. In a three-tier management solution, RMS server is a middle-management layer that monitors the status of the servers and other components in the replication environment. Replication Manager connects to the servers in the environment through RMS.

The Replication Manager provides the client interface that displays the information provided by RMS.

For more information about RMS and its functionality, see Chapter 5, “Introducing Replication Monitoring Services.”

Replication Manager plug-in replaces Replication Server plug-in

With Replication Server 15.0, the Replication Manager plug-in replaces the Replication Server plug-in as the complete management tool for developing, managing, and monitoring a Sybase Replication Server environment for a two-tier management solution.

In earlier versions, the Sybase Central Replication Server Manager included a Replication Server plug-in and a Replication Server Manager Server (RSM Server) as the software tools to monitor, analyze, troubleshoot, and administer a replication system.

The Replication Server plug-in included a graphical user interface (GUI) integrated with Sybase Central. The Replication Manager plug-in has an interface that is similar to the Replication Server plug-in, but that runs within the Sybase Central, Java Edition framework.

In addition, Replication Manager does not require the RSM Server to manage servers in a replication environment. It communicates directly with the Replication Server and Adaptive Servers in a two-tier management solution.

Online help

Online help for Replication Manager is now available. The online help contains extensive topic-level help that provides a quick reference for all Replication Manager concepts and tasks.

Select Help from the Sybase Central main menu and then select Replication Manager online help.

Visual monitoring of status

The state of each object displays on the object icon, in the parent object Details list, and on the Properties dialog box for that object. You can monitor the status of servers, connections, routes, and queues.

Using the Details list

When you select an object in the left pane that contains subcomponents or function components, one or more tabs display in the right pane of the Sybase Central window with lists of information. For most objects, a single tab called "Details" displays, which contains a list of general information about the object.

The Details list displays:

- Subcomponents, which are other replication or database objects that are contained in another object.
- Function components, which are components that invoke a wizard when double-clicked. For example, the Add Connection object is a function component that invokes the Add Connection wizard.

Event Log pane

Replication Manager displays an event log in a pane at the bottom of the Sybase Central window. The event log displays:

- Component state changes for connections, routes, and queues
- Server availability changes
- Background thread completion
- RMS event trigger execution

Note For more information about RMS and event triggers, see Chapter 5, "Introducing Replication Monitoring Services."

To display or hide the event log, select Event Log from the View menu.

Background processing

Several tasks performed by Replication Manager can be very time-consuming, such as creating a subscription that also materializes the table. These tasks are now performed in the background, allowing Sybase Central to continue to function. When you start a time-consuming task, Replication Manager displays a message window indicating that a process is running. You can click Stop Process in this window to stop the process.

Note The process continues even if the Background Process window is closed.

When a background task completes, the Replication Manager places an event entry in the event log.

Using the Background Processes dialog box

To see the status of a background process, you can open the Background Processes dialog box, which displays a list of all of the currently running background threads.

To access the Background Processes dialog box, select Search | Background Processes. The Background Processes dialog box opens, displaying the following:

- Process – the name of the process.
- Start time – the start time of the process.
- Status – the status of the process.

Replication Manager logging enhancement

Replication Manager now uses the Sybase Central message logging feature to provide a log of all commands sent by the Replication Manager to any server. Because the log may contain passwords (for example, passwords are needed to create a connection, and these passwords are saved in the log), the Replication Manager provides the ability to turn off command logging.

❖ **Turning command logging on and off**

1 Right-click the Replication Manager object (at the top of the tree).

2 Select Properties.

3 In the Replication Manager properties dialog box, select Write SQL Commands to Log.

To turn logging off, return to the Replication Manager properties dialog box and unselect the Write SQL Commands to Log check box.

Warning! Turning on command logging can fill up the Sybase Central log, causing it to crash. If you turn command logging on, monitor the log closely.

Script editors

Replication Manager provides two script editors, the Replication Command Language (RCL) script editor and the Structured Query Language (SQL) script editor. These editors operate the same way, except the RCL script editor highlights RCL keywords while the SQL script editor highlights SQL keywords.

You can use the script editor to view generated RCL commands, which include syntax to create connection and configuration parameters that can be used to create connections outside the Replication Manager.

With the script editors, you can:

- Select several Replication Server objects and generate RCL for all objects selected.
- Edit and save the generated RCL script.
- Load an RCL script from a file and add it to the current script.
- Resubmit an RCL script to Replication Server.

❖ **Accessing the script editor**

1 Select the Replication Server object for which you want to generate RCL.

2 Right-click that object.

3 Select Generate RCL from the context menu. The RCL Script Editor opens and contains the commands needed to create the object.

Replication Manager features supported

This section describes the Replication Server 15.0 features that the Replication Manager also supports.

Support for new datatypes

Replication Manager supports the following new datatypes supported by Replication Server 15.0:

- bigint
- unsigned integer datatypes
  - unsigned tinyint
  - unsigned smallint
  - unsigned int
  - unsigned bigint
- unitext

For more details of the datatypes supported by Replication Server 15.0, see Chapter 2, “New Features in Replication Server Version 15.0.”

Support for DirectConnect

In version 15.0, the Replication Manager manages a component that represents a Sybase DirectConnect™ data access server. The DirectConnect server acts as an Open Server gateway by converting the Open Client/Server protocol used by Replication Server to the native communication protocol used by the non-Sybase replicate database.

In Sybase Central, the DirectConnect data access server is managed as any other object in the replication environment. The Replication Manager displays the state on the icon of the DirectConnect data access server and on the parent object’s Detail list. The state of the server reflects whether the server is available and the state of the back-end data servers.

Limitations

In version 15.0, the Replication Manager neither displays the error log nor sets configuration parameters for DirectConnect.

Replication support

Replication Server provides replication at the database, table, and stored procedure levels. Replication Manager allows you to create a replication definition for a database, a table, or a stored procedure. Note that a replication definition for a stored procedure is called a function replication definition. You can create, alter, and delete function replication definitions and function subscriptions.

For more information about creating, altering, and deleting replication definitions and subscriptions, see the Replication Manager online help.

Routes

A route is a one-way message stream from a source Replication Server to a destination Replication Server. Routes carry replication data. You can use Replication Manager to manage routes.

For more information on managing routes using Replication Manager, see the Replication Manager online help.

Upgrading routes

With Replication Server 15.0, you can upgrade routes using the Replication Manager. In earlier versions, the RSM Server provided the route upgrade capability.

The Replication Manager displays a set of all routes in the Replication Server that are eligible to be upgraded. You can select one single route at a time for upgrade. Upon selection of a route, the Replication Manager runs the upgrade process in the background and reports any errors or warnings in the event log.

The new route upgrade implementation provides the following additional functionalities:

- Canceling a route upgrade process – the ability to cancel a route upgrade process. However, canceling can leave the replicate Replication Server in an unstable state.
- Recovering from an unsuccessful upgrade – a recovery procedure to reset the replicate RSSD. An unsuccessful upgrade can result from either an abnormal termination of the upgrade process or a cancellation of the upgrade process.

❖ **Upgrading a route using Replication Manager**

1 In the object tree, select a Replication Server object.

2 Right-click the Replication Server object and select Upgrade Route. A dialog box lists all the routes that can be upgraded.

3 From the list, select a route and click the Upgrade button.

A background process starts, which on completion indicates whether the upgrade is successful or not.

Limitations

The route upgrade implementation has these limitations:

- You cannot currently upgrade multiple routes at the same time.
- Route upgrade is supported only in an environment domain using the two-tier solution.

Troubleshooting tools

Occasionally, an environment stops replicating data. This can happen when a transaction is not formatted correctly, or when a server generates an error. To troubleshoot the situation, you can view the Replication Server queue data and exceptions log.

Accessing the exception log

Use the exceptions log to troubleshoot a problem with replication.

❖ **Using the exceptions log**

When a transaction is not processed because of a SQL error, you can:

1 Issue a resume command with a skip transaction clause to the queue.

2 Right-click the connection and select View Exceptions.

The View Exceptions dialog box displays the exceptions log in the Exceptions Log table.

3 View the skipped transaction and the erroneous SQL in the exceptions log.

4 Filter the display of the exceptions log by selecting the columns in the table: Origin Data Server, Origin Database, and Previous Hours.

5 To see the command language associated with a transaction, select it from the Exceptions Log table. The Command Editor text box displays the SQL commands in the transaction.

6 Edit the transaction in the text box.

7 To resubmit the transaction to the replicate data server, select Resubmit.

Accessing queue data

Data that is passed between servers (Adaptive Server, Replication Server, and so on) is stored in stable queues within Replication Server. The Replication Manager displays the statistics of queue usage and displays the content of the queues.

Using the View Queue Data dialog box

The View Queue Data dialog box lets you filter and sort the data from a queue as an aid in troubleshooting transactions in the queue. You can also edit, delete, or undelete a given command, or purge the first transaction in the queue.

The View Queue Data dialog box contains the following fields:

- Filter fields, which let you select the type of filters that the Replication Manager uses to display data from the queue. These filters include:
  - Column
  - Column value
  - Segment
  - Starting block
  - Number of blocks displayed
  - Number of rows displayed
  - Whether to start at the first active segment or not
  - Whether to include all data to the end of the segment
  - Whether to include all rows or not
  - Whether you want to show deleted data
  - Whether to view all data to the end of the queue
- General buttons, which let you:
  - Display the queue data with the current filters
  - Close the dialog box
  - Purge the first transaction from the queue

  - Edit transactions
  - Delete transactions
  - Undelete transactions
  - Group transactions, which returns the Queue Data scrolling list display back to grouped transactions
- Queue Data scrolling list, which contains rows of data from the current queue. Each column contains specific information about the command and transaction contained in each row. For example, to sort the queue data by a specific column, select that column name. The Queue Data scrolling list refreshes, sorting the data according to that column. An arrow displays next to the column name to show that you have sorted the data by that column. The columns you can sort by include:
  - Segment
  - Transaction Name
  - Command
  - Origin Site
  - Origin Commit Time
  - Origin User
  - Transaction ID
  - Origin QID

Note You can only delete, undelete, or purge queue transactions when Replication Server is in standalone mode. For more information, see Chapter 3, "Managing Replication Server with Sybase Central," in the *Replication Server Administration Guide Volume 1*.

❖ **Viewing queue data**

1 Right-click the queue whose data you want to view.

2 Select View Data. The View Queue Data dialog box opens.

3 To filter data shown, select one of the filter fields.

For more information, see "Using the View Queue Data dialog box" on page 43.

4 To sort the data, select segment, transaction, origin, size, status, commit time, or user.

Connection status hide options

You can hide (or filter out) the status of connections if you do not want to see the connection status either on the individual connection icon or as part of the rollup status for Replication Server.

Because the filtering state of the connection status is stored locally by the Replication Manager, different instances of the Replication Manager do not share filtering states. For example, if you create a connection using one instance of the Replication Manager, and then set the Replication Agent status to "hide" for that connection, another Sybase Central plug-in instance monitoring the same environment does not filter the connection status because the filtering information is available only to the original Replication Manager instance.

In addition, any connection created outside of Sybase Central (by rs_init or from the command line) is not filtered automatically by the Replication Manager. You must set the filtering manually from within Sybase Central.

Filtering connection status in warm standby environments

If you are creating a warm standby environment, the Replication Manager automatically sets the filtering state for the active Data Server Interface (DSI) thread and standby RepAgent thread connections. You must set filtering for the physical connection manually by selecting one of the connection status hide options from the context menu.

Using connection status hide options

The options for hiding connection status are as follows:

- Hide the State of the Replication Agent – hides the state of the Replication Agent thread in the Details list, on the Connection Properties dialog box, and in the rollup status for the Replication Server to which that Replication Agent thread is connected.
- Hide the State of the DSI Thread – hides the state of the DSI thread in the Details list, on the Connection Properties dialog box, and in the rollup status for the Replication Server to which the DSI thread is associated.

❖ **Hiding connection status**

1 Right-click the connection whose status you want to hide.

2 Select Hide Connection Status from the drop-down menu.

A dialog box shows options for hiding the connection status.

3 Select an option.

The state for that connection now reads “Hidden.” The state on the Connection Properties dialog box and in the rollup status for the Replication Server is also hidden. The Event Log records this change.

Warm standby wizards

To create a warm standby environment, you must create the following components in order:

- A logical connection
- A connection to the active database
- A connection to the standby database

Previously, these steps were only part of the Configure Replication Environment wizard, which enables you to build a warm standby environment in one step. With this version of Replication Manager, you can use a separate wizard for each step in the process, which lets you drop and re-create connections as needed.

The three wizards are:

- Add Logical Connection – add a logical connection whether you are basing this connection on an existing physical connection or not.
- Add Active Database – add an active database connection to an existing logical connection.
- Add Standby Database – add a standby database to an existing logical connection that already has an active database connection.

Using the Add Logical Connection wizard

Using the Add Logical Connection wizard, you can add a logical connection whether you are basing this connection on an existing physical connection or not.

❖ **Creating a logical connection**

1 In the tree view, select the Logical Connection folder under the Replication Server object.

2 In the Details view, double-click Add Logical Connection. The Add Logical Connection wizard starts.

3 On the Convert Existing Connection wizard page, verify that the Use an Existing Connection as the Active Connection check box is not selected.

4 Enter the database name and data server name for the logical connection.

5 Review the summary information for the logical connection.

6 If everything looks correct, click Finish. Otherwise, click Back to return to an earlier page in the wizard and change the logical connection information. Then return to the final wizard page and click Finish.

Replication Manager creates the logical connection object.

Note You must create the connections to the active and standby databases before you have a working warm standby environment. For more information, see "Using the Add Active Database wizard" on page 47 and "Using the Add Standby Database wizard" on page 48.

Using the Add Active Database wizard

Using the Add Active Database wizard, you can add an active database connection to an existing logical connection.

❖ **Creating an active database connection**

1 In the Details view, right-click the logical connection object you created using the Add Logical Connection wizard.

2 Select Add Active Database. The Add Active Database wizard starts.

3 Select the active server and active database.

4 Select the Replication Server that will manage the database connections.

5 Enter the user name and password of the maintenance user.

6 Select the user name and password that the RepAgent will use to connect to Replication Server.

If the RepAgent user does not exist, the wizard creates one for you and gives it a default name and password. Accept the defaults, or enter your own values.

7 Select the Enable all objects in the active database check box to enable database objects.

8 Review the summary information about the replication environment.

9 If everything looks correct, click Finish. Otherwise, click Back to return to an earlier page in the wizard and change the replication environment information. Then return to the final wizard page and click Finish.

Replication Manager creates the active database connection.

Using the Add Standby Database wizard

Using the Add Standby Database wizard, you can add a standby database to an existing logical connection which already has an active database connection.

❖ **Creating a standby database connection**

1 In the Details view, right-click the logical connection object you created using the Add Logical Connection wizard.

2 Select Add Standby Database. The Add Standby Database wizard starts.

3 Select the standby server and standby database.

4 Select the Replication Server that will manage the database connections.

5 Enter the user name and password of the maintenance user.

6 Select the user name and password that the RepAgent will use to connect to Replication Server.

If the RepAgent user does not exist, the wizard creates one for you and gives it a default name and password. Accept the defaults, or enter your own values.

7 Select the materialization method.

8 Select one of the following options:

- Initialize Standby Database with Dump Load, which initializes the standby database with the current data dump from the active database.
- Use Dump Marker in Transaction Log, which replicates transactions that are executed between the time the active database is enabled and the time the data is dumped.

9 Review the summary information about the replication environment.

10 If everything looks correct, click Finish. Otherwise, click Back to return to an earlier page in the wizard and change the replication environment information. Then return to the final wizard page and click Finish.

Replication Manager creates the standby database connection.

Note Before you use your warm standby environment, resume the DSI thread on the replicate database connection, if necessary.

Thread management

Replication Manager displays the state of the threads in the Replication Server. If the thread is directly related to another Replication Server component, such as a connection, route, or queue, then the thread component is represented in Sybase Central by the related component and its features. For example, the Replication Agent thread (RepAgent) or DSI threads are represented by the associated connection component. The DSI EXEC thread, however, is not represented by another component.

Using thread context menus

The Thread context menus give you access to the menus of a related component, if a related component exists for a particular thread. For example, if you right-click a DIST thread, the context menu displays Connection and Copy. If you select Connection, a context submenu displays, showing you the commands from the connection context menu.

Note Some thread objects are not related to a component and, therefore, do not give access to a submenu. For example, if you right-click a DSI EXEC thread, the context menu displays only the Copy command.

The following table shows the mapping of thread objects to related components:

Thread	Related component
DSI	Connection
DIST	
REP AGENT	

Thread	Related component
SQM	Queue
SQT	
RSI	Route
USER	User

Viewing the Details list and other thread information

When you select the Threads folder in the left pane, the Details list and tab display, as well as several other tabs. These other tabs give you access to additional information about each type of thread, as follows:

- DIST – displays information about distributor threads.
- DSI – displays information about DSI threads.
- RSI – displays information about RSI threads.
- SQM – displays information about SQM threads.
- SQT – displays information about SQT threads.

See the *Replication Server Administration Guide Volume 1* for more information about Replication Server threads.

Using thread information

You can copy information about a thread from any of the thread information lists to the system clipboard. The columns of data in each list are separated by tabs so that you can paste the columns directly into a columnar format such as a spreadsheet.

❖ **Copying thread information**

1 Right-click the thread in the right pane.

2 Select Copy.

3 Go to the destination file and paste the information as you would any text.

For example, if you select the DIST thread number 19 for a Replication Server called “myRepServer,” then select Copy from the menu and paste that information into a file, the results look similar to the following:

```
19 Awaiting Wakeup 102 myRepServer.emb2 102 P Normal 0 1 0 21787 43856 0 0 0 14
```

CHAPTER 5

Introducing Replication Monitoring Services

This chapter introduces a new Replication Server 15.0 component called the Replication Monitoring Services (RMS). RMS is the new middle-management monitoring layer that replaces the existing Replication Server Manager and provides monitoring services for large and complex replication environments.

Name	Page
Introducing Replication Monitoring Services	51
Monitoring a replication environment using RMS	55

Introducing Replication Monitoring Services

Replication Monitoring Services (RMS) replaces the functionality of the existing Replication Server Manager Server (RSM Server). RMS monitors the servers and components in a replication environment, provides the ability to control the flow of data in the replication environment, and sets the configuration parameters.

RMS functionality is available through the Replication Manager plug-in to Sybase Central and the command line API. Replication Manager provides commands to create, modify, or delete replication objects, while RMS provides an API to monitor, manage, and configure the replication environment.

RMS is applicable only for a three-tier management solution.

Three-tier management solution

A three-tier management solution is for large and complex replication environments consisting of ten or more Replication Servers. The Replication Manager connects to the servers in the environment through RMS. RMS provides the monitoring capabilities for the replication environment.

RMS monitors the status of servers and other components in the replication environment, and the Replication Manager plug-in provides the client interface that displays the status information provided by the RMS.

Table 5-1 lists the features supported in the Replication Server Manager, a component in the earlier versions, and in the Replication Monitoring Services, the new component in version 15.0.

Table 5-1: Difference in features supported in RSM and RMS

Feature	Replication Server Manager	Replication Monitoring Services
Allows you to manage, monitor, and configure replication system components	X	X
Monitors the availability of servers and the state of all connections and routes	X	X
Manages a warm standby environment	X	X
Supports multi-site availability		X
Supports the Embedded RSSD	X	X
Provides support for SSL and network-based security (DCE/Kerberos) security between servers in a replication environment	X	
Runs on Microsoft Windows platforms as well as all UNIX platforms supported by Replication Server		X
Provides a server-centric view of the replication environment	X	X
Enables administration of a logical group of servers		X
Sets configuration parameters of Replication Servers, Replication Agents, the Adaptive Server Enterprise Replication Agent thread (RepAgent), connections, and routes	X	X
Monitors the latency and the state and performance of a replication path	X	X
Provides commands to create, alter, and delete replication objects	X	
Provides logging and tracing of server commands	X	X
Generates a rollup status for Replication Servers and Adaptive Servers	X	X
Executes user-defined scripts for events in the replication environment	X	X

Monitoring servers in the replication environment

Using RMS, you can monitor the following servers in your replication environment:

- Adaptive Server Enterprise
- Adaptive Server Anywhere and IQ

- Replication Agent
- Mirror Replication Agent
- DirectConnect
- Open Server
- Replication Server
- Remote RMS servers

Software requirements and compatibilities

Replication Monitoring Services requires the following software:

- JRE version 1.4.2.03
- Sybase Unified Agent Framework (UAF) version 1.5.0.97
- jConnect™ for JDBC™ version 6.0

Installation

You can use InstallShield to install the Replication Monitoring Services component at the same time you install Replication Server.

If you specify a Typical or Full installation when you install Replication Server 15.0, InstallShield automatically installs the Replication Monitoring Services component along with Sybase Central and Replication Manager.

Starting and stopping RMS

RMS is a Java application built using the Sybase Unified Agent Framework (UAF). To start or stop RMS you must start or stop UAF.

❖ **Starting RMS**

1 Navigate to the Replication Server installation directory, *%SYBASE%* on Windows, or *$SYBASE* on UNIX.

2 Set the environment variables by executing SYBASE.bat (Windows) or by sourcing SYBASE.csh (UNIX).

3 Change to the *%SYBASE_UA%\bin* directory (Windows) or the *$SYBASE_UA/bin* directory (UNIX).

4 In Windows, execute agent.bat or add this command to run Unified Agent (UA) in the background. Similarly on UNIX, execute agent or add this command to run UA in the background.

❖ **Stopping RMS**

1 Navigate to the Replication Server installation directory, *%SYBASE%* on Windows, or *$SYBASE* on UNIX.

2 Set the environment variables by executing SYBASE.bat (Windows) or by sourcing SYBASE.csh (UNIX).

3 Stop RMS:

- If an Adaptive Server Unified Agent is running, enter one of these commands:

  ```
  shutdown [-port rmi-port] [-U username]
  [-P password]
  ```

  Or

  ```
  shutdown [-port rmi-port] [-user username]
  [-password password]
  ```

  where:

  - *rmi-port* value is 9999.
  - *username* and *password* are the values configured for the UA.
- If RMS is the only Unified Agent running, enter shutdown.bat (Windows) or shutdown (UNIX).

Connecting to RMS in Sybase Central

To connect to RMS in Sybase Central:

1 Select the Connect icon from the toolbar. The Connect to a Replication Domain window opens.

2 Select the RMS Server radio button.

3 Enter the user name and password required to connect to RMS.

4 Select RMS from the list of servers in the drop-down list or click the Options button to provide the connection information for the RMS.

5 Enter a server name, host, and port number.

6 Click OK.

The RMS server is added to the object tree in Sybase Central.

Monitoring a replication environment using RMS

To monitor a replication environment, you must set up a RMS domain. This domain consists of the servers in your environment you want to monitor using the RMS. In a three-tier management solution, the Replication Manager connects to the RMS, which then connects to the various servers in your replication environment.

To set up the RMS domain and monitor servers in this domain, you can either:

- Use the Replication Manager plug-in graphical interface, or
- Use command line API

In a three-tier replication environment, you can perform some monitoring through Replication Manager, and additional monitoring through the RMS API at the command line.

The following sections provide information on how to monitor RMS using the Replication Manager plug-in and the API. For those tasks for which you can use both interfaces, only the Replication Manager plug-in is described in detail, while the name of API is identified. For those tasks that can be performed using only the API, detailed API information is given.

Adding and dropping servers for monitoring in Sybase Central

Add and drop servers in a three-tier environment in the same way as in a two-tier environment.

Adding a server

The servers that you add for monitoring in a three-tier solution can be from anywhere in your network. This allows you to monitor replication systems that are distributed worldwide.

You can use Replication Manager wizard to add a server to a RMS domain. Depending on the type of server you are adding (Adaptive Server, Replication Server, RepAgent, Open Server™, Sybase Replication Agents, DirectConnect), the wizard prompts for different information.

Before you add a server, have this information available:

- The server's name.
- The type of server you are adding.
- The user name and password used to administer the server. The login must have System Administrator privileges on the server.
- *Replication Server only* – the user name and password of the RSSD primary user.

To add servers to the RMS server folder in Sybase Central:

1 Select the RMS.

2 Select File | New | Servers, or double-click the Add Servers icon in the right pane of the main window.

3 Enter the required information in the wizard dialog boxes.

The servers are added to the Sybase Central viewer under the RMS server in the object tree.

API command to add server to monitor

Use add *server* to add a server to be monitored by RMS. For detailed information about this API, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Dropping a server

To drop a server from the RMS domain:

1 Select the server you want to drop.

2 Do one of the following:

- Click the Delete icon from the toolbar.

- Right-click the selected server and select Delete from the context menu.

Note Although Sybase Central removes the server from the RMS server folder, the server is not actually removed from your replication system. Therefore, the server name may still appear in the dialog boxes because there are routes or database connections associated with it.

API command to drop server from being monitored

Use drop *server* to drop a server that is being monitored by RMS. For detailed information about this API, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Viewing monitored objects in Sybase Central

Viewing monitored objects in Replication Manager is the same as viewing objects in a two-tier environment.

In the object tree, double-click or expand the RMS icon to view the replication objects managed by RMS. Under RMS, you can view the monitored servers and its components such as connections, routes, queues, and threads. When you select a particular replication object such as the Routes folder, you can view the list of created routes. You can manage these replication objects using Replication Manager.

Setting configuration parameters for monitored replication objects

You can set the configuration parameters for the following replication objects monitored by RMS:

- Replication Server
- Replication Agent thread
- Remote RMS server
- Database connections and logical connections
- Routes

To set the configuration properties of an object, in Replication Manager

1 Select the object and choose File | Properties. The property sheet for the object opens.

2 Select the parameters tab in the Properties dialog box. The Parameters page displays a list of all configuration parameters for the selected server or component.

3 Select the parameter you want to modify and click the Edit Parameter button.

4 In the Edit dialog, enter the new parameter and click OK.

5 You can modify any other parameter in the list. When you are finished, click OK in the Properties dialog box.

API commands to set configuration parameters

RMS API commands to set configuration parameters for the replication objects are:

- configure server – returns configuration parameter information for a Replication Server or Replication Agent, or sets the value of a specified configuration parameter. configure server also retrieves and sets RMS-specific parameters.
- configure component – returns configuration parameters for a component or sets the value of the configuration parameter specified.
- configure *RMS* – returns the configuration parameter information for the RMS, or sets the value of a specified RMS configuration parameter

For detailed information about these APIs, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Monitoring a logical group of servers

RMS enables you to define a set of servers as a logical group and monitor the group as a single entity. You can have many logical groups in your replication environment. The servers in a group can belong to different logical groups. However, all servers in one logical group must be of the same type. For example, you can have one logical group of all Replication Servers and another for all Adaptive Servers. When you issue any command to a group, it affects all the servers contained in this logical group.

You can delete logical groups. When you delete a logical group, only the group is deleted, not the servers within the group.

RMS also returns a rollup status for each group or each server in a group. Roll-up status shows the lowest status reported, for example, if any server in a group is not UP, then the group status is reported as SUSPECT.

API commands to create/delete/list logical groups

RMS API commands to manage logical groups are:

- create group – defines a logical group of servers.
- delete group – deletes a logical group that was added using the create group command.
- get group – returns a result set that contains either a list of the groups and a rollup status for each group or for each server.

For detailed information about these APIs, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Suspending or resuming components in the replication environment

You can suspend or resume the following monitored components in the replication environment:

- RepAgent in Adaptive Server
- Replication Agent and DSI connection in Replication Server
- Routes
- Queues (resume only)
- Replication in a Replication Agent

In Sybase Central, select the monitored connection or route from the RMS Server folder in the object tree:

- To suspend a monitored connection or route, select the connection or route object from the RMS Server folder and select Suspend from the context menu.
- To resume a monitored connection or route, select the connection or route object from the RMS Server folder and select Resume from the context menu.

API commands to suspend/resume replication components

RMS API commands to suspend or resume components in the replication environment are:

- suspend *component* – suspends a component in a specified server.
- suspend *replication* – suspends replication in a Replication Agent.
- resume *component* – resumes a component in a specified server.
- resume *replication* – resumes replication in a Replication Agent.

For detailed information about these APIs, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Shutting down monitored servers

RMS enables you to shut down any monitored Replication Server, Replication Agent, or Mirror Replication Agents.

In Sybase Central, select the monitored server or RMS in the object tree and select Shutdown from the context menu. To shut down a group of servers, select the logical group folder and select Shutdown All from the context menu.

API command to shut down Replication Server, group, or RMS

Use shutdown *server* to shut down one or more, or group of monitored Replication Servers or RMS. For detailed information about this API, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Generating rollup status for servers

The RMS status shows the lowest status reported, for example, if the status of any server in the list is not UP, then the status for the RMS is reported as SUSPECT.

For example, the status of the Adaptive Server reflects the status of its Replication Agent threads. If any of the monitored Replication Agent threads in the Adaptive Server are not UP, then the state of the Adaptive Server is set to SUSPECT.

For more information about RMS server and component states, see Appendix C, "RMS Server and Component States," in the *Replication Server Reference Manual*.

API commands to view rollup status information

RMS API commands to view rollup status information and the various server states are:

- get servers – returns a list of servers that are monitored by the RMS, and the status of the RMS environment.
- get status description – retrieves the list of status descriptions for a server or component.

For detailed information about these APIs, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Generating latency and heartbeat information

The RMS heartbeat feature uses the stored procedure rs_ticket to generate latency information, which is the amount of time it takes for a transaction to move from the primary database to the replicate database. You can use this information to monitor Replication Server performance, module heartbeat, replication health, and table-level quiesce.

At a specified interval, RMS executes rs_ticket at the primary database. The generated latency information is stored in a table in the replicate database. You can use RMS to set up the heartbeat process to retrieve the latency information from the replicate database.

For more information about rs_ticket and its functionality, see Chapter 6, "Adaptive Server Stored Procedures," in the *Replication Server Reference Manual*.

API commands to monitor heartbeat

RMS API commands to monitor heartbeat are:

- get heartbeat – retrieves a list of the heartbeat processes that have been defined in the RMS.
- get heartbeat ticket – retrieves a set of tickets from the rms_ticket_history table, for the heartbeat process and date and time range specified.
- start heartbeat – sets up and starts a heartbeat process from a specified primary connection to a specified replicate connection.
- stop heartbeat – stops the heartbeat process between the primary and replicate databases. Optionally truncates the rms_ticket_history table.

For detailed information about these APIs, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

Adding event triggers

Replication Monitoring Services is designed to monitor the replication environment. When something happens in your environment, server and component status changes. These changes are displayed in the event log. RMS allows you to create event triggers to monitor these changes.

Event triggers notify you when some events occur in the replication environment. RMS executes the script when the specified event occurs. For example, a user can set up a script to be notified with an e-mail message when a connection suspends. You can create an event trigger for any server or component that the RMS monitors.

❖ **Creating an event trigger for a Replication Server**

1 In the object tree, select the Replication Server.

2 On the right side of the desktop, select the event log pane.

3 Double-click the Add Server Event Trigger icon.

4 Select the status change that will trigger the event.

5 As an option, enter a "Wait before executing" value. This notifies RMS to wait for the event to change before executing the trigger.

6 To execute a trigger at each monitoring interval, rather than only once, select "Execute at Each Interval."

7 Enter the name of the script for RMS to execute when the event occurs.

8 Click OK. The new event displays in the Event Log pane.

API commands to add/drop/get triggers

RMS API commands for trigger-related tasks are:

- add *event* trigger – sets up a trigger, such as a process or a script, that is executed by the RMS when a specific event occurs.
- drop *event* trigger – removes a trigger that the RMS is monitoring.
- get triggers – displays information about the triggers that are monitored by the RMS.

For detailed information about these APIs, see Chapter 9, "Replication Monitoring Services API," in the *Replication Server Reference Manual*.

CHAPTER 6

New Features in Replication Server Version 12.6

This chapter describes the new features introduced with Sybase Replication Server version 12.6 and 12.5 EBF.

Topic	Page
Multi-Site availability (MSA)	63
Support for symmetric multiprocessors (SMP)	64
The embedded RSSD (ERSSD)	66
Performance enhancements	66
date and time datatypes	69
Support for sending encrypted passwords	69
New bulk materialization method	70
Chinese character set (GB18030) support	71

Multi-Site availability (MSA)

MSA extends Replication Server replication capabilities and can make the process of setting up a replication system both faster and easier.

Some of the features that MSA provides include:

- A simple replication methodology that requires only one replication definition for the primary database and only one subscription for each subscribing database.
- A replication filtering strategy that lets you choose whether or not to replicate individual tables, transactions, functions, system stored procedures, and data definition language (DDL).
- Replication of DDL to any replicate database—including non–warm standby databases.
- Replication to multiple replicate sites—for standby as well as nonstandby databases.

You can overlay MSA scenarios onto your existing replication structure. The procedures for implementing MSA are similar to those you already use to replicate to warm standby or replicate databases.

Database replication

When you use table and function replication, you describe each piece of data that is to be replicated using individual table and function replication definitions and subscriptions. This methodology allows you to transform data and provides fine-grained control over the information being entered in the replicate database. However, you must mark each table or function to be replicated, create a replication definition for each replicated table or function, and create subscriptions for each replication definition at each replicate database.

MSA lets you identify specific database objects: tables, functions, transactions, DDL, and system stored procedures in a single replication definition. You can choose to replicate the entire database; or you can choose to replicate—or not replicate—specific tables, functions, transactions, DDL, and system stored procedures in that database. If you do not need to replicate partial tables, MSA can provide replication while affording the advantages of simple setup and maintenance.

When the replicate is a warm standby database

In the non-MSA warm standby scenario, changes to the primary database are copied directly to the warm standby database without alteration. This methodology allows replication of DDL. To change or qualify the data sent, you must add table and function replication definitions. Each primary database can have one, and only one, standby database.

MSA provides all the features of traditional Sybase warm standby. In addition, MSA:

- Enables replication to multiple standby databases
- Provides the option to replicate or not replicate specific database objects

MSA supports the use of logical connections.

Support for symmetric multiprocessors (SMP)

Replication Server 12.6 lets you run Replication Server on either symmetric multiprocessor (SMP) or single-processor platforms. Replication Server's multithreaded architecture supports both hardware configurations.

On a single processor platform, Replication Server threads run serially. On a multiprocessor platform, Replication Server threads can run in parallel, thereby improving performance and efficiency.

Replication Server support for multiple processors is based on Open Server™ support for multiple processors. Both servers use the POSIX thread library on UNIX platforms and the WIN32 thread library on Windows NT platforms. For detailed information about Open Server support for multiple processing machines, see the *Open Server Server-Library/C Reference Manual.*

When Replication Server is in single-processor mode, a server-wide mutual exclusion lock (mutex) enforces serial thread execution. Serial thread execution safeguards global data, server code, and system routines, ensuring that they remain thread-safe.

When Replication Server is in multiprocessor mode, the server-wide mutex is disengaged and individual threads use a combination of thread management techniques to ensure that global data, server code, and system routines remain secure.

Configuring SMP

To enable SMP on a multiprocessor machine, use configure replication server with the smp_enable option. For example:

```
configure replication server set smp_enable to 'on'
```

Monitoring thread status

You can verify Replication Server thread status using these commands and procedures:

- admin who – provides information on all Replication Server threads
- admin who_is_up or admin who_is_down – lists Replication Server threads that are running, or not running.
- sp_help_rep_agent – provides information on the RepAgent thread and the RepAgent User thread. This is an Adaptive Server stored procedure.

Monitoring performance

Replication Server provides monitors and counters for monitoring performance. See Chapter 15, “Using Counters to Monitor Performance,” in the *Replication Server Administration Guide Volume 2*.

See the *Replication Server Administration Guide* for more information about increasing Replication Server performance using SMP.

The embedded RSSD (ERSSD)

Replication Server can run either on an Adaptive Server Replication Server System Database (RSSD) or on an embedded RSSD (ERSSD). ERSSD is designed for users who do not want to manage the Replication Server RSSD in Adaptive Server. Replication Server is easier to install and manage with ERSSD. If you select embedded RSSD when you install Replication Server, ERSSD is automatically installed, configured, and started in the background. It is self-maintained. Backup procedures are automatic and pre-configured.

Limitations

Currently, you cannot create a route originating from Replication Server with ERSSD. Nor can you migrate between RSSD and ERSSD.

To use the ERSSD, you must select it when you install Replication Server. For more details, see the *Replication Server Installation Guide* and the *Replication Server Administration Guide*.

Performance enhancements

Replication Server 12.6 includes new performance enhancements. See Chapter 16, "Performance Tuning," in the *Replication Server Administration Guide Volume 1* for detailed information about these enhancements.

Better management of empty transactions

Transactions that contain only a begin and a commit statement can degrade the performance of warm standby connections. To enhancement performance, these transactions are now deleted from the inbound queue as they are read.

To further enhance performance, Sybase recommends that you also tune your application to eliminate as many of these empty transactions as possible.

Internal commit control for parallel processing

To resolve conflicting updates when using parallel processing, Replication Server must maintain transaction commit order and resolve commit consistency deadlocks.

Replication Server introduces a new method to maintain commit control using the function string rs_dsi_check_thread_lock. Replication Server uses rs_dsi_check_thread_lock to check whether the current DSI executor thread is blocking another replicate database process. This new method handles commit control within Replication Server thus requiring less network I/O than other methods, and may result in the rollback of only one transaction instead of many.

New database connection parameters used with internal commit control are:

- dsi_commit_check_locks_intrvl
- dsi_commit_check_locks_max
- dsi_commit_control

New Replication Server configuration parameters

You can use new Replication Server configuration parameter to fine-tune Replication Server performance.

- sqt_init_read_delay – the amount of time an SQT thread sleeps while waiting for an SQM read before checking for new instructions in its command queue.
- sqt_max_read_delay – the maximum amount of time an SQT thread sleeps while waiting for an SQM read before checking for new instructions in its command queue.

New database configuration parameters

You can use new database connection configuration parameters to fine-tune Replication Server performance.

- dsi_commit_check_locks_intrvl – specifies the number of milliseconds (ms) the DSI executor thread waits between executions of the rs_dsi_check_thread_lock function string. Used with parallel DSI.
- dsi_commit_check_locks_max – specifies the maximum number of times a DSI executor thread checks whether it is blocking other transactions in the replicate database before rolling back its transaction and retrying it. Used with parallel DSI.

- dsi_commit_control – specifies whether commit control processing is handled internally by Replication Server using internal tables (on) or externally using the rs_threads system table (off). Used with parallel DSI.

Changed database configuration parameters

The database connection parameters dsi_serialization_method and dsi_partitioning_rule have changed.

dsi_serialization_method specifies how and when parallel DSI threads can start. It includes these new or changed options:

- no_wait – specifies that a transaction can start as soon as it is ready—without regard to the state of other transactions.
- wait_for_start – specifies that a transaction can start as soon as the transaction scheduled to commit immediately before it has started. Replaces the none option. When used with the origin partitioning parameter, this option replaces the single_transaction_per_origin serialization method.
- isolation_level_3 – is the same as wait_for_start, except that DSIs will specify isolation level 3 when connecting to the replicate database.
- wait_for_commit – specifies that a transaction can start only when the previous transaction is ready to commit.
- none – maintained for backward compatibility. Replaced by wait_for_start.
- single_transaction_per_origin – maintained for backward compatibility. Replaced by dsi_serialization_method set to wait_for_start and dsi_partitioning_rule set to origin.

dsi_partitioning_rule specifies the partitioning rules the DSI uses to partition transactions among available parallel DSI threads. It includes these new options:

- origin – specifies that transactions with the same origin must be serialized when applied to the replicate database.
- origin_sessid – specifies that transactions with the same origin *and* the same process ID (SPID in Adaptive Server) must be serialized when applied to the replicate database. The LTL version must be 600 or later.

  Sybase recommends that you try setting dsi_partitioning_rule to origin_sessid, time as this setting may provide the most efficient partitioning.

date and time datatypes

There are two new datatypes, date and time, in Replication Server. These datatypes extend the existing datetime and smalldatetime datatypes, providing date and time columns to replicate and standby databases. Both are fixed-width 4-byte datatypes that support rs_subcmp, and mixed-version environments.

Replication Server version 12.6 with date and time datatype support is backward-compatible with earlier versions of Adaptive Server. However, earlier versions of Adaptive Server do not recognize date and time, and thus can send only datetime and smalldatetime data.

The new columns generated by date and time datatypes allow you to replicate date and time data to both standby and replicate databases. These columns can be part of the primary key in a replication definition, and are searchable columns in a replication definition. You can use date and time columns in the where clause of define subscription, create subscription, or create article. In the same way, the date and time columns are searchable parameters in a function replication definition, again used in the where clause of define subscription, create subscription, or create article.

Table 6-1: Range and storage needs for date/time datatypes

Datatype	Range	Storage needed
date	January 1, 0001 to December 31, 9999	4
time	12:00:00 AM to 11:59:59.999 PM	4
smalldatetime	January 1, 1900 to June 6, 2079; 12:00:00AM to 11:59:59:999 PM	4
datetime	January 1, 1753 to December 31, 9999; 12:00:00AM to 11:59:59.999 PM	8

Support for sending encrypted passwords

Replication Server 12.6 supports the -X option in isql that sends encrypted passwords through the network when making a client connection.

To ensure that all Replication Server client connections—except the first connection to the RSSD—send encrypted passwords, set the Replication Server configuration parameter send_enc_password to on. For example, enter:

```
configure replication server
 set send_enc_password to 'on'
```

To ensure that all Replication Server client connections, *including* the first connection to the RSSD, send encrypted passwords, set the configuration parameter RS_enc_pw to on in the *rs_name.cfg* file using a text editor.

If RS_enc_pw is on, all Replication Server connections to the RSSD send encrypted passwords, even if send_enc_password is off.

New bulk materialization method

Replication Server 12.6 supports a new bulk materialization method for copying or moving a database from a source Adaptive Server to a destination Adaptive Server without shutting down the source Adaptive Server. The Adaptive Server quiesce database ... to *manifest_file* and mount commands let you quiesce the server and copy or move the database.

To use this bulk materialization method, both the source and destination database servers must be Adaptive Server version 12.5.1 or later.

You can use quiesce database ... to *manifest_file* to generate all the data storage information and then use mount to mount the data to a new database—see Chapter 10, "Managing Subscriptions," in the *Replication Server Administration Guide*. You can use quiesce database ... to *manifest_file* and mount when you add a warm standby database—see Chapter 13, "Managing Warm Standby," in the *Replication Server Administration Guide Volume 1*.

For information about mount and unmount, see Chapter 22, "Database Mount and Unmount," in the *Adaptive Server Enterprise System Administration Guide*.

Chinese character set (GB18030) support

Replication Server supports all character sets supported by Adaptive Server Enterprise. Accordingly, Replication Server 12.6 supports the Chinese character set (GB18030).

CHAPTER 7

System Changes in Replication Server 12.6

This chapter describes system changes introduced with Sybase Replication Server version 12.6 and 12.5 EBF.

Topic	Page
Configuration parameters	73
RCL commands	75
System stored procedures	76
Function strings	76
System tables	77
Keywords	77

Configuration parameters

This section summarizes new and changed configuration parameters for Replication Server 12.6.

Replication Server parameters

Table 7-1 describes new Replication Server configuration parameters stored in the *rs_config* system table.

Table 7-1: New Replication Server parameters for rs_config table

Parameter	Description
smp_enable	Enables SMP on Replication Server on a multiprocessor machine.
send_enc_password	Enables encrypted password to be sent through the network for all Replication Server client connections, except the first connection to the RSSD.

Parameter	Description
sqt_init_read_delay	The length of time an SQT thread sleeps while waiting for an SQM read before checking to see if it has been given new instructions in its command queue. With each expiration, if the command queue is empty, SQT doubles its sleep time up to the value set for sqt_max_read_delay. Default: 2000 ms
sqt_max_read_delay	The maximum length of time an SQT thread sleeps while waiting for an SQM read before checking to see if it has been given new instructions in its command queue.

Table 7-2 describes new Replication Server configuration parameters stored in the *rs_name.cfg* file.

Table 7-2: New Replication Server parameters for rs_name.cfg file

Parameter	Description
RS_enc_pw	Enables encrypted password to be sent through the network for all Replication Server client connections, including the first connection to the RSSD.

Database connection parameters

Table 7-3 describes the new database connection configuration parameters.

Table 7-3: New configuration parameters

Parameter	Description
dsi_commit_check_locks_intrvl	The number of milliseconds (ms) the DSI executor thread waits between executions of the rs_dsi_check_thread_lock function string. Default: 1000ms (1 second) Minimum: 0 Maximum: 86,400,000 ms (24 hours)
dsi_commit_check_locks_max	The maximum number of times a DSI executor thread checks whether it is blocking other transactions in the replicate database before rolling back its transaction and retrying it. Used with parallel DSI. Default: 200 Minimum: 1 Maximum: 1,000,000

Parameter	Description
dsi_commit_control	Specifies whether commit control processing is handled internally by Replication Server using internal tables (on) or externally using the rs_threads system table (off). Default: off
rep_as_standby	When rep_as_standby is on, table subscriptions replicate tables marked by sp_reptostandby. For rep_as_standby on to succeed, the RepAgent parameters send maint xacts to replicate must be false and send warm standby xacts must be true. Default: off

Table 7-4 describes changed configuration parameters.

Table 7-4: Changed database configuration parameters

Parameter	Description of enhancements
dsi_partitioning_rule	Provides new options: • origin – specifies that transactions with the same origin must be serialized when applied to the replicate database. • origin_sessid – specifies that transactions with the same origin and the same process ID must be serialized when applied to the replicate database.
dsi_serialization_method	Provides new options: • no_wait – specifies that a transaction can start as soon as it is ready. Replaces the none option. • wait_for_start – specifies that a transaction can start as soon as the transaction scheduled to commit immediately before has started.

RCL commands

This section summarizes new and changed RCL commands for Replication Server 12.6.

Table 7-5 describes new commands.

Table 7-5: New Transact-SQL® commands

Command	Description
alter database replication definition	Changes a database replication definition.

Command	Description
create database replication definition	Creates a replication definition that describes a database or database objects for replication.
drop database replication definition	Deletes a database replication definition.

Table 7-6 describes changed commands.

Table 7-6: Changed commands

Command	Description of enhancements
activate subscription	Starts the distribution of updates from the primary to the replicate database and sets the database subscription status to ACTIVE.
check subscription	Checks the materialization status of a subscription to a database replication definition.
create subscription	Creates a subscription for a database replication definition.
define subscription	Defines a subscription for a database replication definition.
drop subscription	Drops a subscription for a database replication definition.
validate subscription	Sets the status of a database subscription to VALID.

System stored procedures

Table 7-7 describes new stored procedures.

Table 7-7: New system stored procedures

System procedure	Description
rs_helpdbrep	Displays information about database replication definitions associated with the current Replication Server.
rs_helpdbsub	Displays information about database subscriptions associated with the replicate Replication Server.

Function strings

Table 7-8 describes new function strings.

Table 7-8: New function strings

Function string	Description
rs_dsi_check_thread_lock	Checks whether the current DSI executor thread is blocking another replicate database process. Used during parallel processing.

System tables

Table 7-9 describes new system tables.

Table 7-9: New system tables

System table	Description
rs_dbreps	Stores all information about database replication definitions except name sets. It is replicated to all sites with a version number of 12.6 or later.
rs_dbsubsets	Stores the name sets for database replication definitions. It is replicated to all sites with a version number of 12.6 or later.

Keywords

There are three new keywords in Replication Server 12.6:

- functions
- osid
- procedures
- transactions

Index

A

activate subscription command 76
Adaptive Server Anywhere (ASA)
  ERSSD support 20
Adaptive Servers
  communication through Replication Manager 36
add *event* trigger
  enhancement 4–6
add *server*
  enhancement 8
adding event triggers in Replication Manager 61
adding servers to RMS server 55
alter database replication definition command 75
alter partition command 18

B

background processes
  running in background 38
  stopping a background process 38
  viewing in Background Processes dialog box 38
Background Processes dialog box 38
batching of commands
  for non-ASE servers 25
  using DSI connection and configuration parameters 26
  using function strings 26
bulk materialization 70

C

character set
  Chinese (GB18030) 71
check subscription command 76
Chinese character set support 71
commit control
  internal 66
computed columns
  materialized 15
  replicating 15
  virtual 15
configuration parameters
  setting with Replication Manager 57
  stats_sampling 23
connection latency trigger 3
connection status
  filtering in warm standby environments 45
  hide options 45
conventions, document style x
counter statistics
  displaying on screen 23
  reporting 23
  saving to RSSD 23
create database replication definition command 76
create partition command 18
create route command 19
create subscription command 76

D

database connection parameter
  dsi_ignore_commit_order 30
  dsi_isolation_level 31
  dsi_replication_ddl 31
  dsi_rs_ticket_report 31
  dsi_serialization_method 31
  stats_show_zero_counter 31
datatype
  bigint 12, 34
  date, time 69
  unitext 14, 34
  unsigned 13
  unsigned bigint 34
  unsigned int 34
  unsigned smallint 34
  unsigned tinyint 34

datatypes
supported by Replication Manager 40
date, database 69
dbltm 7
adding to RMS 8
displaying connection name 8
define subscription command 76
Details list 37
dialog boxes
Background Processes 38
Connection Properties 45
View Queue Data 43
drop database replication definition command 76
drop *event* trigger
enhancement 6–7
drop subscription command 76
dropping servers from RMS server 56
dsi_commit_check_locks_intrvl configuration parameter 74
dsi_commit_check_locks_max configuration parameter 74
dsi_commit_control configuration parameter 75
dsi_partitioning_rule configuration parameter 75
dsi_serialization_method configuration parameter 75

E

editors
RCL script editor 39
SQL script editor 39
Embedded Replication Server System Database
routing 19
Embedded Replication Server System Database (ERSSD) 19, 66
creating route from 19
embedded Replication Server System Database (ERSSD)
limitations 66
empty transactions 66
encrypted passwords
sending 70
ERSSD
See Embedded Replication Server System Database
Event Log pane 37
displaying or hiding 37
viewing events 37
event triggers
adding in Replication Manager 61
creating in RMS 61
event triggers, new 2–4
connection latency 3
partition size 3
queue latency 4
queue size 3
exceptions log
troubleshooting 42
viewing 42

F

filter *connection*
enhancement 9
function string
rs_batch_end 33
rs_batch_start 33
rs_dsi_check_thread_lock 67
rs_repl_on 33
rs_set_ciphertext 34
rs_set_dml_on_computed 15, 34
rs_set_isolation_level 34
rs_set_isolation_level_3 34
rs_ticket_report 34
rs_truncate 34
rs_writetext 34

G

get connections
enhancement 9
get triggers
enhancement 7

I

internal commit control 66
isolation level
kinds of 24
setting 24

J

Java JRE
  compatibility with Replication Monitoring Services 53
jConnect for JDBC
  compatibility with Replication Monitoring Services 53

K

keywords
  functions 77
  procedures 77
  transactions 77

L

license management
  with SySAM 26
list
  details 37
log
  event 37
longer identifiers 11

M

message logging
  disabling 38
  enabling 38
mixed-version
  enhanced support for 21
monitoring of status 37
monitoring replication environment
  for three-tier environment 51
  using heartbeat and latency information 61
  using Replication Manager graphical interface 55
  using RMS 55
  using RMS API 55
monitoring replication environment using RMS 55
monitoring status
  in Details list 37
  replication objects 37
  visual display 37
mount command 70
multiprocessors
  enabling 65
  monitoring 65
multi-site availability (MSA) 63
multithreaded architecture 64

N

new features
  Replication Manager 15.0 35
  Replication Server 15.0 1, 11

O

online help
  in Replication Manager plug-in 36
  invoking 37

P

parallel processing 64
  empty transactions 66
  internal commit control 66
partition size trigger 3
password encryption
  Advanced Encryption Standard algorithm 20
  FIPS-certified 20
  for maintenance user passwords 21
  for route user passwords 21
  for user passwords 21
  for user passwords in configuration file 21
POSIX thread library 65

Q

queue data
  troubleshooting 43
  viewing 43
queue latency trigger 4
queue size trigger 3

quiesce database command 70

R

RCL command
admin statistics, flush_status 32
admin stats 31
admin stats, {tps | tps | bps} 31
admin stats, backlog 31
admin stats, flush status 32
admin stats, reset 31
admin stats, status 31
admin stats_config_connection 32
admin stats_config_module 32
admin stats_config_route 32
admin stats_intrusive_counter 32
alter connection 32
alter logical connection 32
alter partition 31
configure replication server 32
create partition 31
sysadmin erssd 32
rep_as_standby configuration parameter 75
replicating
database objects 64
encrypted columns 16
partitioned tables 16
to standby databases 64
Replication Agent
connection consolidation 7–8
connection status 7–8
Replication features
supported by Replication Manager 40
Replication Manager
adding event triggers 61
new features 15.0 35
support of Replication Server/Replication features 40
using to set configuration parameters 57
Replication Monitoring Services
accessing using Replication Manager or API 55
adding servers 55
compatibilities 53
connecting to 54
description 51
dropping servers 56
in three-tier management solution 51
installing 53
monitoring logical group of servers 58
monitoring replication environment 55
setting configuration parameters 57
shutting down monitored servers 60
software requirements 53
starting and stopping 53
suspending/resuming components 59
using Unified Agent Framework 53
viewing monitored objects 57
Replication Server
communication through Replication Manager 36
configuring using Replication Manager 57
function strings 33
RCL new commands 31
system tables changes 32
Replication Server configuration parameters
errsd_backup_interval 29
errsd_backup_path 30
errsd_backup_start_time 30
erssd_backup_start_date 30
erssd_ra 30
send_enc_pw 30
stats_daemon_sleep_time 30
stats_flush_rssd 30
stats_reset_afterflush 30
Replication Server features
supported by Replication Manager 40
Replication Server System Database (RSSD) 66
RMS
See Replication Monitoring Services
routes
upgrading 41
rs_dbreps system table 77
rs_dbsubsets system table 77
rs_dsi_check_thread_lock function string 67, 77
RS_enc_pw configuration parameter 74
rs_helpdbrep system procedure 76
rs_helpdbsub system procedure 76

S

send_enc_password configuration parameter 69, 73
send_enc_pw configuration parameter 30

setting configuration parameters
using Replication Manager 57
smp_enable configuration parameter 73
software compatibilities of Replication Monitoring Services 53
software requirements for Replication Monitoring Services 53
sqt_init_read_delay configuration parameter 74
sqt_max_read_delay configuration parameter 74
starting and stopping RMS 53
status monitoring using RMS 37
stored procedure
rs_dump_stats 33
rs_ticket 33
rs_ticket_report 33
support
contacting Sybase Technical Support xii
for DirectConnect 40
symmetric multiprocessors (SMP) 64
syntax conventions xi
system tables
rs_articles 32
rs_dbrep 77
rs_dbsubsets 77
rs_funcstrings 32
rs_publications 32
rs_statdetail 33

T

threads
in Replication Server 49
using context menu 49
viewing details 50
three-tier management solution 51
time, database 69
transactions
empty 66
triggers, new 2–4
connection latency 3
partition size 3
queue latency 4
queue size 3
two-tier management solution 35, 36

U

unsigned datatypes 13
unsigned bigint 13
unsigned int 13
unsigned smallint 14
unsigned tinyint 14
user documentation, for Replication Server viii
user interface features, Replication Manager 35

V

validate subscription command 76
visual monitoring of status 37

W

warm standby
adding active database connection 47
adding logical connection 46
adding standby database connection 48
filtering connection status 45
wizards in 46
WIN32 thread library 65
wizards
Add Active Database 47
Add Logical Connection 46
Add Standby Database 48