



Installation and Configuration Guide

Risk Analytics Platform

3.0

DOCUMENT ID: DC00244-01-0300-01

LAST REVISED: September 2006

Copyright © 2005-2006 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, 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, Advantage Database Server, Afaia, Answers Anywhere, Applied Meta, Applied Metacomputing, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, ASEP, Avaki, Avaki (Arrow Design), Avaki Data Grid, AvantGo, Backup Server, BayCam, Beyond Connected, Bit-Wise, BizTracker, Certified PowerBuilder Developer, Certified SYBASE Professional, Certified SYBASE Professional Logo, ClearConnect, Client-Library, Client Services, CodeBank, Column Design, ComponentPack, Connection Manager, Convoy/DM, Copernicus, CSP, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, Dejima, Dejima Direct, Developers Workbench, DirectConnect Anywhere, DirectConnect, Distribution Director, Dynamic Mobility Model, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, eFulfillment Accelerator, EII Plus, Electronic Case Management, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise Portal (logo), Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, eremote, Everything Works Better When Everything Works Together, EWA, ExtendedAssist, Extended Systems, ExtendedView, Financial Fusion, Financial Fusion (and design), Financial Fusion Server, Formula One, Fusion Powered e-Finance, Fusion Powered Financial Destinations, Fusion Powered STP, Gateway Manager, GeoPoint, GlobalFIX, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InstaHelp, Intelligent Self-Care, InternetBuilder, iremote, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Legion, Logical Memory Manager, lLite, M2M Anywhere, Mach Desktop, Mail Anywhere Studio, Mainframe Connect, Maintenance Express, Manage Anywhere Studio, MAP, M-Business Anywhere, M-Business Channel, M-Business Network, M-Business Suite, MDI Access Server, MDI Database Gateway, media.splash, Message Anywhere Server, MetaWorks, MethodSet, mFolio, Mirror Activator, ML Query, MobiCATS, MobileQ, MySupport, Net-Gateway, Net-Library, New Era of Networks, Next Generation Learning, Next Generation Learning Studio, O DEVICE, OASIS, OASIS logo, ObjectConnect, ObjectCycle, OmniConnect, OmniQ, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Business Interchange, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, Partnerships that Work, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, Pharma Anywhere, PhysicalArchitect, Pocket PowerBuilder, PocketBuilder, Power++, Power Through Knowledge, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, Powering the New Economy, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Pylon, Pylon Anywhere, Pylon Application Server, Pylon Conduit, Pylon PIM Server, Pylon Pro, QAnywhere, Rapport, Relational Beans, RemoteWare, RepConnector, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RFID Anywhere, RW-DisplayLib, RW-Library, SAFE, SAFE/PRO, Sales Anywhere, Search Anywhere, SDF, Search Anywhere, Secure SQL Server, Secure SQL Toolset, Security Guardian, ShareLink, ShareSpool, SKILS, smart.partners, smart.parts, smart.script, SOA Anywhere Trademark, 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, Stage III Engineering, Startup.Com, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Financial Server, Sybase Gateways, Sybase Learning Connection, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase Synergy Program, Sybase Virtual Server Architecture, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SybFlex, SybMD, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, The Enterprise Client/Server Company, The Extensible Software Platform, The Future Is Wide Open, The Learning Connection, The Model For Client/Server Solutions, The Online Information Center, The Power of One, TotalFix, TradeForce, Transact-SQL, Translation Toolkit, Turning Imagination Into Reality, UltraLite, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viafone, Viewer, 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, XP Server, XTNDAccess and XTNDConnect are trademarks of Sybase, Inc. or its subsidiaries. 07/06

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 Sybase Risk Analytics Platform	1
Overview	1
Assumptions	1
Installation components	2
Core installation components	2
Optional installation components	2
Support tools	3
SySAM licensing requirements	4
Obtaining your SySAM license	4
Planning your Sybase RAP installation	5
Installation media	6
Preinstallation tasks	6
CHAPTER 2 Installing RAP Core Components	9
Overview	9
Installation layout	10
Run the RAP installer	11
Postinstallation tasks	12
Activate your SySAM license	12
Start the RAPCache server	14
Reload the RAPCache sample data	15
Verify that all servers are running	15
Use the ftp utility to move the data models to Windows	16
Postinstallation task for Linux users	17
Managing servers with Sybase Central	17
Starting Sybase Central	17
Registering VLDBServer with Sybase Central	18
Setting the System Administrator password	18
Creating RAP users	20
Installing client tools	22
Adaptive Server Client tools	22

- Sybase IQ Network Client tools..... 24
- Managing client connections 27
 - Adaptive Server..... 27
 - Sybase IQ..... 27
 - Distribute the connection information to users 28
- Uninstalling RAP components..... 28

CHAPTER 3 Installing Sybase OpenSwitch..... 29

- Overview 29
 - Minimum requirements 29
 - Implementation options 30
 - Where to find more information 30
- Preinstallation tasks 31
 - Install Sybase RAP core components 31
 - Check the operating system requirements 31
 - Check for operating system patches 31
 - Read the OpenSwitch documentation..... 32
- Installing OpenSwitch..... 32
- Configuring OpenSwitch 33
 - What the configuration tool does..... 33
 - Before you begin configuring OpenSwitch 34
 - Starting the configuration tool..... 35
 - Configuring OpenSwitch manually 36
- Postinstallation tasks..... 37
 - Updating environment variables..... 37
 - Starting OpenSwitch..... 37
 - Verifying client connections..... 38
 - Stopping OpenSwitch..... 39

CHAPTER 4 Loading TAQ Data 41

- Overview 41
 - How data loading works 42
 - Task summary..... 42
- Load process directories 43
- RAP load scripts..... 44
- Setting the environment variables 45
- Starting and stopping the load process..... 47
 - Loading the TAQ sample data..... 47
 - Loading additional TAQ sample data 49
- Checking the load process 51
- Clearing the cache 52
- Adding new instruments 52

APPENDIX A	Sample ASE Configuration Settings	53
	Introduction	53
	General information.....	54
	Backup/Recovery	54
	Named Cache:c_log	55
	Named Cache:default data cache	56
	Disk I/O.....	57
	Unicode	57
	Network Communication	58
	Physical Resources.....	58
	Physical Memory	59
	Processors	59
	SQL Server Administration.....	60
	User Environment.....	61
	Lock Manager.....	62
	Component Integration Services	63
	DTM Administration.....	64
	Cache memory configuration example.....	65
APPENDIX B	Data Replication.....	67
	Overview	67
	Replication environment.....	67
	Sybase Central.....	68
	Source database setup	69
	Preliminary steps.....	69
	Creating the schema for the ASE source database	69
	Replication Server Configuration	70
	Configuration tasks.....	70
	RSM configuration.....	71
	Target database configuration.....	71
	Replication Server Manager	74
	Replicating test data.....	77
	Troubleshooting	78
	Data fails to replicate	78
	While trying to replicate data from an ASE table to an IQ table owned by DBA, the connection to IQ comes DOWN.....	78
	Creating Publications within RSM	79
	IQ connection remains in a DOWN state upon creation.....	80
	Enabling RepServer transact SQL tracing	80
	Enabling Sybase IQ tracing	81
	Rounding with the FLOAT data type	81
	Defining a new connection using RSM may fail with execute permission error.....	81
	Sybase IQ runs out of IQ STORE dbspace.....	81

Sybase IQ database status DOWN and lower LTL version error reported	82
Sybase IQ permission denied error.....	82
Index	85

About This Book

Audience	<i>Risk Analytics Platform Installation and Configuration Guide</i> is intended for Sybase® Professional Services and other technical personnel who need to install and configure the components of Sybase Risk Analytics Platform. Familiarity with Sybase Adaptive Server® Enterprise, Sybase Replication Server®, Sybase IQ, data warehousing, and other related topics is assumed.
How to use this book	Before following the instructions in this book to install and configure Risk Analytics Platform, refer to the <i>Release Bulletin Risk Analytics Platform</i> for any last minute information regarding this product.
Related documents	Refer to the following documents for more information: <ul style="list-style-type: none">• <i>Release Bulletin Risk Analytics Platform</i>• <i>Risk Analytics User's Guide</i>• <i>Risk Analytics Platform Administration Guide</i>• Sybase IQ 12.6 product documentation• Adaptive Server Enterprise 15.0 product documentation• OpenSwitch 15.0 product documentation• PowerDesigner® 11.1 product documentation• Replication Server 12.6 product documentation• White paper titled Time Series in finance: the array database approach at http://cs.nyu.edu/shasha/papers/jagtalk.html• White paper titled FinTime --- a financial time series benchmark at http://www.cs.nyu.edu/cs/faculty/shasha/fintime.html

Note This product includes software developed by The Apache Software Foundation at <http://www.apache.org/>.

Other sources of information

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

- The Getting Started CD contains the release bulletin, installation and configuration guide, administration guide, and user's guide in PDF format. 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 Sybase Infocenter Web site is an online version of the product manuals that you can access using a standard Web browser.

To access the Infocenter Web site, go to Sybooks Online Help at <http://infocenter.sybase.com/help/index.jsp>

- 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.

Note The Product Family for Risk Analytics Platform 3.0 is Sybase IQ.

- 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.

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.

The Risk Analytics Platform 3.0 documentation complies 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.

For information about accessibility support in the Sybase IQ plug-in for Sybase Central, see “Using accessibility features” in Chapter 1, “Introducing Sybase IQ” in *Introduction to Sybase IQ*. The online help for Sybase IQ, which you can navigate using a screen reader, also describes accessibility features, including Sybase Central keyboard shortcuts.

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.

Sybase Risk Analytics Platform

About this Chapter

This chapter tells you how to install and configure Sybase Risk Analytics Platform.

Contents

Topic	Page
Overview	1
Installation components	2
SySAM licensing requirements	4
Planning your Sybase RAP installation	5

Overview

Sybase Risk Analytics Platform (RAP) is a consolidated risk and trade data repository, which manages high-volume real-time data feeds as well as massive historical and corporate reference data sets.

Sybase RAP captures high-frequency inbound data in one or more cache databases; long-term trend data resides in a scalable, high-capacity data repository.

Assumptions

The RAP installation instructions assume you are familiar with:

- The components of your platform, such as directories, files, and environment variables
- All platform-specific commands used to manipulate the software and hardware, such as those for changing directories and mounting the CD
- Sybase servers
- Failover systems

Installation components

Sybase RAP is delivered with all the components you need to install and configure sample databases, including data models, DDL scripts to create the databases, database configurations, sample data, and a set of queries you can use to test installation and fine-tune performance.

Core installation components

Sybase RAP is built on an infrastructure of core components, which includes the in-memory data cache, historical data repository, scripts, and documentation.

Component	Description	Disk Space
RAPCache Database	RAPCache is a fully functional, multi-threaded, T-SQL compliant cached RDBMS database built with Sybase Adaptive Server Enterprise 15.0 (ASE). ASE features an asynchronous commit option that allows inbound data to be stored in an in-memory cache, ensuring extremely fast load and retrieval times.	6GB
RAP VLDB Database	RAP VLDB is a disk-based, very large capacity (~petabyte range) database for consolidated historical data. Built with Sybase IQ 12.6, RAP VLDB includes a high-performance query engine, which provides fast access to very large data sets.	3GB

Optional installation components

Optional installation components include tools that manage high-availability deployment options and monitor activities within Sybase RAP.

Component	Description	Disk Space
OpenSwitch	<p>OpenSwitch is an application placed between client connections and two or more Sybase Adaptive Server Enterprise servers, which transfers incoming connections to any Sybase server product in response to an administrative request or failure in a high-availability environment.</p> <p>OpenSwitch requires a minimum of 256MB RAM and 450MB hard disk space. Verify that you have 100MB free space in your temporary directory.</p> <p>See Chapter 3, “Installing Sybase OpenSwitch” for additional information.</p>	450MB

Component	Description	Disk Space
Operations Console	<p>RAP Operations Console provides tools that monitor and manage activities within the Risk Analytics Platform, including:</p> <ul style="list-style-type: none"> • Database status • Query performance • Configuration management • Exception handling • Data loading <p>See <i>Risk Analytics Platform Administration Guide</i> for more information.</p>	400MB

Support tools

Support tools include Sybase PowerDesigner and Sybase RAP product documentation.

Component	Description	Disk Space
PowerDesigner	<p>Risk Analytics Platform includes a single seat license of the Physical Architect module of Sybase PowerDesigner 11.1. PowerDesigner is a Windows application that allows you to develop, import, and enhance existing data models as well as manage versions across the instances of the RAP Cache and the RAP VLDB databases.</p> <p>PowerDesigner requires a minimum of 256MB RAM, an SVGA or higher-resolution graphics adapter and compatible color monitor (800x600), and 100MB to 160MB disk space. Depending on installation options, more disk space may be required.</p>	160 MB
Data Models	<p>Sybase RAP includes separate data models for RAPCache and VLDBServer. The installer installs the data model and DDL scripts when you install the core components.</p>	
Documentation	<p>Documentation is included as part of the core components. The installer installs the RAP documentation in <i>/Documentation</i> in the root installation directory when you install RAPCache and VLDBServer. Check the RAP Getting Started CD and the Web site Product Manuals at http://www.sybase.com/support/manuals/ for the latest version of the documentation.</p>	

SySAM licensing requirements

Sybase Risk Analytics uses Sybase Software Asset Management (SySAM) to perform license administration and asset management tasks.

The Sybase Product Download Center (SPDC) enables authorized Revenue Order Contacts and Technical Support Contacts with current support contracts with download access to Sybase product releases. You must use SPDC to dynamically generate license keys for the Risk Analytics Platform.

Accessing SPDC requires a separate login and password; authorized contacts automatically receive SPDC login and password information through e-mail.

Refer to the SySAM online documentation at <http://infocenter.sybase.com/help/index.jsp> or the *User's Guide Sybase Software Asset Management 2.0* on the Getting Started CD for more information on the use of SySAM.

Obtaining your SySAM license

- 1 Start your browser and log in to the Sybase Product Download Center (SPDC) at <https://sybase.subscribenet.com>.
- 2 On the Product List page, choose Sybase IQ.
- 3 On the Product Information page, choose RAP 3.0 as your platform.
- 4 On the Product Download page, click the License Keys link.
- 5 On the Sybase Software Asset Management License Information page, select the radio button to the left of the applicable license. Click Select to Generate.
- 6 On the Generate Licenses page in Step 1 - License Model, choose the appropriate license model. Click Next.

Questions about your license type? Click the more information link to learn more about the licenses Sybase offers.
- 7 On the Generate Licenses page in Step 2 - License Quantity, specify the number of licenses you want to generate. Click on the links for more information.
- 8 On the Generate Licenses page in Step 3 - Host Information, add the Server Host ID, Host Name, and (optional) Port Number. Click Generate.

Click on "What's My Host ID," if you need help determining your host ID.

The View Licenses page displays your license information.

9 Click Download License File to save your license file to disk.

You activate your SySAM license after you install Sybase Risk Analytics Platform 3.0 on your host machine.

Planning your Sybase RAP installation

Sybase Risk Analytics Platform loads inbound market data to an in-memory cache database and a historical data repository. Users can track quotes and trade data as the information arrives, or construct queries that leverage this information with historical market data.

RAPCache

The in-memory cache database server, RAPCache, is built with Adaptive Server Enterprise. RAPCache captures and loads inbound tick data into temporary databases, which provides immediate access to this information. Time-critical trading decisions can be based on the results returned by queries to the RAPCache. You can have multiple RAPCache databases on multiple servers, but there is only one historical database.

VLDBServer

The historical data repository, VLDBServer, is built with Sybase IQ. The inbound tick data is loaded in parallel into the VLDBServer database. End-of-day market transactions are also loaded and stored in the VLDBServer database.

High Availability

In an High Availability (HA) environment, you need a primary RAPCache server and a secondary RAPCache server. You can configure Sybase OpenSwitch to ensure a failover from a primary to a secondary server without any loss of data.

Operations Console

Operations Console monitors the availability, performance, and status of the Risk Analytics Platform. A Web-based management console interfaces with the central repository that contains information collected about specific system components. Operations Console can be deployed on a UNIX, Linux, or Windows system.

Installation media

Sybase RAP is distributed on several installation CDs.

CD	Description
1 – 4	Install Sybase RAP infrastructure and support components, including RAPCache, VLDBServer, scripts, and other support tools. Also includes OpenSwitch.
Operations Console	Installs Sybase RAP Operations Console. You can deploy Operations Console on UNIX, Linux, or Windows platforms. See <i>Risk Analytics Platform Administration Guide</i> for installation instructions.
ASE 15.0 PC-Client	Installs PowerDesigner Physical Architect 11.0 on Windows. After you install PowerDesigner, you must download PowerDesigner EBF 13831: 11.1.0 ESD #10 (or later) from the Sybase Downloads site.
Sybase IQ 12.6 Network Client	Installs the components required for connection to a Sybase IQ network server.

Preinstallation tasks

- Plan your deployment. Planning your deployment for production before you install Sybase RAP greatly simplifies postinstallation tasks.
- Read the release bulletins for the latest information on the products (Adaptive Server Enterprise, Sybase IQ, and so on) that you are installing. See the special installation instructions section in the release bulletin.
- Install operating system patches, if required. If you are having trouble starting the installer, make sure you have the required operating system patches for the Java Runtime Environment (JRE).
- Use the Adaptive Server Configuration Guide to review SySAM procedures and plan your client/server configuration.
- Make sure that there is sufficient available disk space and the path name of the directory does not contain any spaces.
- Check to ensure that the operating system meets the version-level, RAM, and network protocol requirements for your platform.
- Verify that your network software is configured. Sybase software uses network software even if Adaptive Server and Sybase IQ client applications are installed on a machine that is not connected to a network.

If you are having connection problems or need to verify your network configuration, ping the host.

- Allocate a shared memory segment that is at least as large as the Adaptive Server total logical memory configuration parameter.

The default total logical memory parameter on Sun is 47,104 2K pages (92MB) on the 64-bit operating system. To adjust the shared memory value of the operating system on Sun Solaris, add the following line to the operating system configuration file */etc/system*:

```
set64 shmsys:shminfo_shmmax = <new_value_in_bytes>
```

To adjust the shared memory value of the operating system on Linux, add the new value in bytes to the file */proc/sys/kernel/shmmax*, so that the new value takes effect immediately. Also add a line to the file */etc/sysctl.conf*, so that the new value is retained the next time the machine is rebooted.

Installing RAP Core Components

About this Chapter

This chapter describes installation and configuration procedures for Sybase RAP core components.

Contents

Topic	Page
Overview	9
Postinstallation tasks	12
Managing servers with Sybase Central	17
Installing client tools	22
Installing PowerDesigner PhysicalArchitect	23
Managing client connections	27
Uninstalling RAP components	28

Overview

Use these instructions to install the Sybase RAP core components. RAP core components include RAPCache and VLDB servers, data models, scripts, and support tools. There are two installation options: a *typical install* and *custom install*.

What a typical install does

A typical install creates a main level directory as specified by the user (for example, */RAP30*), and installs RAPCache, VLDBServer, and other components into separate subdirectories on a single machine.

What a custom install does

A custom install allows you to choose the components you want to install and where you want to install them. Like a typical install, a custom install creates a common directory structure and installs the appropriate components into separate subdirectories on one or more target machines.

Note You must run the install for each RAPCache server you want to install. If you intend to deploy Sybase RAP in a High Availability (HA) environment, or install optional components on different machines, you must run the installer multiple times.

Installation layout

In a typical install, you specify the location of the installation directory, for example, */RAP30*.

Note Throughout this document, the environment variable \$RAP30 refers to the RAP 3.0 installation directory.

The default installation directory contains the following subdirectories.

Directory	Contents
Data	Contains subdirectories that store the sample data for RAPCache and VLDB
Model	Contains subdirectories that store the data models for RAPCache and VLDB
Scripts	Contains subdirectories that store data scripts for RAPCache and VLDB
VLDBServer	Contains a fully deployed Sybase IQ server, IQ support tools, and a historical database populated with a limited set of sample data
RAPCache	Contains a fully deployed ASE server, ASE support tools, and cache databases populated with a limited set of sample data
Auto_File_Load	Contains the load process directories and data loading scripts
Documentation	Contains Risk Analytics Platform 3.0 documentation in Adobe Acrobat PDF format and Microsoft Excel format
_jvm	Java Virtual Machine (jvm)
_unist	Uninstall information

In a custom install, you choose the server you want to install and where you want to install this server. The subdirectory structure remains the same, but includes only the server you chose to install.

Run the RAP installer

- 1 Log in as root. Stop and shut down all programs before running the installer.
- 2 Insert the Risk Analytics Platform installation CD #1 into the CD drive.
If Risk Analytics Setup does not start automatically, mount the drive, then use this command to start the installer:

```
<CD_mount_directory>/rap_cd1/setup<Platform>.bin
```

where *<Platform>* is the name of your operating system. For example, use the command `/cdrom/rap_cd1/setupSolaris.bin` on Sun Solaris.

When you run the setup command, your current directory can be any directory that is *not* on the CD.

- 3 Click Next on the Welcome screen.
- 4 Choose a location and read and accept the license agreement. Click Next.
If you do not find a license agreement that matches your location, or if the license agreement is unreadable on your system, you can read all available license agreements on the Sybase Web site at <http://www.sybase.com/softwarelicenses>.
- 5 Choose an installation directory and click Next.

You can accept the default target directory or choose another installation directory.

- If \$SYBASE is set, the default target directory is `$SYBASE/RAP30`.
- If \$SYBASE is not set, the default target directory is `/opt/sybase/RAP30`.

- 6 Choose an installation type:

If you choose...	The installer...
Typical install	Installs all core components in subdirectories of the installation directory on a single machine.
Custom install	Allows you to choose the components you want to install and where you want to install them. Choose this option if you want to install components on more than one machine.

- 7 Follow the instructions on the screen to complete your installation.

Notes

- You must run the install for each RAPCache server you install. If you install multiple servers or deploy components on different machines, you must run the install multiple times.
- If you transfer installed files to another UNIX or Linux machine using the ftp utility, execute this command:

```
chmod -R 777 *
```

after the transfer to set executable permission on the files.

Postinstallation tasks

The VLDBServer and RAPCache servers are running when the installation completes. After you install the Sybase RAP core components, shut down the RAPCache server, activate your SySAM license, restart the RAPCache server, and check the status of your servers.

Activate your SySAM license

- 1 After installation of the RAP core components, check the SySAM license type (the default is CP) and change the type, if necessary.

To check the license type, run the stored procedure `sp_lmconfig` on your RAPCache server:

```
1> sp_lmconfig
2> go
```

To change the license type to SF or DT to match your license file, run the appropriate command:

```
1> sp_lmconfig 'license type', 'SF'
2> go
```

or

```
1> sp_lmconfig 'license type', 'DT'
2> go
```

- 2 Shut down the RAPCache server.

- 3 Copy the license file saved from SPDC to the `$RAP30/RAPCache/SYSAM-2_0/licenses` directory on the host machine. (See “Obtaining your SySAM license” on page 4.)

If you use the ftp utility to transfer the license file, be sure to specify ASCII mode.

- 4 If you are using a served license, use the following command in the `$RAP30/RAPCache/SYSAM-2_0/bin` directory to start the license server:

```
sysam start
```

You must have at least one served license in the `licenses` directory `$RAP30/RAPCache/SYSAM-2_0/licenses`, or the error “License Manager: Can’t initialize...” is reported when you attempt to start the license server.

- 5 Restart the RAPCache server as described in “Start the RAPCache server” on page 14.

To verify your license, you can run the `sp_lmconfig` stored procedure in Interactive SQL (isql) or check the ASE error log. In the ASE error log, a line beginning with “Checked out license...” indicates that the license configuration is successful. If you see a “Sysam: FLEXnet Licensing error:” message, check with your Sybase representative to resolve the issue.

See also

For more information on the use of SySAM, refer to:

- “SySAM Configuration” in the *Adaptive Server Enterprise Version 15.0 Configuration Guide* for your platform
- SySAM online documentation at <http://infocenter.sybase.com/help/index.jsp>
- *User’s Guide Sybase Software Asset Management 2.0* on the Getting Started CD
- “Troubleshooting SySAM Issues” in the *Adaptive Server Enterprise 15.0 Installation Guide* for your platform

Start the RAPCache server

After you install the core components and activate your SySAM license, restart the RAPCache server. To start RAPCache, you execute a *RUN_<server_name>* file that contains the information required to start or restart the server.

- 1 Open a command window.
- 2 Change directory to *\$RAP30/RAPCache*.
- 3 Use the appropriate shell command to source the environment variables:

```
source SYBASE.csh
```

or

```
./SYBASE.sh
```

- 4 Change directory to *\$RAP30/RAPCache/ASE-15_0/install*.
- 5 Use this command to start the RAPCache server:

```
./startserver -f RUN_<server_name>
```

where *<server_name>* is the name of the server.

For example, the *RUN_<server_name>* file name for an Adaptive Server named *rapserver* is *RUN_rapserver*.

Reload the RAPCache sample data

Every time the RAPCache server is stopped, the TAQ sample data contained in the in-memory cache is lost. You need to reload this data in order to run the RAP sample queries.

Before reloading the TAQ data, you can verify that your RAPCache server is running by following the procedure in “Verify that all servers are running” on page 15.

To reload the TAQ sample data into the RAPCache, run the script `$RAP30/Data/RAPCache/RAP_Load/load_unix_ASE_TAQ_tables.sh` from a command window on the machine that hosts the ASE RAPCache server. This script uses the bulk copy utility `bcp` to load the sample data into the RAPCache `rap` and `rapq` databases. Errors are written to the files `STOCK_TRADE.err` and `STOCK_QUOTE.err`.

To check the status of the data loads, review the log file for Adaptive Server Enterprise. The default location of the ASE log file is `$RAP30/RAPCache/ASE-15_0/install/<server_name>.log`.

Verify that all servers are running

After you start RAPCache, check to see that your servers are running and that you can connect to the servers.

RAPCache

The `showserver` command displays all Adaptive Server-related processes on the system. Use this command to determine whether the RAPCache server is running.

- 1 Open a command window.
- 2 Change directory to `$RAP30/RAPCache`.
- 3 Use the appropriate shell command to source the environment variables:

```
source SYBASE.csh
```

or

```
./SYBASE.sh
```

- 4 Run this command:

```
$RAP30/RAPCache/ASE-15_0/install/showserver
```

- 5 Check the message for references to: `...$RAP30/RAPCache/ASE-15_0/bin/dataserver...`

VLDBServer

VLDBServer should be up and running after installation. You can use the `stop_asiq` utility to verify the current status of the VLDBServer.

- 1 Open a command window.
- 2 Change directory to `$RAP30/VLDBServer/ASIQ-12_6`.
- 3 Use one of these commands to source the environment variables:

```
source ASIQ-12_6.csh
```

or

```
./ASIQ-12_6.sh
```

- 4 Run this command:

```
$RAP30/VLDBServer/ASIQ-12_6/bin/stop_asiq
```

VLDBServer displays a message that indicates the current status of your server, and prompts you to decide whether you want to shut down the server.

- 5 Type `N` to respond to the prompt.
- 6 If the VLDBServer is not running, change directory to `$RAP30/VLDBServer/RAPDB` and enter:

```
start_asiq @<VLDB_server>.cfg <VLDB_server>.db
```

where `<VLDB_server>` is the name of your VLDB server.

Use the ftp utility to move the data models to Windows

The Sybase RAP installer installs the RAPCache and VLDB data models in separate subdirectories in the `$RAP30/Model` directory. If you want to view the data models with PowerDesigner, use the ftp utility to move the data models from UNIX to Windows.

See “Installing PowerDesigner PhysicalArchitect” on page 23 for information on installing PowerDesigner.

Postinstallation task for Linux users

As you tune the RAPCache server by adding named caches or increasing the sizes of existing caches, the server may stop rebooting at some point. The memory allocated to the server is large enough to accommodate the additional cache size, and the operating system level parameter SHMMAX (the maximum size in bytes for a shared memory segment) is also large enough, but the server does not restart. The problem may be that the operating system level parameter SHMALL, the system-wide maximum of shared memory pages, is not large enough. To resolve this issue, increase the value of SHMALL in the file `/proc/sys/kernel/shmall`. Modifying `/proc/sys/kernel/shmall` requires root permission.

Managing servers with Sybase Central

Sybase Central is a management console for Sybase products. Sybase Central lets you manage multiple servers from a single location, and lets you perform complex administration tasks without the need to remember the Transact-SQL commands or system stored procedures.

Starting Sybase Central

- 1 Open a terminal window, and change directory to `$RAP30/RAPCache`.
- 2 Use one of these commands to source the environment variables:

```
source SYBASE.csh
```

or

```
./SYBASE.sh
```

- 3 Use this command to start Sybase Central:

```
$RAP30/RAPCache/ASEP/bin/aseplugin
```

- 4 Select Tools | Connect, then choose a server from the drop-down box.

- 5 On the Connect dialog, enter the user name and password of the System Administrator in the appropriate text boxes.

Warning! The first time you log in to Adaptive Server, use the default `sa` user name and leave the password blank. After you log in for the first time, change the password of the System Administrator. See “Setting the System Administrator password” on page 18.

- 6 Choose the name of your RAPCache server as the Server Name.
- 7 Click OK.

Registering VLDBServer with Sybase Central

To access the VLDBServer from Sybase Central, you must register the Sybase IQ plug-in.

- 1 Click Tools | Plug-ins.
- 2 On the Plug-ins panel, click Register.
- 3 On the plug-in registration dialog, click Browse.
- 4 Change the path to `$RAP30/VLDBServer/ASIQ-12_6/java`.
- 5 Choose `iq.prg` and click OK.
- 6 Click Finish to close the plug-in registration dialog.
- 7 Close the Plug-ins panel.

Setting the System Administrator password

RAPCache and VLDBServer include different administrator default logins. Sybase recommends that you change these logins to protect your system.

RAPCache
administrator

The RAPCache installation creates an administrator called `sa` for the Sybase System Administrator. There is no initial password set for this account. Because an `sa` user can use any database on Adaptive Server with full privileges, Sybase recommends that you create a password to protect your production environment.

- 1 In the left pane, click Adaptive Server Enterprise | Default | *<RAPCache_server>*, where *<RAPCache_server>* is the name of your RAPCache server.
- 2 Click the Logins folder.
- 3 In the right pane, right-click *sa* and choose Properties.
- 4 On the Login Properties dialog, click the Parameters tab, then click Change Password.

The Parameters tab also includes options that allow you to set the minimum password length, maximum login attempts, and expiration period.

- 5 On the Change Password dialog:
 - Type your password in New Password box.
 - Type your password in New Password box, then retype the password in the Confirm New Password box.
- 6 Click OK.

Note Sybase recommends that you create passwords with at least six characters that include both letters and numbers.

VLDBServer administrator

The VLDBServer installation creates an administrator account with a default user ID (DBA) and password (SQL). The DBA user ID identifies a user with full administration and resource creation rights. The first time you log in to the VLDBServer, change the DBA password to protect your production environment.

- 1 In the left pane, right-click the Sybase IQ icon and choose Connect.
- 2 On the connection dialog, click the Database tab.
- 3 From the Server name drop-down box, enter *host_name:port_number*, where *host_name* is the name of the machine that hosts the VLDB server, and *port_number* is the server port number.
- 4 Click the Browse button next to the Database file drop-down box and choose *<VLDB_database_name>.db*, where *<VLDB_database_name>* is the name of your RAP VLDB database. Click OK.
- 5 Click the Identification tab.
- 6 Type *DBA* in the User id box, type *SQL* in the Password box, and click OK.

You are now connected to the RAP VLDB database on the VLDB server. Expand Sybase IQ in the left panel to see the server name. Expand the server name to see the database name.

- 7 Expand the database folder, then expand the Users & Groups folder.
- 8 Right click DBA and select Properties in the pop-up menu.
- 9 In the DBA User Properties dialog, enter and confirm the new DBA password.
- 10 Click OK to save and exit the dialog. The new password takes effect immediately for new DBA connections.

Creating RAP users

You must create separate user and group accounts for RAPCache and VLDBServer. Before you create these accounts, be aware that user and group models in Adaptive Server Enterprise and Sybase IQ differ.

How users connect

- All ASE users require a server login ID and password and a user ID for each database they want to access on that server.
- Sybase IQ users do not require a server login ID. All Sybase IQ users receive a user ID and password for a database.

User groups

ASE and IQ groups differ:

- ASE allows each user to be a member of only one group.
- In IQ, users and groups are database objects. Groups can contain users and other groups. Users can belong to multiple groups, and group hierarchies are allowed.

Database object permissions

GRANT and REVOKE statements for granting permissions on individual database objects are very similar.

- Both ASE and Sybase IQ allow SELECT, INSERT, DELETE, UPDATE, and REFERENCES permissions on database tables and views, and UPDATE permissions on selected columns of database tables.

For example, the following statement is valid in both products:

```
GRANT INSERT, DELETE
ON TITLES
TO MARY, SALES
```

This statement grants permission to use the INSERT and DELETE statements on the TITLES table to user MARY and to the SALES group.

- Both products allow EXECUTE permissions to be granted on stored procedures.
- ASE also supports GRANT and REVOKE on additional items:
 - Objects: columns within tables, columns within views, and stored procedures
 - User abilities: Create Database, Create Default, Create Procedure, Create Rule, Create Table, Create View
- Sybase IQ requires a user to have RESOURCE authority to create database objects. (A closely corresponding Adaptive Server Enterprise permission is GRANT ALL, used by a Database Owner.)
- Both products support the WITH GRANT OPTION clause, allowing the recipient of permissions to grant them in turn, although IQ does not permit WITH GRANT OPTION to be used on a GRANT EXECUTE statement.

Database-wide permissions

Adaptive Server Enterprise and Sybase IQ use different models for database-wide permissions.

- Sybase IQ employs DBA permissions to allow a user full authority within a database.
- ASE System Administrators enjoy database-wide user permission for all databases on the server, but the database owner must use a SETUSER statement to gain permissions on objects owned by other users.

Adding users

ASE and Sybase IQ use different models to create users.

- ASE requires separate server and database accounts. sp_addlogin creates a server login account for a new user. sp_adduser gives database and group privileges.
- Sybase IQ adds users in a single step, but does not assign permissions beyond connecting to the database and viewing the system tables. Permissions to access database tables must be explicitly assigned.
- Sybase IQ does not require stored procedures to add or drop users, but does allow DBAs to add or drop user accounts. When a DBA enables IQ User Administration, these accounts let DBAs control user connections and password expirations.

Where to find additional information

For details about managing Adaptive Server login accounts and database users, see “Managing Adaptive Server Logins, Database Users, and Client Connections” in the *Adaptive Server Enterprise System Administration Guide, Volume 1*.

For details about Sybase IQ User Administration, see “Managing User IDs and Permissions” and “Sybase IQ as a Data Server,” in the *Sybase IQ System Administration Guide*.

Installing client tools

You must install the appropriate tools on each client machine to connect to the Sybase RAP databases. Adaptive Server client tools are required for a client to connect to RAPCache. Sybase IQ client tools are required for a client to connect to the VLDBServer.

Adaptive Server Client tools

The ASE 15 PC-Client CD contains several products packaged with their own installers. Before installing any products, review the *readme.txt file*, which contains a brief description of each of the products, lists dependencies, and contains any last-minute instructions or changes.

See “Installing PowerDesigner PhysicalArchitect” on page 23 for specific instructions on installing PowerDesigner 11.1.

Installing the client tools on Windows

Because some portion of the PC-Client components are required by the other products, install the PC-Client components first.

- 1 Use an account with Administrator privileges to log in to Windows.
- 2 Exit any Windows programs running on your machine.
- 3 Insert the PC-Client CD into the CD drive.

If PC-Client Setup does not start automatically:

- Click Start and choose Run.
 - Browse to select your CD drive letter and choose *setup.exe*.
- 4 Follow the instructions on your screen to install the PC-Client tools.

Configuring network connections for client products

To configure client network connections to Adaptive Server, see *Configuring Adaptive Server Enterprise* for your platform.

Adaptive Server communicates with other Adaptive Servers, Open Server applications (such as Backup Server), and client software on your network. Clients can talk to one or more servers, and servers can communicate with other servers by remote procedure calls.

For Sybase products to interact with one another, each product must know where the others reside on the network. This information is stored in the interfaces file (*sql.ini* on Windows), or in a Lightweight Directory Access Protocol (LDAP) server.

For additional information, see “Installing Sybase PC-Client Products” in the *Adaptive Server Enterprise Installation Guide* for your platform. This document includes detailed installation instructions for Adaptive Server client products like the Adaptive Server plug-in to Sybase Central.

Installing PowerDesigner PhysicalArchitect

PowerDesigner 11.1 provides full support for the Sybase RAP data models. To install PowerDesigner 11.1, you must first install PowerDesigner 11.0 from the ASE 15 PC-Client CD included with your Sybase RAP installation package, then download and install PowerDesigner EBF 13831: 11.1.0 ESD #10 (or later).

- 1 Insert the ASE 15 PC-Client CD into your CD drive on a Windows system.
If Setup does not start automatically:
 - Click Start and choose Run.
 - Browse to select your CD drive letter and choose *setup.exe*.
- 2 Follow the instructions on your screen to install PowerDesigner. For additional information, see the PowerDesigner product documentation that is installed with the software.

Installing EBF 13831: 11.1.0 ESD #10

Use these instructions to download and install PowerDesigner EBF 13831: 11.1.0 ESD #10 (or later).

- 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 Choose PowerDesigner.
- 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 next to EBF 13831: 11.1.0 ESD #10 (or the latest EBF available) to display the EBF/Maintenance report, or click the product description to download the software.
- 6 Use the instructions in the EBF/Maintenance report to install the maintenance release.

Sybase IQ Network Client tools

This section includes procedures that tell you how to install the Sybase IQ Network Client on Windows and Linux.

You can also install the IQ Network Client on many end-user machines remotely, without displaying dialog boxes or requiring interactive responses. For directions, see “Installing without human interaction” in the *Sybase IQ Installation and Configuration Guide* for your platform.

Installing client tools on Windows

- 1 Use an account with Administrator privileges to log in to Windows.
- 2 Exit any Windows programs running on your machine.
- 3 Insert the IQ Network Client CD into the CD drive.

If Network Client Setup does not start automatically:

 - Click Start and choose Run.
 - Browse to select your CD drive letter and choose *setup.exe*.
- 4 Follow the instructions on your screen to install the IQ Network Client tools.

Installing client tools
on Linux

If you installed Sybase Central, see the *Sybase IQ System Administration Guide* for instructions on configuring and running the IQ Agent. You must run the IQ Agent in order to use Sybase Central.

Sybase IQ Network Client for Linux contains the components required for connection to a network server, and is compatible with IQ servers on all supported server platforms. Avoid installing Sybase IQ Network Client in the same directory as an IQ server.

- 1 Use this command to change to the installation directory:

```
% cd $RAP30/VLDBServer
```

- 2 Start the install utility sybinstall.

You can run this utility as a series of menus with prompts or bypass menus using the sybinstall command line parameters:

Parameter	Function
-add_agent	Install standalone 12.6 IQ Agent only
-autoinstall	Install all defaults
-help	Display all parameters and usage
-I_accept_sybase_license	Bypass license agreement
-info version	Display information about this product
-y	Assume yes for all questions, warnings, and errors

- 3 Do one of the following:

- Run sybinstall with all menus and prompts. This allows you to choose the components you want to install:

```
% /cdrom/sybinstall
```

- Run sybinstall to install the default products:

```
% /cdrom/sybinstall -autoinstall  
-I_accept_sybase_license
```

- 4 When the Welcome screen appears, press Return.
- 5 Type the number that corresponds to your location.

If you cannot read the license agreement or find a license agreement that matches your location, you can read all available license agreements on the Sybase Web site at <http://www.sybase.com/softwarelicenses> and rerun sybinstall, passing it the parameter `-I_accept_sybase_license`. For example:

```
% /cdrom/sybinstall -I_accept_sybase_license
```

The first screen of the Software Test and Evaluation License Agreement appears.

- 6 Read and accept the license agreement. As you read, hold down the Return key until you reach the end of the agreement.

The script next lists the amount of free space available in your *\$RAP30/VLDBServer* directory and the amount of space required for the three components it installs.

- 7 Do one of the following:
 - To install all four products, type *s*.
 - To deselect or change any of the installed products, type the option number at the prompt.
- 8 Type *s* to start the installation.

An installation procedure log is created in *\$RAP30/VLDBServer/syinstall.log*. If the log file cannot be created in the *\$RAP30/VLDBServer* directory, the path defaults to */tmp/syinstall.log*.

Setting Linux environment variables

After you install Sybase IQ Network Client for Linux, use these instructions to set the environment variables:

- 1 Open a command window.
- 2 Change directory to *\$RAP30/VLDBServer/ASIQ-12_6* and use one of the following commands to source the environment variables:

```
% source ASIQ-12_6.csh
```

or

```
% . ./ASIQ-12_6.sh
```

Managing client connections

Clients and servers communicate through a connection. For a client application to connect to a server application, the server application must be listening for the client connection request.

Adaptive Server

Clients, such as Sybase Central and Open Client, are used to access Adaptive Server. Some ASE client applications use ODBC drivers to connect to Adaptive Server. For example, PowerDesigner connects through an ODBC driver. Other third-party applications may also require an ODBC driver.

You can also install the driver separately on other client workstations on which you will be running third-party or developed products. For more information about the ODBC driver, see the *Adaptive Server Configuration Guide* for your platform and the *ODBC Driver Reference Guide* in Technical Library.

Sybase IQ

Sybase IQ requires an ODBC data source on the client computer for each database you want to access using ODBC. A data source describes how to get to data on a network. Configuration requirements are platform specific.

Windows

On Windows, Sybase recommends that you use a system DSN, which offers faster performance than a file DSN. You can also share system DSNs with all users on the network.

For information on creating ODBC data sources, see the chapter “Configuring Sybase IQ” in the *Sybase IQ Installation and Configuration Guide for Windows*.

UNIX and Linux

On UNIX or Linux operating systems, ODBC data sources are held in a file named *odbc.ini*. When you create an *odbc.ini* file, you must use the long form of each identifier, for example:

```
[My Data Source]
EngineName=myserver
CommLinks=tcpip(port=1870)
Userid=DBA
Password=SQL
```

For information about creating ODBC data sources, see the chapter “Configuring Sybase IQ” in the *Sybase IQ Installation and Configuration Guide* for your server platform.

Distribute the connection information to users

If you created File DSNs, you can make them available to users on the network or distribute the files. Otherwise, make the connection information available to users, so they can create the data sources on their machines.

Uninstalling RAP components

Before uninstalling the RAP components, shut down all RAPCache and VLDB servers. See “Verify that all servers are running” on page 15 for information on determining the current status of your servers.

To uninstall RAP, run `$RAP30/_uninst/uninstall.bin`.

Installing Sybase OpenSwitch

About this Chapter

This chapter tells you how to install and configure Sybase OpenSwitch.

Contents

Topic	Page
Overview	29
Preinstallation tasks	31
Installing OpenSwitch	32
Configuring OpenSwitch	33
Postinstallation tasks	37

Overview

Sybase OpenSwitch can provide near-continuous database access during planned downtime and unplanned outages. OpenSwitch is placed between client connections and two or more Sybase RAPCache servers. During a failover or switchover, OpenSwitch transfers access from the primary RAPCache server to the next available server in the pool.

Minimum requirements

Deploying Sybase RAP in a High Availability environment with OpenSwitch requires at least two nodes to provide an Active/Active standby capability. RAP read/write applications apply inbound transactions in two parallel streams to both the Primary and Standby RAPCache databases. After failover, data flow to the Standby database meets the same latency conditions that existed before failover.

To ensure high availability in a production environment, Sybase strongly recommends that you install OpenSwitch on a host other than the host where Adaptive Server and Replication Server are installed.

Implementation options

Several configurations are possible, depending on whether you use Replication Server, whether the Replication Server is configured for warm-standby, and whether your Adaptive Servers are configured for failover.

Name	Software requirements
Simple	<ul style="list-style-type: none"> • 1 OpenSwitch • 1 Optional Coordination Module (CM) • 2 Adaptive Servers
High availability, warm-standby	<ul style="list-style-type: none"> • 1 OpenSwitch server and 1 Replication Coordination Module (RCM) configured to coordinate failover through the OpenSwitch server OpenSwitch servers in this configuration do not communicate with each other, unlike a mutually aware configuration. • 2 Adaptive Servers configured for high availability • 1 Replication Server configured for warm-standby
Redundant high availability, warm-standby	<ul style="list-style-type: none"> • 2 OpenSwitch servers (1 primary and 1 secondary) • 2 RCMs configured to coordinate failover through the OpenSwitch server • 2 Adaptive Servers configured for high availability • 1 Replication Server configured for warm-standby
Mutually aware	<ul style="list-style-type: none"> • 2 mutually aware, companion OpenSwitch servers. Both mutually aware OpenSwitch servers within the same cluster regard each other as companions and both are aware of each other's state and the state of the other servers. • 2 CMs or 2 RCMs (optional) • 2 Adaptive Servers, which may be configured for high availability

Where to find more information

- See the *OpenSwitch Administration Guide* and the *OpenSwitch Coordination Module Reference Manual* for details about possible OpenSwitch configurations.
- See the Replication Server documentation for information about using warm-standby. See the Adaptive Server Enterprise documentation for information about high availability failover.

Preinstallation tasks

Read this section to learn about tasks you perform before you install OpenSwitch 15.0.

Install Sybase RAP core components

Install the RAPCache servers, then install OpenSwitch. Use the Sybase RAP installer to create a custom install, and install RAPCache for each server in your cluster.

Check the operating system requirements

OpenSwitch requires a minimum of 256MB RAM and 450MB hard disk space. The OpenSwitch installer requires 100MB free space in your temporary directory. See the table below for other hardware and software requirements.

Hardware	Operating system	Supported protocols
Sun Solaris (SPARC)	Solaris 2.8 with patch 108528-20 or higher	TCP
Linux	Red Hat Enterprise AS release 3	TCP

Check for operating system patches

Apply any operating system patches before you install OpenSwitch. See the table below to identify the patches on your system.

OS	Command
Sun Solaris (SPARC)	showrev -p
Linux	uname -a
AIX	lspp -h

Note Do not use a patch that is earlier than the version suggested for your operating system. Use the patch recommended by the operating system vendor, even if it supersedes the patch listed here.

Read the OpenSwitch documentation

Deploying OpenSwitch in an HA environment is not difficult but does require planning. Read the *OpenSwitch 15.0 Administration Guide* to learn about OpenSwitch functionality, features, and deployment issues. Refer to the *Installation Guide OpenSwitch 15.0* for specific installation and configuration issues.

Installing OpenSwitch

Follow these steps to run the RAP installer to install OpenSwitch.

- 1 Log in as `root`. Stop and shut down all programs before running the installer.

- 2 Insert the Risk Analytics Platform installation CD #1 into the CD drive.

If Risk Analytics Setup does not start automatically, mount the drive, then use this command to start the installer:

```
<CD_mount_directory>/rap_cd1/setup<Platform>.bin
```

where *<Platform>* is the name of your operating system. For example, use the command `/cdrom/rap_cd1/setupSolaris.bin` on Sun Solaris.

When you run the `setup` command, your current directory can be any directory that is *not* on the CD.

- 3 Click Next on the Welcome screen.

- 4 Choose a location and accept the license agreement.

If you do not find a license agreement that matches your location, or if the license agreement is unreadable on your system, you can read all available license agreements on the Sybase Web site at <http://www.sybase.com/softwarelicenses>.

- 5 Choose an installation directory.

You can accept the default target directory or choose another installation directory:

- If `$SYBASE` is set, the default target directory is `$SYBASE/RAP30`.
- If `$SYBASE` is not set, the default target directory is `/opt/sybase/RAP30`.

- 6 When the installer prompts you for the type of installation you want to perform, choose Custom. Click Next.
- 7 When the installer prompts you for the features you want to install, choose Open Switch (Cache DB HA Option). Deselect any other options. Click Next.
- 8 Read the installation summary. Click Install.
- 9 When the OpenSwitch Installer starts, follow the instructions on the screen.

Select Custom as the type of installation. Select Open Client, as well as other default components, on the list of features to install.

After installation, the installer identifies the OpenSwitch installation directory, and prompts you to decide whether you want to run the graphical configuration tool. If you want to configure OpenSwitch now, click Next. If not, click Cancel.

See “Configuring OpenSwitch” on page 33 for additional information about OpenSwitch Configuration.

Configuring OpenSwitch

OpenSwitch 15.0 includes a configuration wizard tool you can use anytime after you install the product. You can access the configuration tool from the OpenSwitch installer or as a standalone application after installation.

You must run the configuration utility for each OpenSwitch server that you install. A mutually aware implementation, for example, includes two separate OpenSwitch installations. You must run the configuration tool for each installation.

What the configuration tool does

The configuration tool creates several files that store OpenSwitch connection and configuration parameters.

Throughout this document, the environment variable \$RAP30 refers to the RAP 3.0 installation directory. As used in this chapter, \$SYBASE is the directory *\$RAP30/RAPCache*, unless OpenSwitch is installed in a different installation directory.

File name	Location	Description
<i>interfaces</i>	<i>\$\$SYBASE</i>	Identifies the servers in your implementation
<i><primary_OpenSwitch_server_name>.cfg</i>	<i>\$\$SYBASE/OpenSwitch-15_0/config</i>	Stores OpenSwitch configuration settings
<i>_rcm.cfg</i>	<i>\$\$SYBASE/OpenSwitch-15_0/config</i>	Stores Replication Coordination Module (RCM) configuration settings
<i>oswConfig.log</i>	<i>\$\$SYBASE/OpenSwitch-15_0/logs</i>	Logs OpenSwitch configuration errors
<i>osw.err</i>	<i>\$\$SYBASE/OpenSwitch-15_0/logs</i>	Logs OpenSwitch errors

Before you begin configuring OpenSwitch

- Identify all server names and port numbers for each OpenSwitch server and each Adaptive Server in your installation.

If you plan to install two OpenSwitch servers and two Adaptive Servers, you must provide the server name, host name, and port number for each server. If you plan to implement mutually-aware support, you must also include the server name, host name, and port number for the companion OpenSwitch server as the last line in the entry of the other OpenSwitch server.
- Identify an OpenSwitch Connection Monitor (CMON). CMON monitors Adaptive Servers and asynchronously notifies threads as soon as connectivity to the remote server is lost.
- Define your server pool. A pool is a group of servers within OpenSwitch. A pool can contain one or more servers that are treated as a self-contained failover group, so all connections within the group fail over only to servers defined within the group.
- Determine what actions to take when a failure occurs. OpenSwitch fails over to the next available server in the pool by default. If you plan to implement a failover solution in mutually-aware support, you can decide what action to take when a specific failover event occurs.

Starting the configuration tool

You can run the configuration utility as a standalone application after you install OpenSwitch. The configuration tool requires a JRE version of 1.4 or later. Check your environment to see that `JAVA_HOME` is set to a JRE version of 1.4 or later.

- 1 Change to the `$SYBASE` directory.

- 2 Use one of these commands to source the environment variables:

```
source SYBASE.csh
```

or

```
./SYBASE.sh
```

- 3 Set the `JAVA_HOME` environment variable. For example:

```
setenv JAVA_HOME $SYBASE/uninstall/OSWSuite/JRE-1_4
```

- 4 Start the OpenSwitch configuration tool:

```
$OPENSWITCH/CFG-1_0/bin/oswcfg.sh
```

- 5 Follow the instructions on the screen.

For step-by-step instructions, see the *OpenSwitch 15.0 Installation Guide*.

OpenSwitch configuration recommendations specific to Risk Analytics Platform are listed in the following table:

Configuration section	Recommended value
Coordination Module	None. As part of the default installation, you do not need to select a CM or RCM.
Authorization	Admin User Name: sa Admin Password: sa Leave other information deselected.
Failure Action	Accept all default options.

- 6 After exiting the configuration tool, check the log file(s) in the OpenSwitch `logs` directory and the `<install_directory>/oswInstall.log` to be sure the installation was successful and error free.

Configuring OpenSwitch manually

If you need to change settings, you can use a text editor to modify the interfaces and configuration files without restarting OpenSwitch. This allows you to change OpenSwitch behavior without interrupting user connections.

Interfaces file

The *interfaces* file stores information about the servers in your installation. There is one entry (server name, host name, and port number) for each OpenSwitch server and Adaptive Server in your installation. Verify that all of your servers are listed in the interfaces file and add the information for any missing servers.

Note When you reconfigure an OpenSwitch server, the existing interfaces file is copied to a backup file and appended with an *.001* extension (for example, *interfaces.001*) and saved in the logs directory. Similarly, the OpenSwitch configuration file is copied to *<OpenSwitch ServerName>.cfg.001* and the RCM config file is copied to *<OpenSwitch ServerName>_rcm.cfg.001* and saved in the config directory.

Configuration file

The OpenSwitch configuration file *<install_directory>/OpenSwitch/OpenSwitch-15_0/config/<OS_server_name>.cfg* contains these sections:

- [CONFIG] – defines OpenSwitch command line and configuration options, which are entered in this format: `NAME=VALUE`

Add the following option or modify the existing option in the [CONFIG] section:

```
RESPONSE_TIMEOUT=3600
```

Increasing `RESPONSE_TIMEOUT` to 3600 from the default value of 60 seconds can help query response time and avoid timeout errors.

- [SERVER] – defines the settings and status of Adaptive Servers available for use within OpenSwitch. Each remote server includes the following parameters:
 - `SERVER_NAME` identifies the remote server in the OpenSwitch interfaces file. `DEFAULT` sets the default values for any server not explicitly listed in this section.
 - `OPTION=VALUE` represent parameters specific to `SERVER_NAME`. See “Using the Configuration File” in the *OpenSwitch Administration Guide* for a list of valid parameters.
- [POOL] – defines a group of Adaptive Servers, including any attributes, servers, and connections that the pool serves.

- [LIMIT_RESOURCE] – specifies the connections that are candidates for resource governing and which resources to limit.
- [COMPANION] – identifies the name of the OpenSwitch companion, and the administrator user name and password used to make a connection.

Note If you manually edit the OpenSwitch configuration file to change configuration values, you may also need to edit the interfaces UNIX file. See the chapter “Configuring OpenSwitch” in the *OpenSwitch 15.0 Installation Guide* for more information.

Postinstallation tasks

After you install and configure OpenSwitch, perform the following tasks to update your environment and verify your client connections.

Updating environment variables

After you install and configure OpenSwitch, you may need to update the environment variables. To do so, use these instructions:

- 1 Open a command window.
- 2 Change to the `$$SYBASE` directory.
- 3 Use one of these commands to source the environment variables:

```
source SYBASE.csh
```

or

```
./SYBASE.sh
```

Starting OpenSwitch

If you need to start OpenSwitch, use these instructions:

- 1 Open a command window.
- 2 Change to the `$$SYBASE` directory.

- 3 Use one of these commands to source the environment variables:

```
source SYBASE.csh
```

or

```
. ./SYBASE.sh
```

- 4 Change to the `$OPENSWITCH/bin` directory.

- 5 Use this command to start OpenSwitch:

```
./OpenSwitch -c ../config/OpenSwitch.cfg
```

where `-c` specifies the OpenSwitch configuration file to use during start-up.

Verifying client connections

You can perform a sanity test to confirm that your Open Switch is running and can accept client connections.

- 1 Open a command window.

- 2 Use one of these commands to source the environment variables:

```
source SYBASE.csh
```

or

```
. ./SYBASE.sh
```

- 3 Use `isql` to log in to the OpenSwitch server:

```
isql -Usa -P -S<OS_server_name>
```

where `<OS_server_name>` is the name of the OpenSwitch server.

- 4 Execute the following command:

```
1> select @@servername  
2> go
```

The name of the Primary Cache server should display.

- 5 Use `isql` to log in to the OpenSwitch server as an Administrator:

```
isql -Usa -Psa -S<OS_server_name>
```

where `<OS_server_name>` is the name of the OpenSwitch server.

- 6 At the command prompt, enter:


```
1> rp_set TEXTSIZE
2> go
```

You should see:

```
parameter value-----TEXTSIZE 1048576 (1 row affected)
(return status = 0)
```

- 7 Use isql to log in to OpenSwitch as a client:

```
isql -Username -Ppassword -S<OS_server_name>
```

where:

- *username* is the client user name
- *password* is the client password
- *<OS_server_name>* is the name of the OpenSwitch server

- 8 At the prompt, issue the select @@version command to see the version information of the Adaptive Server to which OpenSwitch is connected. For example:

```
1> select @@version
2> go
```

```
Adaptive Server Enterprise/15.0/EBF 13444
ESD #2/P/Sun_svr4/OS 5.8/ase150/2193/64-bit
/FBO/Wed May 17 22:39:59 2006
```

Stopping OpenSwitch

Follow these steps to shut down OpenSwitch:

- 1 Use isql to log in to OpenSwitch as the user admin:

```
isql -S<OS_server_name> -Usa -Psa
```

where *<OS_server_name>* is the name of the OpenSwitch server.

- 2 At the command prompt, enter:

```
1> rp_shutdown
2> go
```


About this Chapter

This chapter tells you how to load the sample data.

Contents

Topic	Page
Overview	41
Task summary	42
Load process directories	43
RAP load scripts	44
Setting the environment variables	45
Starting and stopping the load process	47
Checking the load process	51
Clearing the cache	52
Adding new instruments	52

Overview

Sybase RAP uses scripts to load the inbound data to the RAPCache and VLDBServer databases. For these scripts to function correctly, you must create an appropriate directory structure and source all environment variables to allow the bcp and dbisql utilities to run.

Note Throughout this chapter, the environment variable \$RAP30 refers to the RAP 3.0 installation directory, and the environment variable \$RAP_HOME refers to the *\$RAP30/Auto_File_Load* directory.

How data loading works

Running *rap_flap_start.sh* starts a process that continually monitors an inbound data directory (*\$RAP_IN*). When new data arrives, load scripts compare the contents of the inbound directory against the archive (*\$RAP_ARCH*), and link any new files to the input (*\$RAP_DATA* and *\$IQ_DIR*) and archive directories.

A corresponding monitor moves the data to active directories (*\$RAP_DATA/active* and *\$IQ_DIR/active*), where a child process loads the data into the in-memory cache of the RAPCache database and the VLDB database. Log and error files generated by the load operation are placed in the appropriate log directories (*\$RAP_DATA/log* and *\$IQ_DIR/log*).

Task summary

- Set up the environment for the bcp and dbisql utilities.
- Run *\$RAP_HOME/rap_bfap_start.sh* to move the incoming data files to the *\$RAP_DATA* and *\$IQ_DIR* staging areas.
- Run *\$RAP_HOME/rap_flap_start.sh* to start the load process.
- Check the load process.
- Run *\$RAP_HOME/rap_flap_stop.sh* and *\$RAP_HOME/rap_bfap_stop.sh* to stop the load process.
- Clear the data from the RAPCache during the testing process.

Load process directories

The following is the directory structure that is created during the install process for the automated load process:

```

$RAP30/Auto_File_Load
  data
    RAPCache_in
      active
      log
    VLDB_in
      active
      log
    archive
    rap_in
    status
    mon_logs
  VLDB_Load_Scripts

```

The values of the environment variables defined in the *rap_flap_env.sh* and *rap_bfap_env.sh* scripts are:

Variable name	Value
\$RAP_HOME	<i>\$RAP30/Auto_File_Load</i>
\$RAP_DATA	<i>\$RAP30/Auto_File_Load/data/RAPCache_in</i>
\$IQ_DIR	<i>\$RAP30/Auto_File_Load/data/VLDB_in</i>
\$RAP_ARCH	<i>\$RAP30/Auto_File_Load/data/archive</i>
\$RAP_IN	<i>\$RAP30/Auto_File_Load/data/rap_in</i>
\$RAP_STAT	<i>\$RAP30/Auto_File_Load/data/status</i>
\$LOAD_IQ_SCRIPT_DIR	<i>\$RAP30/Auto_File_Load/VLDB_Load_Scripts</i>

The two directories *active* and *log* must exist under both the *\$RAP_DATA* and *\$IQ_DIR* directories.

RAP load scripts

This section describes the scripts that control the loading process.

Shell scripts

Shell scripts in the *\$RAP_HOME* directory control the RAPCache loading process.

Some of the shell scripts in the *\$RAP_HOME* directory are used by the RAP File Load Agent and Batch File Agent, which are controlled by the RAP Operations Console. For more information on the Operations Console agents, see Chapter 2, “Administering Operations Console” in the *Risk Analytics Platform Administration Guide*.

Script name	Description
<i>rap_link.sh</i>	Links new data source files in the <i>\$RAP_IN</i> directory to the <i>\$RAP_DATA</i> , <i>\$IQ_DIR</i> , and <i>\$RAP_ARCH</i> directories.
<i>rap_bfap_start.sh</i>	Launches the file moving process. The incoming data files are moved to the <i>\$RAP_DATA</i> and <i>\$IQ_DIR</i> staging areas. Used by the Operations Console Batch File Agent.
<i>rap_flap_start.sh</i>	Launches the loading process. Used by the Operations Console File Load Agent.
<i>rap_bfap_stop.sh</i>	Creates a tickler file in the <i>\$RAP_STAT</i> directory. The file moving process stops, once the tickler file exists.
<i>rap_flap_stop.sh</i>	Creates a tickler file in the <i>\$RAP_STAT</i> directory. The loading process stops, once the tickler file exists.
<i>rap_bfap_env_pl.sh</i>	Sets environment variables. Called by <i>rap_bfap_start.sh</i> and <i>rap_bfap_stop.sh</i> .
<i>rap_flap_env_pl.sh</i>	Sets environment variables. Called by <i>rap_flap_start.sh</i> and <i>rap_flap_stop.sh</i> .
<i>rap_tbl_bcp_pl.sh</i>	Runs the bcp utility to load data into the Primary Cache database and determines completion status. Issues error messages if errors occurred, and logs trail files. Called by <i>rap_tbl_mon_pl.sh</i> .
<i>rap_tbl2_bcp_pl.sh</i>	Runs the bcp utility with arguments to load data into the Secondary Cache database and determines completion status. Issues error messages if errors occurred, and logs trail files. Called by <i>rap_tbl_mon_pl.sh</i> .

Script name	Description
<i>rap_tbl_dbisql_pl.sh</i>	Runs the dbisql utility to load data into the VLDB database and determines completion status. Issues error messages if errors occurred, and logs trail files. Called by <i>rap_tbl_mon_pl.sh</i> .
<i>rap_tbl_mon_pl.sh</i>	Looks for new data source files. This is a middle-level process called by <i>rap_flap_start.sh</i> .
<i>rap_truncate.sh</i>	Truncates the STOCK_QUOTE and STOCK_TRADE tables in the RAPCache. This script is used during the testing process, so that you can reuse the same data files.

SQL scripts

SQL scripts in the `$LOAD_IQ_SCRIPT_DIR` directory load the VLDB database.

Script name	Description
<i>load_iq_stock_quote.sql</i>	Loads the STOCK_QUOTE table
<i>load_iq_stock_trade.sql</i>	Loads the STOCK_TRADE table

Setting the environment variables

Setting up the RAP environment

The *rap_bfap_env.sh* and *rap_flap_env.sh* scripts set the RAP environment variables for the shell. These scripts are called by other load process scripts as necessary.

Variable name	Description
RAP_HOME	Base directory for RAP data loading. All other paths should be relative to this directory.
RAP_DATA	Input directory for loading files into RAPCache. Monitored by the loading process.
RAP_STAT	Used for placing a tickler file to stop the load process.
RAP_IN	General input directory. Files are linked to the \$RAP_DATA, \$IQ_DIR and \$RAP_ARCH directories for loading. Monitored by loading process.
RAP_ARCH	Archive directory. The loading process compares files between this directory and the \$RAP_IN directory to determine whether to link files to other input directories for loading.
IQ_DIR	Input directory for loading files into VLDBServer. Monitored by loading process.
LOAD_IQ_SCRIPT_DIR	Directory where sql scripts used by the dbisql utility during VLDBServer loading are stored.
TICKLER_FILE	Name of the tickler file that stops the loading script.

Variable name	Description
MAX_PTN_QUOTE	Maximum number of partitions for the STOCK_QUOTE table. Set to 3 by the installer.
MAX_PTN_TRADE	Maximum number of partitions for the STOCK_TRADE table. Set to 3 by the installer.
ASE_SERVER	Name of the RAPCache server.
ASE_INTERFACE	Path to the RAPCache interfaces file.
ASE_UID	The RAPCache user ID.
ASE_PWD	The RAPCache user password.
ASE_DB_QUOTE	RAPCache database name for storing STOCK_QUOTE table data.
ASE_DB_TRADE	RAPCache database name for storing STOCK_TRADE table data.
IQ_ENGINE	Name of the VLDBServer. See the note below.
IQ_HOST	Host machine name of the VLDBServer.
IQ_PORT	Port number of the VLDBServer.
IQ_UID	The VLDBServer user ID.
IQ_PWD	The VLDBServer user password.
IQ_DB_QUOTE	VLDBServer database name for storing STOCK_QUOTE table data.
IQ_DB_TRADE	VLDBServer database name for storing STOCK_TRADE table data (should be the same as IQ_DB_QUOTE).

Note If the VLDB Server is installed on a different machine than the RAPCache Server, you must manually set the value of the IQ_ENGINE environment variable.

Setting up the utilities environment

Before starting the load process, you must set up the environment for the bcp and dbisql utilities if you have not already done so.

To set up the environment for the bcp utility used by RAPCache, run the *SYBASE.sh* or *SYBASE.csh* script, which is located in the *\$RAP30/RAPCache* directory.

To set up the environment for the dbisql utility used by VLDBServer, run the *ASIQ-12_6.sh* or *ASIQ-12_6.ch* script, which is located in the *\$RAP30/VLDBServer/ASIQ-12_6* directory.

Starting and stopping the load process

Use the following commands to start and stop the RAP automated load process.

- Run `$RAP_HOME/rap_bfap_start.sh` to launch the file moving process. The incoming data files are moved from the `$RAP_IN` directories to the `$RAP_DATA` and `$IQ_DIR` staging areas. This script is used by the Operations Console Batch File Agent.
- Run `$RAP_HOME/rap_flap_start.sh` to launch the loading process. This script is used by the Operations Console File Load Agent.
- Run `$RAP_HOME/rap_flap_stop.sh` and `$RAP_HOME/rap_bfap_stop.sh` to stop the loading process.

You can also create a tickler file to stop the loading process, instead of using the `rap_flap_stop.sh` script. The file name should be the same as `$TICKLER_FILE`, specified in the `rap_flap_env.sh`. When the loading process stops, the tickler file is removed.

Note Before running the `rap_bfap_start.sh` and `rap_flap_start.sh` scripts to start the load process, set executable privileges on all of the files in the `$RAP_HOME` directory. Otherwise, "Execution permission denied" errors will occur.

If the input directory (`$RAP_IN`) contains the TAQ sample data files, the data in these files is loaded into the RAPCache and RAP VLDBServer databases the first time you start the automated load process.

Loading the TAQ sample data

Sybase RAP includes two sets of sample data for the RAPCache and VLDBServer databases. A smaller set of sample data is loaded during the RAP installation process. Follow these instructions, if you need to reload this smaller set of sample data.

You can either run a script or use the automated load process to load the sample data. Before loading the sample data, unzip the files *STOCK_QUOTE.zip* and *STOCK_TRADE.zip* located in the *\$RAP30/Data/RAPCache/RAP_Input* directory.

Note The installed files *STOCK_QUOTE.csv* and *STOCK_TRADE.csv* contain a subset of the data in the *STOCK_QUOTE.zip* and *STOCK_TRADE.zip* files. If you want to preserve the smaller *STOCK_QUOTE.csv* and *STOCK_TRADE.csv* files, rename these files before unzipping the *STOCK_QUOTE.zip* and *STOCK_TRADE.zip* files.

Using a script to load
RAPCache only

To load the TAQ (Trades and Quotes) sample data only into the RAPCache, run the script

\$RAP30/Data/RAPCache/RAP_Load/load_unix_ASE_TAQ_tables.sh from a command window on the machine that hosts the ASE RAPCache database. This script uses the bulk copy utility *bcp* to load the sample data into the RAPCache *rapt* and *rapq* databases. Errors are written to the files *STOCK_TRADE.err* and *STOCK_QUOTE.err*.

To check the status of the data loads, review the log file for Adaptive Server Enterprise. The default location of the ASE log file is *\$RAP30/RAPCache/ASE-15_0/install/<servername>.log*.

Loading automatically

Use the automated load process to load the TAQ sample data into both the RAPCache and the RAP VLDB databases.

Remove any TAQ sample data already loaded in the RAPCache, or you will receive "Duplicate row" errors. See the section "Clearing the cache" on page 52 for information on how to remove the data from the RAPCache.

If the TAQ sample data is not in the RAPCache, follow these instructions before you start the automated load process.

- 1 Divide the *STOCK_QUOTE.csv* file into smaller batch files using the *split* command:

```
split -l 500000 STOCK_QUOTE.csv STOCK_QUOTE.csv.
```

This command creates several files called *STOCK_QUOTE.csv.nn* (where *nn* is an automatically generated two-character suffix beginning with aa, ab, ac, and so on) that contain 500,000 rows each. This optional step is recommended, to better demonstrate how the automated load process works.

Note To improve load performance, you may need to adjust the number of rows specified in the `split` command to an appropriate value for the physical resources on your system.

- 2 Divide the *STOCK_TRADE.csv* file into smaller batch files using the `split` command:

```
split -l 250000 STOCK_TRADE.csv STOCK_TRADE.csv.
```

This command creates several files called *STOCK_TRADE.csv.nn* (where *nn* is an automatically generated two-character suffix beginning with aa, ab, ac, and so on) that contain 250,000 rows each. This step is optional, but is recommended.

- 3 Copy all of the *STOCK_QUOTE.csv.nn* and *STOCK_TRADE.csv.nn* files into the `$RAP_IN` input directory of the automated load process.
- 4 If the automated load process is not already running, run `$RAP_HOME/rap_bfap_start.sh` and `$RAP_HOME/rap_flap_start.sh` to start the process.

Loading additional TAQ sample data

This optional procedure describes how to load a larger set of sample data into the VLDBServer database. The automated load process that loads RAPCache also loads the VLDBServer database. Sample data that contains millions of additional rows of data for VLDBServer is, however, also available. Running the queries against the sample data loaded during installation returns limited results, so you may want to load this larger set of additional data into the VLDBServer database.

To load the larger set of sample data, first unzip the files *STOCK_QUOTE.zip*, *STOCK_TRADE.zip*, and *STOCK_HISTORY.zip* in the *\$RAP30/Data/VLDB/RAP_Input* directory.

Note The installed files *STOCK_QUOTE.csv*, *STOCK_TRADE.csv*, and *STOCK_HISTORY.csv* contain a subset of the data in the *STOCK_QUOTE.zip*, *STOCK_TRADE.zip*, and *STOCK_HISTORY.zip* files. To preserve the smaller *STOCK_QUOTE.csv*, *STOCK_TRADE.csv*, and *STOCK_HISTORY.csv* files, rename these files before unzipping the *STOCK_QUOTE.zip*, *STOCK_TRADE.zip*, and *STOCK_HISTORY.zip* files.

To run the script that loads the larger set of TAQ sample data into VLDBServer, you must have Sybase IQ 12.6 client tools installed on your system.

Run the *load_unix_IQ_TAQ_tables.sh* script on the machine that hosts the Sybase IQ database. This script truncates existing data and loads the larger set of TAQ sample data into the VLDB database.

To load the larger set of sample data for the *STOCK_HISTORY* table, edit the file *load_unix_IQ_TAQ_tables.tst*. Append the following lines, replacing *<RAP30_installation_directory>* with the path name of your RAP 3.0 installation directory:

```
truncate table DBA.STOCK_HISTORY;
commit;
load table DBA.STOCK_HISTORY
(
INSTRUMENT_ID    null (blanks, 'NULL') ,
TRADE_DATE       null (blanks, 'NULL') ,
TRADING_SYMBOL  null (blanks, 'NULL') ,
OPEN_PRICE       null (blanks, 'NULL') ,
CLOSE_PRICE      null (blanks, 'NULL') ,
LOW_PRICE        null (blanks, 'NULL') ,
HIGH_PRICE       null (blanks, 'NULL') ,
VOLUME           '\x0a'
)
from
'<RAP30_installation_directory>/Data/VLDB/RAP_Input/
STOCK_HISTORY.csv'
quotes off
escapes off
preview on;
```

```
commit;
```

Note To preserve the TAQ data already in VLDB, edit the file *load_unix_IQ_TAQ_tables.tst* and comment out the following commands:

```
truncate table DBA.STOCK_QUOTE;
commit;
truncate table DBA.STOCK_TRADE;
commit;
```

load_unix_IQ_TAQ_tables.tst is called by *load_unix_IQ_TAQ_tables.sh*.

To check the status of the TAQ data load, review the log file for the VLDBServer database, *<database_name>.iqmsg*. The default location of the log file is the directory where the VLDBServer database file (*<database_name>.db*) is located.

Checking the load process

If a loading error occurs, the RAP load process sends a message to *stdout* and creates an error file.

RAPCache load

If an error occurs loading RAPCache, the following message appears in *stdout*:

```
ERROR in BCP operation
```

An *.err* error file is created in the *\$RAP_DATA/log* directory. Check this file for error messages to determine the cause of the RAPCache load failure.

VLDBServer load

If an error occurs loading the VLDBServer IQ database, the following message is displayed in *stdout*:

```
ERROR in DBISQL operation
```

An *.err* error file is created in the *\$IQ_DIR/log* directory. Check this file for error messages to determine the cause of the VLDBServer database load failure.

File Load and Batch File Agents

The RAP File Load Agent and Batch File Agent, which are controlled by the RAP Operations Console, also monitor and report on the data load processes. For more information on the Operations Console agents, see Chapter 2, “Administering Operations Console” in the *Risk Analytics Platform Administration Guide*.

Clearing the cache

Clear the data from the RAPCache on a regular schedule, as the data is permanently stored in the VLDB database and is redundant in the cache. The script *rap_truncate.sh*, located in the *\$RAP_HOME* directory, clears the data by truncating the STOCK_TRADE and STOCK_QUOTE tables in the RAPCache.

The script *rap_truncate.sh* can be run manually or called from your own scheduling program. This script depends on the environment variables set by *rap_flap_env.sh* as described in “Setting the environment variables” on page 45.

If errors occur during the truncate process, a message is displayed and error messages are written to the file *truncate.err*.

Note The RAP Operations Console also has an option to delete data from the RAPCache database using the Data Synch agent. See the *Risk Analytics Platform Administration Guide* for more information on the Operations Console and the Data Synch agent.

Adding new instruments

Two stored procedures for adding new instruments to the RAPCache *rapt* and *rapq* databases are included in the RAPCache model database. The names of the stored procedures are *proc_GetInstruments* and *proc_GetInstrumentIdBySymbol*. These procedures are used by the third-party software that loads streaming data into the *rapq* and *rapt* RAPCache databases.

The scripts that create the new instrument stored procedures are installed in the *\$RAP30/Model/RAPCache* directory.

Sample ASE Configuration Settings

About this Appendix

This appendix contains information about the settings of the options in the configuration file for the RAPCache Adaptive Server Enterprise server that are not default values.

Contents

Topic	Page
Introduction	53
General information	54
Physical Resources	58
Cache memory configuration example	65

Introduction

A sample Adaptive Server Enterprise configuration file *RAP_RAPCache_Server.cfg* is included with Risk Analytics Platform in the installation folder *\$RAP30/Scripts/RAPCache*. This appendix contains information about the settings of the options in the configuration file for the RAPCache Adaptive Server Enterprise server that are not default values.

An example of configuring memory for cache is also included at the end of this appendix. See the section “Cache memory configuration example” on page 65.

Configuration parameters are grouped according to the area of Adaptive Server Enterprise behavior they affect. This grouping makes it easier to identify all parameters that you might need to tune to improve a particular area of Adaptive Server Enterprise performance.

Note If a group of configuration parameters is not listed in this appendix, then the options in that group all have the default values.

For more information on Adaptive Server Enterprise configuration parameters, refer to the *Adaptive Server Enterprise System Administration Guide* or consult your ASE Administrator.

General information

Parameters in this group are not related to any particular area of Adaptive Server behavior:

- “Backup/Recovery” on page 54
- “Named Cache:c_log” on page 55
- “Named Cache:default data cache” on page 56
- “Disk I/O” on page 57
- “Unicode” on page 57
- “Network Communication” on page 58

Backup/Recovery

Backup/Recovery parameters configure Adaptive Server for backing up and recovering data.

```
[Backup/Recovery]
recovery interval in minutes = 32767
print recovery information = DEFAULT
tape retention in days = DEFAULT
max concurrently recovered db = DEFAULT
number of checkpoint tasks = DEFAULT
```

recovery interval in minutes = 32767 This option controls how often checkpoint flushes dirty (modified) pages to disk and is set to the maximum value to delay checkpointing for as long as possible.

Named Cache:c_log

Parameters in this group specify properties in a named cache. Each configured data cache on the server has this block of information in the configuration file. In this case, `c_log` is the configuration file entry for a separate named cache for log devices.

Note The Named Cache:c_log section, including the related [4K I/O Buffer Pool] and [16K I/O Buffer Pool] sections, is not included in the default ASE configuration file and must be added to your configuration file. Use a text editor to cut this section from the *RAP_ASE_Server.cfg* configuration file included with Risk Analytics Platform and paste it into your own configuration file.

```
[Named Cache:c_log]
  cache size = 800M
  cache status = log only
  cache status = HK ignore cache
  cache replacement policy = DEFAULT
  local cache partition number = 1
```

cache size = 800M The value of 800M is a good size for the RAP sample database. This option is tunable.

cache status = log only This named cache is used only for the log.

cache status = HK ignore cache Setting this option equal to HK ignore cache prevents the housekeeper process from washing the log cache and avoids contention between the housekeeper and cache manager spinlock.

local cache partition number = 1 The value of 1 is the standard recommendation for this option and specifies 1 partition for the cache.

```
[4K I/O Buffer Pool]
  pool size = DEFAULT
  wash size = DEFAULT
  local async prefetch limit = 0
```

local async prefetch limit = 0 The value of 0 disables prefetch activities, which are not needed.

```
[16K I/O Buffer Pool]
  pool size = 780M
  wash size = DEFAULT
  local async prefetch limit = 0
```

pool size = 780M The value of 780M is a starting point. The remainder of the 800M is allocated to the 4K pool. When changing cache size, give the most to the 16K buffer pool, because the log I/O size is set to 16K for the databases using this pool. This option is tunable and should be increased if necessary.

local async prefetch limit = 0 The value of 0 disables prefetch activities, which are not needed.

Named Cache:default data cache

Parameters in this group specify properties in a named cache. Each configured data cache on the server has this block of information in the configuration file. In this case, this group contains the configuration file entries for a separate default data cache.

```
[Named Cache:default data cache]
  cache size = 1700M
  cache status = default data cache
  cache replacement policy = DEFAULT
  local cache partition number = 4
```

cache size = 1700M The value of 1700M is a starting point for the RAP sample database. The remainder of the 1700M goes to the 4K buffer pool. When changing cache size, give the least to the 16K pool. This option is tunable and can be increased to improve the speed of queries.

cache status = default data cache This named cache is the default data cache.

local cache partition number = 4 The value of 4 is for a four-CPU machine. This value can be greater for a greater number of CPUs. Increasing the number of cache partitions for a cache can reduce spinlock contention and improve performance.

```
[4K I/O Buffer Pool]
  pool size = DEFAULT
  wash size = DEFAULT
  local async prefetch limit = DEFAULT

[16K I/O Buffer Pool]
  pool size = 100M
  wash size = DEFAULT
  local async prefetch limit = DEFAULT
```

pool size = 100M To avoid using a large buffer pool and to have fewer large I/Os, the value of 100M for buffer pool size is relatively small. The remainder of the 1700M is allocated to the 4K buffer pool. When changing cache size, allocate the least amount to the 16K pool.

Disk I/O

Parameters in this group configure the disk I/O of Adaptive Server.

```
[Disk I/O]
disk i/o structures = 2048
number of large i/o buffers = DEFAULT
page utilization percent = DEFAULT
number of devices = 50
disable disk mirroring = DEFAULT

allow sql server async i/o = 0
```

disk i/o structures = 2048 This option specifies the initial number of disk I/O control blocks ASE allocates at start-up and is tunable. Start with a value of 2048.

number of devices = 50 This option controls the number of database devices ASE can use. This value should be greater than or equal to the number of devices specified in disk init commands.

allow sql server async i/o = 0 Setting this option to 0 disables asynchronous disk I/O for ASE and prevents prefetching of pages unnecessarily.

Unicode

Parameters in this group configure Unicode features:

```
[Unicode]
enable unicode normalization = 0
enable surrogate processing = 0
enable unicode conversions = 0
size of unilib cache = DEFAULT
```

The Unicode options are set equal to 0 to disable Unicode functions, as the processing of Unicode characters is not needed.

Network Communication

Parameters in this group are used to configure communication between Adaptive Server and remote servers, and between Adaptive Server and client programs.

```
[Network Communication]
default network packet size = 32768
max network packet size = 32768
remote server pre-read packets = DEFAULT
number of remote connections = DEFAULT
number of remote logins = DEFAULT
number of remote sites = DEFAULT
max number network listeners = DEFAULT
tcp no delay = DEFAULT
send doneinproc tokens = DEFAULT
allow sendmsg = DEFAULT
syb_sendmsg port number = DEFAULT
allow remote access = DEFAULT
```

default network packet size and max network packet size = 32768 These options control the amount of data in one packet. This option is tunable. Start with 32K and tune for performance; also try 16K and 64K. Keep in mind that the amount of memory specified is allocated to each connection.

Physical Resources

Parameters in this group determine the optimal setup for physical resources at the Adaptive Server server and operating system levels.

- “Physical Memory” on page 59
- “Processors” on page 59
- “SQL Server Administration” on page 60
- “User Environment” on page 61
- “Lock Manager” on page 62
- “Component Integration Services” on page 63
- “DTM Administration” on page 64

Physical Memory

Parameters in this group configure your machine's physical memory resources.

```
[Physical Memory]
max memory = 2000000
additional network memory = DEFAULT
shared memory starting address = DEFAULT
allocate max shared memory = DEFAULT
dynamic allocation on demand = DEFAULT
lock shared memory = 1
heap memory per user = DEFAULT
```

max memory = 2000000 The value of this option is the maximum amount of total physical memory (in 2K pages) that ASE can use and is platform dependent. Set this value high, even though the server may not use all of this memory.

lock shared memory = 1 Setting this option to 1 locks memory, so that ASE does not swap pages to disk. This reduces disk reads, which are expensive. This parameter is platform dependent and works only on platforms that support locking shared memory.

Processors

Parameters in this group configure processors in an SMP environment.

```
[Processors]
max online engines = 8
number of engines at startup = 3
statement cache size = DEFAULT
```

max online engines = 8 This tunable option specifies the maximum number of ASE engines that can be online at any one time in an SMP environment. The value should be set equal to the number of logical CPUs or to the number of logical CPUs - 1.

number of engines at startup = 3 This tunable option is used during start-up to set the number of engines brought online. The value must be less than or equal to the number of CPUs on your machine and not greater than the specified value of max online engines. You need CPUs for both the client and the server, and you should allow at least 1 CPU for BCP processes. You need to allow 1 CPU for other processes and the operating system as well.

SQL Server Administration

Parameters in this group are related to general Adaptive Server administration:

```
[SQL Server Administration]
procedure cache size = 32000
default database size = DEFAULT
identity burning set factor = DEFAULT
allow nested triggers = DEFAULT
allow updates to system tables = DEFAULT
default fill factor percent = DEFAULT
default exp_row_size percent = DEFAULT
number of mailboxes = DEFAULT
number of messages = DEFAULT
number of alarms = DEFAULT
number of pre-allocated extents = DEFAULT
event buffers per engine = DEFAULT
cpu accounting flush interval = 2147483647
i/o accounting flush interval = 2147483647
sql server clock tick length = DEFAULT
runnable process search count = DEFAULT
i/o polling process count = DEFAULT
time slice = DEFAULT
cpu grace time = DEFAULT
number of sort buffers = DEFAULT
size of auto identity column = DEFAULT
identity grab size = DEFAULT
housekeeper free write percent = 5
enable housekeeper GC = 0
sysstatistics flush interval = DEFAULT
allow resource limits = DEFAULT
number of aux scan descriptors = DEFAULT
SQL Perfmon Integration = DEFAULT
license information = DEFAULT
text prefetch size = DEFAULT
enable HA = DEFAULT
i/o batch size = DEFAULT
enable semantic partitioning = 0
enable xml = DEFAULT
enable webservices = DEFAULT
enable job scheduler = DEFAULT
job scheduler tasks = DEFAULT
job scheduler interval = DEFAULT
percent database for history = DEFAULT
percent history free = DEFAULT
percent database for output = DEFAULT
percent output free = DEFAULT
```

maximum job output = DEFAULT

procedure cache size = 32000 This option specifies the size of the procedure cache in 2K pages. The procedure cache is used while running stored procedures and also to compile queries while creating stored procedures. A value lower than 32000 (such as 16000) might be sufficient, depending on the performance of queries.

cpu accounting flush interval = 2147483647 This option specifies the amount of time in machine clock ticks that ASE waits before flushing to disk CPU usage statistics for each user. This option is set to the maximum value to reduce the number of times statistical information is written to disk.

i/o accounting flush interval = 2147483647 This option specifies the amount of time in machine clock ticks that ASE waits before flushing to disk I/O statistics for each user. This option is set to the maximum value to reduce the number of times statistical information is written to disk.

housekeeper free write percent = 5 This option specifies the maximum percentage by which the housekeeper task can increase database writes. The default value is 1, so setting this option to 5 allows the housekeeper to do more writes when idle.

enable housekeeper GC = 0 Setting this option to 0 disables garbage collection by the housekeeper. Garbage collection is not necessary, since there are mostly inserts and few deletes.

enable semantic partitioning = 0 Disables partitioning other than round-robin (for example, range partitioning) in ASE.

User Environment

Parameters in this group configure user environments.

```
[User Environment]
number of user connections = 40
stack size = DEFAULT
stack guard size = DEFAULT
permission cache entries = DEFAULT
user log cache size = 65536
user log cache spinlock ratio = 1
max native threads per engine = DEFAULT
messaging memory = DEFAULT
enable real time messaging = DEFAULT
histogram tuning factor = DEFAULT
rtm thread idle wait period = DEFAULT
```

number of user connections = 40 This option sets the maximum number of user connections that can be connected to Adaptive Server at the same time. The default is 25.

user log cache size = 65536 This option specifies the size in bytes of each user's log cache. The size of this private log cache is determined by the logical page size of the server. A large value of 65536 (65K) avoids having to write to the real log until this cache is full, which is more efficient when performing many loads.

user log cache spinlock ratio = 1 For ASE with multiple engines, this option specifies the ratio of user log caches per user log cache spinlock. There is one user log cache for each configured user connection. A value of 1 reduces contention and locking of the real log.

Lock Manager

Parameters in this group configure locks:

```
[Lock Manager]
  number of locks = 100000
  deadlock checking period = 5000
  lock spinlock ratio = 20
  lock address spinlock ratio = DEFAULT
  lock table spinlock ratio = DEFAULT
  lock hashtable size = 131072
  lock scheme = DEFAULT
  lock wait period = DEFAULT
  read committed with lock = DEFAULT
  print deadlock information = DEFAULT
  deadlock retries = DEFAULT
  page lock promotion HWM = DEFAULT
  page lock promotion LWM = DEFAULT
  page lock promotion PCT = DEFAULT
  row lock promotion HWM = 20000
  row lock promotion LWM = 20000
  row lock promotion PCT = DEFAULT
```

number of locks = 100000 This option sets the total number of available locks for all users on Adaptive Server. The value of 100000 accommodates six streams of BCPs with a batch size of 10,000.

deadlock checking period = 5000 This option specifies the minimum amount of time in milliseconds before Adaptive Server initiates a deadlock check for a process that is waiting on a lock to be released. Since few deadlocks are expected, the value of 5000 reduces the cost of the overhead of deadlock checking.

lock spinlock ratio = 20 The value of 20 for this option avoids spinlock contention. This option is tunable.

lock hashtable size = 131072 This option specifies the number of hash buckets in the lock hash table. A value of 131072 reduces lock contention. This parameter is tunable.

row lock promotion HWM = 20000 This option specifies the high water mark for row lock promotion. A high value of 20000 disables lock promotion, because the data is loaded in batches of 10000. This parameter should be tuned according to the batch size of the data load.

row lock promotion LWM = 20000 This option specifies the low water mark for row lock promotion. A high value of 20000 essentially disables lock promotion.

Component Integration Services

Parameters in this group configure Adaptive Server for Component Integration Services:

```
[Component Integration Services]
enable cis = 0
cis connect timeout = DEFAULT
cis bulk insert batch size = DEFAULT
max cis remote connections = DEFAULT
cis packet size = DEFAULT
cis cursor rows = DEFAULT
enable snmp = DEFAULT
enable file access = DEFAULT
cis bulk insert array size = DEFAULT
enable full-text search = DEFAULT
cis rpc handling = DEFAULT
```

enable cis = 0 The value of 0 for this option disables Component Integration Service.

DTM Administration

Parameters in this group configure Distributed Transaction Management (DTM) facilities:

```
[DTM Administration]
enable DTM = DEFAULT
enable xact coordination = 0
xact coordination interval = DEFAULT
number of dtx participants = DEFAULT
strict dtm enforcement = DEFAULT
txn to pss ratio = DEFAULT
dtm lock timeout period = DEFAULT
dtm detach timeout period = DEFAULT
```

enable xact coordination = 0 This option is set to 0 to disable Adaptive Server transaction coordination services and prevent unnecessary processing. (The Distributed Transaction Management feature is disabled by default.)

Cache memory configuration example

In this configuration example, 4GB of memory has been allocated for the RAPCache ASE server.

- 1 Make the configuration changes for items other than cache that use memory.
- 2 Make all configuration changes except the changes for `c_log` and the default data cache.
- 3 Shut down and restart the server.
- 4 Run `sp_configure memory`

The output is similar to this:

Parameter Name	Default	Memory Used	Config Value	Run Value	Unit	Type
additional network memory	0	0	0	0	bytes	dynamic
allocate max shared memory	0	0	0	0	switch	dynamic
heap memory per user	4096	0	4096	4096	bytes	dynamic
locked shared memory	0	0	1	1	switch	static
max memory	53248	4000000	2000000	2000000	memory pages(2k)	dynamic
memory alignment boundary	4096	0	4096	4096	bytes	static
memory per worker process	1024	4	1024	1024	bytes	dynamic
messaging memory	400	0	400	400	memory pages(2k)	dynamic
shared memory starting address	0	0	0	0	not applicable	static
total logical memory	53248	174190	87095	87092	memory pages(2k)	read-only
total physical memory	0	174192	0	87096	memory pages(2k)	read-only

`sp_configure memory` returns this information: "An additional 3825816 K bytes of memory is available for reconfiguration. This is the difference between 'max memory' and 'total logical memory'."

3825816K bytes of memory are available and some of this will be used for memory for cache. The rest of the memory is reserved for future changes. The objective is to hold as much data as possible in memory. 800M is allocated for the log and 1700M for the data cache.

- 1 Make the configuration changes for `c_log` and the default data cache.
- 2 Shut down and restart the server.
- 3 Run `sp_configure memory`

The output is similar to this:

Parameter Name	Default	Memory Used	Config Value	Run Value	Unit	Type
additional network memory	0	0	0	0	bytes	dynamic
allocate max shared memory	0	0	0	0	switch	dynamic
heap memory per user	4096	0	4096	4096	bytes	dynamic
lock shared memory	0	0	1	1	switch	static
max memory	53248	4000000	2000000	2000000	memory pages(2k)	dynamic
memory alignment boundary	4096	0	4096	4096	bytes	static
memory per worker process	1024	4	1024	1024	bytes	dynamic
messaging memory	400	0	400	400	memory pages(2k)	dynamic
shared memory starting address	0	0	0	0	not applicable	static
total logical memory	53248	2873302	1436651	1436649	memory pages(2k)	read-only
total physical memory	0	2873304	0	1436652	memory pages(2k)	read-only

`sp_configure memory` returns this information: "An additional 1126702 K bytes of memory is available for reconfiguration. This is the difference between 'max memory' and 'total logical memory'."

To determine how much memory to allocate to data cache for your application, run the following stored procedure commands:

```
sp_estspace STOCK_TRADE, <nbr of rows>  
sp_estspace STOCK_QUOTE, <nbr of rows>
```

Data Replication

About this Appendix

This appendix tells you how to use Sybase Replication Server 12.6 to replicate data from Adaptive Server Enterprise to the Sybase IQ VLDBServer.

Contents

Topic	Page
Overview	67
Source database setup	69
Replication Server Configuration	70
Target database configuration	71
Replicating test data	77
Troubleshooting	78

Overview

Although Replication Server 12.6 is not packaged as part of the Sybase Risk Analytics Platform, you can use Replication Server to load other reference data from a source database to the RAP VLDBServer database.

Instructions in this appendix assume that you purchased a copy separately and are familiar with the Replication Server 12.6 environment.

Instructions also assume that the source database is Adaptive Server Enterprise.

Replication environment

- Sybase Adaptive Server Enterprise version 15.0 (ASE) serves as the source database and the Replication Server System Database (RSSD).
- Sybase IQ version 12.6 serves as the VLDBServer target database.

- Sybase Replication Server version 12.6 (RepServer) captures and processes transactions on the source database and transfers them to the Sybase IQ VLDBServer target database.
- Sybase Replication Server Manager version 12.6 (RSM) manages the replication environment and is used to configure replication between the ASE source and Sybase IQ VLDBServer target databases.
- Sybase Replication Manager version 12.6 (RepMgr) manages the Multiple Site Availability database replication. Multiple Site Availability enables the replication of entire databases and its table DDL.

Sybase Central

With this configuration, there are three different versions of Sybase Central available because different products install their own versions. The information below should help clarify the use of each version of Sybase Central:

- 1 To manage the IQ server, start Sybase Central Java Edition version 4.3 under Adaptive Server IQ 12.6 from the Windows Start menu.
- 2 To manage the RepServer using RSM, start Sybase Central version 3.2 under Sybase from the Windows Start menu. RSM enables the creation of connections, replication definitions, subscriptions, and publications for configuring a replication environment.
- 3 To manage Multiple Site Availability within RepServer using RepMgr, start Sybase Central Java Edition version 4.1.1 under Sybase, from the Windows Start menu. RepMgr enables the configuration of replication for entire databases, including table DDL, which is called Multiple Site Availability. Additionally, this version of Sybase Central manages the ASE server and databases.

Source database setup

Before configuring the Replication Server, you must create and configure a source database that contains the end-of-day historical market data that is replicated to the RAP VLDBServer Sybase IQ target database. In these instructions, Sybase Adaptive Server Enterprise serves as the source database and the Replication Server System Database (RSSD), although another RDBMS can be used.

Adaptive Server Enterprise 15.0 is included in the Risk Analytics Platform package. See the *Release Bulletin Risk Analytics Platform* for the specific version requirements of Adaptive Server Enterprise.

Install Adaptive Server Enterprise as described in this manual. Sybase strongly recommends that you always check the online support Web site for the latest ASE software updates (ESD or EBF). See “Sybase EBFs and software maintenance” on page ix for more information.

Preliminary steps

You must have appropriate administrative tools installed on your machine and authority to perform the tasks outlined in this section.

You need to create a user with DBA authority and an empty database that identifies that user as the database owner. You should also establish an appropriate ODBC DSN or other connection to the database.

Creating the schema for the ASE source database

PowerDesigner includes all the resources you need to generate a set of data definition language (DDL) statements in a SQL script directly from the data model. You can run this script to generate a schema for your ASE database. The ASE database schema must match the Sybase IQ schema, as Sybase IQ is the Replication Server target database in Risk Analytics Platform.

To create the SQL script with DDL statements for the ASE database schema, follow the steps in “RAPCache database” in Chapter 4, “Generating DDL Scripts” of the *Risk Analytics User's Guide*. Execute the generated script to create the ASE database tables and other database objects.

You can also create the ASE source database using the DDL script `$RAP_HOME/Model/RAPCache/RAP_Table.sql` which is provided with RAP.

Note Risk Analytics Platform does not include indexes for an Adaptive Server Enterprise 15.0 source database.

Replication Server Configuration

Risk Analytics Platform relies on Sybase Replication Server 12.6 to replicate transactions between the Adaptive Server Enterprise source database and the RAP VLDBServer Sybase IQ target database.

Configuration tasks

This section summarizes the tasks you use to install and configure Replication Server. For specific installation instructions, see the *Replication Server Installation and Configuration Guide* for your platform.

- 1 Install and configure Replication Server version 12.6 according to the product documentation.
- 2 Install the RepServer Client for Windows, which installs the required files for Replication Server Manager (RSM). The RSM Client is a plug-in to Sybase Central and is available only on Microsoft Windows platforms.
 - If Setup asks you whether you want to overwrite any Open Client/Open Server DLLs older than those currently installed on your system, choose No.
 - Do not use the `SAMPLE_RS` that can be configured during the RepServer installation. `SAMPLE_RS` uses an ASA database, which is not currently supported by RSM.
- 3 Use `rs_init` to create a RepServer instance. `rs_init` is a utility used to configure Replication Server and add databases to a replication system.
- 4 From the Windows Services management console, start the new RepServer service. This service is named `Sybase REP_ <server_name>`, where `<server_name>` is the name you gave to your server.

RSM configuration

This section summarizes the tasks you use to configure Replication Server Manager. For specific installation instructions, see the *Replication Server 12.6 Configuration Guide* for your platform.

- 1 Configure Replication Server Manager according to the product documentation. For specific configuration instructions, see the *Configuration Guide Replication Server 12.6* for your platform.
- 2 Verify that the `%SYBASE%\ini` directory includes a `sql.ini` file that contains an entry for each server you want RSM to manage, as well as an entry for the RSM Server.

Use `dsedit` to create an RSM server entry and an entry for the Sybase IQ server in the `sql.ini` file. Define the RSM port as any available port number; the Sybase IQ port should be the actual port that the server is using. See the *Replication Server Administration Guide* and the *Replication Server Installation Guide* for Windows for more information.

- 3 Run RSM Server Setup from the Sybase group on the Windows Programs menu. This creates a Windows service for RSM named `Sybase RSMServer_<server name>`, where `<server name>` is the RSM server name. Enter a valid Windows login and password. This user ID will be used to stop and start the new RSM service.
- 4 From the Windows Services management console, start the new RSM service. This service is named `Sybase RSMServer_<server name>`, where `<server name>` is the name you used to define the RSM server.

Target database configuration

Replication Server uses a data server maintenance user for each local data server database that contains replicate tables. This allows Replication Server to maintain and update the replicate tables in the database.

The Database Owner or the System Administrator must create and grant appropriate privileges to the maintenance user in the Sybase IQ server for the RepServer connection definition.

Creating a maintenance user

This procedure creates a maintenance user called *ruser* in the Sybase IQ VLDBServer database. All procedures are performed in Interactive SQL. Separate start-up procedures for Windows and UNIX are provided below.

- 1 To start dbisql, do one of the following:
 - In Windows, click Start | Programs | Sybase | Adaptive Server IQ 12.6 | Interactive SQL.
 - In UNIX, open a command shell and type the following command at a system prompt: dbisql.
- 2 Run the following commands:

```
grant connect to ruser identified by ruser
grant membership in group rs_systabgroup to ruser
```

Granting maintenance access to system tables and procedures

The Database Owner or the System Administrator must grant to the maintenance user the permissions required to insert, delete, and update rows in replicated tables and to execute replicated stored procedures.

Procedure permissions

Grant execute permissions for each of the following procedures:

```
rs_getlastcommit
rs_initialize_threads
rs_maker
rs_update_lastcommit
rs_update_threads
```

All procedures are performed in Sybase Central. Follow these instructions for Windows or UNIX.

- 1 To start Sybase Central, do this:
 - In Windows, click Start | Programs | Sybase | Adaptive Server IQ 12.6 | Sybase Central Java Edition.
 - In UNIX, open a command shell, change directory to `$$SYBASE/sybccentral`, and type: `% scjview`

If you have added `$SYBASE/ASIQ-12_6/bin` or `$SYBASE/bin` to your path as instructed at the end of the installation, you can issue the `scjview` command from any directory.

- 2 In Sybase Central, right-click on Sybase IQ and choose Connect.
- 3 Add the appropriate connection information and click OK.
- 4 Expand Sybase IQ | *Sybase IQ* <target database> | Procedures & Functions.
- 5 Right-click on a procedure, choose Properties and click Permissions | Grant.

Note If these procedures and functions are not listed, do the following in Sybase Central:

- Right-click on the database name.
 - Select Filter Objects by Owner.
 - Check `rs_systabgroup` and click OK.
 - Click View | Refresh Folder.
 - Expand Procedures & Functions and continue with this step.
-

- 6 On the Grant Permissions dialog, choose `rs_systabgroup` and click OK.
The `rs_systabgroup` group on the properties dialog appears with a check mark in the Execute column. Click OK.
- 7 Repeat steps 5 – 6 for each procedure in the list.

Table permissions

This procedure grants Insert, Update, and Delete Permissions for the `rs_lastcommit` and `rs_threads` tables in the Sybase IQ database. `rs_lastcommit` commits a transaction; `rs_threads` detects deadlocks and performs transaction serialization between parallel DSI threads.

- 1 Expand Sybase IQ | <target database> | Tables.
- 2 Right-click on the appropriate table, choose Properties, and click Permissions | Grant.
- 3 On the Grant Permissions dialog, choose `rs_systabgroup` and click OK.
- 4 On the Permissions tab, choose Select, Insert, Update, and Delete permissions. Click OK.

- 5 Repeat steps 2 – 4 for each table in the list.

Replication table permissions

This procedure grants all permissions on the Sybase IQ target replication tables to either the maintenance user or rs_systabgroup group.

- 1 Expand Sybase IQ | *<target database>* | Tables.
- 2 Right-click the replicate table, choose Properties and click Permissions | Grant.
- 3 On the Grant Permissions dialog, choose rsuser (maintenance user) *or* rs_systabgroup and click OK.
- 4 Close Sybase Central.

Replication Server Manager

Replication Server Manager (RSM) allows you to manage, monitor, and troubleshoot most replication system components, including primary and replicate data servers in Replication Server. Replication Server Manager is a client application and is available only in Windows.

Defining Servers within RSM

This section tells you how to use Replication Server Manager to add the following servers to the RSM *<server name>* domain:

- ASE (source database) as Adaptive Server Enterprise (ASE) Server Type
- RepServer as Replication Server Server Type
- IQ (target database) as Adaptive Server Anywhere (ASA) Server Type

This procedure requires the Replication Server 12.5 Sybase Central 3.2 plug-in for Windows.

- 1 Click Start | Programs | Sybase | Sybase Central.
- 2 Click Connect, and log in to the RSMServer.
- 3 Click the RSMServer name | Add Server. The Add Server wizard starts.
- 4 Choose the appropriate server type, then follow the instructions on the screen for the type of server you chose.
- 5 Click Finish.

- 6 Repeat steps 3 – 5 until you add each server in the list.

Creating source and target connections

This procedure tells you how to define the Sybase IQ <target database> as the target connection and the ASE source as a primary and target connection. You must perform this procedure twice—once for the target Sybase IQ <target database> and once for the ASE database.

Replication Server communicates directly with all the data servers by using Open Client connectivity. Replication Server database connections are configured as either replicate-only, or primary-and-replicate (bidirectional) in a Sybase-only replication system.

- 1 Click Rep Server | Connections | Add Database Connection. The Add Database Connection Wizard starts.
- 2 Choose the server name and database for the new connection. Click Next.
- 3 Enter the user name and password for the maintenance user. Click Next.
- 4 Identify the connection type. Choose one of these options, then click Next.
 - Replicate Only (use this option for Sybase IQ)
 - Primary and Replicate (use this option for ASE)
- 5 Review the wizard summary. If all options are correct, click Next.
- 6 Repeat these steps until you define the Sybase IQ <target database> and the ASE database.

Notes

- The maintenance user in the connection definition must match the user defined in Sybase IQ <target database>. For example, if you create a maintenance user called ruser, you must create a maintenance user called ruser as part of the connection definition.
 - The maintenance user for the ASE connection must differ from the IQ maintenance user for replication to work properly.
 - See “Troubleshooting” on page 78 for information on problems that you may encounter after creating the IQ connection.
-

Creating replication definitions

Replication definitions identify specific tables for replication and, in some cases, specify a subset of the columns you want to replicate. This procedure shows you how to create one or more replication definitions for an ASE source table.

- 1 Expand ASE <server name> | <database name> | Tables.
- 2 Click on the table where you want to create a replication definition, and double-click Add Replication Definition. A property sheet for the replication properties appears.
- 3 Click the Columns tab and define a key for the source table.

To define a key, click on the column icon for the column you want to use as the key. The image changes to a key icon. If the scripts were created in PowerDesigner with the primary key option, the key should already be defined.

- 4 Click the Advanced tab and add the target table owner, if necessary. If you created the Sybase IQ database with the default owner, this value should be DBA.
- 5 Click OK.

Note

When you create a replication definition, RSM may display an error message similar to this message: "Replication definition <table name> requires at least one column in the primary key."

Every replication definition requires that one column be defined as a key. Within the replication definition, select the Columns tab and click on the column name that should be defined as the key for the table. The icon should change to the image of a key.

Creating subscriptions within RSM

Subscriptions identify the replication definition or publication to which you are subscribing, the source and destination databases and data servers, and the materialization method by which the initial information is to be copied. This procedure tells you how to add a subscription to the Replication Definition you created in the previous section:

- 1 Expand Replication Definition | *Replication Definition <name>* | Table Subscriptions.

- 2 Right-click Add Subscription and choose Open. A property sheet for the new subscription appears.
- 3 In the Name box, type a name for the new subscription.
- 4 In Replication Information Group, click Browse and choose the target database server and database.
- 5 Choose the appropriate Owner ID and password. For a replicate database on ASE, the owner is typically defined as sa with no password.
- 6 Choose the appropriate materialization Method.
- 7 Click OK. RSM adds the new subscription to the Subscription folder.

Replicating test data

The following commands assume that you set up the MARKET_INDEX table for replication:

- 1 Using isql from a command prompt, insert data into the ASE source table:

```
insert into MARKET_INDEX values
    (10, 'Test1', 'Test1', 'Test1', '2006-09-18')
go
insert into MARKET_INDEX values
    (20, 'Test2', 'Test2', 'Test2', '2006-09-18')
go
insert into MARKET_INDEX values
    (30, 'Test3', 'Test3', 'Test3', '2006-09-18')
go
```

- 2 Log in to the target IQ database using Interactive SQL (dbisql) to select data from the target table using the following SQL statement. Verify that the data was replicated properly:

```
commit;
select * from MARKET_INDEX
```

- 3 Using isql from a command prompt, insert data into the ASE source table:

```
update MARKET_INDEX
    set INDEX_NAME='Update Test'
    where MARKET_INDEX_ID=20
go
delete MARKET_INDEX
```

```
where MARKET_INDEX_ID=30
go
```

- 4 Use Interactive SQL for IQ (dbisql) and select data from the target table using the following SQL statement. Verify that the data was replicated properly:

```
commit;
select * from MARKET_INDEX
```

Troubleshooting

This section contains some basic troubleshooting information.

Data fails to replicate

If no data was replicated, check the connection to the Sybase IQ *<target database>*, then check the RSM, RepServer, ASE and IQ logs for error messages.

Sybase also recommends the *Sybase Replication Server Troubleshooting Guide* and the *Multi-Vendor Replication with Sybase Replication Technology* white paper for more information on troubleshooting replication problems.

While trying to replicate data from an ASE table to an IQ table owned by DBA, the connection to IQ comes DOWN

Symptom After inserting data into the source ASE table, the connection to the Sybase IQ *<target database>* database comes down and replication does not occur. The RS log reports the following error:

```
T. 2006/09/18 11:06:52. (28): Command(s) to 'asiqsample.asiqsample':
T. 2006/09/18 11:06:52. (28): 'insert into test2 (id, fname)
values (1, 'George')'
E. 2006/09/18 11:06:52. ERROR #1028 DSI EXEC(115(1) asiqsample.asiqsample) -
dsiqmint.c(3034)
Message from server: Message: 2706, State 0, Severity 11 --
'ASA Error -141: Table 'test2' not found'.
H. 2006/09/18 11:06:52. THREAD FATAL ERROR #5049 DSI EXEC(115(1)
asiqsample.
```



```
asigsample) - dsiqmint.c(3041)
```

```
The DSI thread for database 'asigsample.asigsample' is being shutdown.
DSI received data server error #2706 which is mapped to
STOP_REPLICATION. See logged data server errors for more information.
The data server error was caused by output command #1 mapped from input
command #1 of the failed transaction.
```

```
I. 2006/09/18 11:06:52. The DSI thread for database 'asigsample.asigsample'
```

Solution

When the target table owner is anyone other than the connection maintenance user `rsuser`, the replication definition must specify the appropriate Replicate Table Owner for the replicate (target) table.

Once the replication definition has been updated appropriately, right-click the target database connection within RSM and select Resume. Check the Resume DSI and Skip first transaction boxes and select OK.

The Target database connection should return to an UP state within RSM. The RSM may need to be refreshed to reflect the state change.

Creating Publications within RSM

Publications are used to replicate multiple source tables as a group. The process creates replication definitions for multiple source tables quickly. These are the steps for creating a publication:

- 1 Using the RepServer's Sybase Central (3.2), connect to RSM with the `sa` user.
- 2 Locate the Publication folder within the source database and select Add Publication.
- 3 Enter a publication name at the top, highlight the desired source tables from the Available Items list, and select Add for each table.
- 4 Highlight each table within the Selected Items list and select Edit to modify the replication definition parameters.
- 5 Modify the replication definition name if desired, and specify the Replicate Table Owner if necessary.
- 6 Define a primary key within the replication definition.
- 7 Select the Column tab, highlight the desired primary key column from Selected Columns, and select the Column Properties button.
- 8 Check off the Primary Key box and select OK to save.

The graphic for the primary key column will be a key image.

Enabling Sybase IQ tracing

To enable Sybase IQ tracing, add the following parameters to the IQ start-up command line:

```
-zr ALL -zo <zr_log_file>
```

where *<zr_log_file>* is the path and name for the *.zr* trace log file. This turns on extensive tracing within the Sybase IQ server. Additionally, you can check the standard Sybase IQ database log *.iqmsg*.

Rounding with the FLOAT data type

Minor rounding occurs after replicating FLOAT data types from ASE to Sybase IQ. For example:

```
0222222 was replicated as .022222199999999997
```

Defining a new connection using RSM may fail with execute permission error

Symptom	While defining a new connection using RSM, you may receive the following error: "Execute permission denied on rs_configure."
Solution	If this error occurs, grant execute permission on the rs_configure stored procedure within the ASE RSSD database. This can be done using Sybase Central or manually using isql.

Sybase IQ runs out of IQ STORE dbspace

Symptom	You need to add another dbspace, if IQ stops processing transactions and the IQ session log (console) reports the following error: "You have run out of IQ STORE dbspace in database /s1/Sybase/ASIQ-12_6/demo/asiqdemo.db. In another session, please issue a CREATE DBSPACE...IQ STORE command and add a dbspace of at least 8MB."
Solution	To add another dbspace, execute the following command in Interactive SQL: <pre>CREATE DBSPACE asiqdemo4 AS '/s1/sybase/asiqdemo4.iq' IQ STORESIZE 500</pre> <p>This command adds 500MB to the database.</p>

Sybase IQ database status DOWN and lower LTL version error reported

Symptom

After you set up your replication environment and try to insert rows of data into ASE, you should be able to see the replicated data within the IQ tables. Connect to the IQ database, issue a COMMIT, and run a SELECT statement to see the data. You may get an error from the Sybase Central session, which puts the DSI on the IQ database status as DOWN with the following error in your ASE log file:

```
"00:00000:00027:2006/09/18 15:46:19.60 server
Started Rep Agent on database, 'rapt' (dbid = 5).
00:00000:00027:2006/09/18 15:46:19.72 server
RepAgent(5): Connecting to a Replication Server that
supports a lower LTL version. Features that are not
supported by the lower LTL version may cause the
RepAgent to shutdown.
Set the 'skip unsupported features' configuration
option to prevent shutdown.
00:00000:00027:2006/09/18 15:46:19.72 server
RepAgent(5): The Replication Server LTL version should
be at least 700 for full support of features.
(Replication server LTL version 600, Rep Agent LTL
version 700)."
```

Solution

Connect to ASE using isql and run the following stored procedure to disable unsupported features:

```
use database
go
sp_config_rep_agent <database>,
'skip unsupported features','true'
go
```

Sybase IQ permission denied error

Symptom

While trying to replicate the inserted data from ASE to IQ, you may get the following error on your IQ tables: "Permission denied: you do not have permission to insert into "<TABLE_NAME>" table."

Solution

Replication Server needs access to update the table in the manner in which data or DDL comes across. Therefore, Sybase recommends granting full permission to the user that the Replication Server is using to apply commands to the target table(s).

To grant full permission on the target Sybase IQ tables, follow these steps:

- 1 Connect to the database using Sybase Central.
- 2 Expand the Tables folder.
- 3 Right-click on the table name.
- 4 Select the Permission tab and click the Grant button.
- 5 Select the user to grant full table access.

Index

A

- Adaptive Server Enterprise
 - as replication source database 69
 - configuration file 53
 - configuration settings 53
 - database setup 9
 - replication schema generation 69
- agents
 - Batch File 44, 51
 - File Load 44, 51
- ASE
 - as replication source database 69
 - configuration file 53
 - configuration settings 53

B

- Batch File Agent 44, 51

C

- configuration
 - OpenSwitch 33
 - settings for ASE 53
 - settings for RAP cache 53
- Core components
 - custom install 9
 - installation 9
 - installation layout 10
 - installer 11
 - postinstallation tasks 12
 - typical install 9
 - uninstalling 28

D

- data loading 41
- data replication
 - ASE source database setup 69
 - environment 67
 - overview 67
 - replicating test data 77
 - replication server configuration 70
 - source database setup 69
 - Sybase Central 68
 - troubleshooting 78
- database
 - building for Adaptive Server Enterprise 9
 - building for RAP cache 9
 - setup for ASE 9
 - setup for RAP cache 9
 - setup for Risk Analytics 9

E

- environment variables
 - for automated load 45

F

- File Load Agent 44, 51
- files
 - automated load errors 51

H

- High Availability
 - OpenSwitch 5
 - RAPCache 5

I

- installation
 - components 2
 - core components 2
 - licensing requirements 4
 - media 6
 - OpenSwitch 29, 32
 - optional components 2
 - overview 1
 - planning 5
 - support tools 3
- instrument
 - adding to RAPCache 52
 - proc_GetInstrumentIdBySymbol 52
 - proc_GetInstruments procedure 52
 - procedures to add 52
- interfaces file
 - OpenSwitch 36

L

- licensing requirements 4
- loading
 - additional TAQ data 49
 - automated load 41
 - automated load directories 43
 - automated load environment variables 45
 - automated load overview 41
 - automated load scripts 44
 - checking automated load process 51
 - error reporting 51
 - monitoring agents 51
 - starting automated load process 47
 - stopping automated load process 47
 - TAQ data 41, 47
 - task summary 42

M

- maintenance access
 - procedures permissions 72
 - replication table permissions 74
 - table permissions 73

O

- ODBC
 - creating connections for Sybase IQ 27
 - creating connections on UNIX 27
 - creating connections on Windows 27
- OpenSwitch
 - configuration file 36
 - configuration options 30
 - configuration tool 33, 35
 - configuring 33
 - documentation 30, 32
 - High Availability 5
 - installing 32
 - interfaces file 36
 - manual configuration 36
 - minimum requirements 29
 - operating system patches 31
 - operating system requirements 31
 - overview 29
 - postinstallation tasks 37
 - preconfiguration tasks 34
 - preinstallation tasks 31
 - RESPONSE_TIMEOUT option 36
 - setting environment variables 37
 - shutting down 39
 - starting 37
 - stopping 39
 - verifying client connections 38
- Operations Console
 - overview 5
 - supported platforms 5

P

- package components, installing
 - PowerDesigner 23
- postinstallation 12
- postinstallation tasks
 - server status 15
 - starting RAPCache 14
- preinstallation tasks 6

R

- RAP cache
 - configuration settings 53
 - database setup 9
- RAPCache
 - configuration file 53
 - configuration settings 53
 - High Availability 5
 - overview 5
- replication definitions 76
- replication server
 - configuration 70
 - configuration tasks 70
 - RSM configuration 71
 - source database setup 69
 - target database configuration 71
- replication server configuration
 - replication server manager 74
- Replication Server Manager
 - defining servers 74
 - source and target connections 75
- RESPONSE_TIMEOUT option 36
- Risk Analytics Platform
 - core components 2
 - installation components 2
 - installation planning 5
 - optional components 2
 - overview 1
 - support tools 3
 - uninstalling 28
- RSM
 - creating subscriptions 76
 - defining servers 74
 - replication definitions 76

S

- source and target connections 75
- subscriptions, creating 76
- Sybase Central 68
- SySAM
 - activating your license 12
 - licensing requirements 4
 - obtaining your license 4

T

- target database
 - configuration 71
 - maintenance access 72
 - maintenance user 72
- troubleshooting data replication 78

U

- uninstalling
 - core components 28
 - RAP components 28

V

- VLSBServer
 - overview 5

