SYBASE®

Installation and Configuration Guide

Risk Analytics Platform

2.0

DOCUMENT ID: DC00244-01-0200-01

LAST REVISED: September 2005

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

Contents

About This Book

Audience

Risk Analytics Platform Installation and Configuration Guide is intended for Sybase® Professional Services, Customer IT/Back Office engineers, 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 *Release Bulletin Risk Analytics Platform* for any last minute information regarding this product.

Related documents

Refer to the following documents for more information:

- *Release Bulletin Risk Analytics Platform*
- *Risk Analytics User's Guide*
- Sybase IQ 12.6 product documentation
- Adaptive Server Enterprise 15.0 product documentation
- PowerDesigner® 11.1 product documentation
- Replication Server 12.6 product documentation

Other sources of information

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

- The Getting Started CD contains the release bulletin, installation 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 time frame and then click Go.

Note The Product Family for Risk Analytics Platform 2.0 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 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 2.0 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.

CHAPTER 1

Sybase Risk Analytics Platform

About this chapter

This chapter describes Sybase Risk Analytics Platform installation procedures.

Contents

Overview

The Sybase Risk Analytics Platform (RAP) is a software infrastructure based on existing Sybase products that captures, stores, and manages real-time market data.

By integrating high volume data feeds with deep historical data in a common data model and data management platform, Sybase Risk Analytics delivers query results in a time frame required for near real-time trading decisions.

Installing and configuring Risk Analytics Platform involves working with two types of data, two types of databases, and two types of loading processes.

- Two types of data:
  - real-time TAQ (Trades and Quotes) data
  - historical end-of-day trading data
- Two types of databases:
  - RAP in-memory cache for TAQ data
  - RAP repository for historical data
- Two types of loading:
  - an automated loading process monitors input of TAQ data and loads this data in parallel to the RAP cache and the RAP repository
  - an end-of-day process loads historical data into the RAP repository

This book describes how to configure Risk Analytics Platform. There are four major tasks:

1 Set up the RAP in-memory cache database

2 Set up the RAP repository database

3 Set up and run the load process for the TAQ data

4 Set up and load the sample data

Accessing the data using queries and customizing the RAP cache and repository databases are discussed in the *Risk Analytics User's Guide*.

What's in the package?

Sybase Risk Analytics Platform includes these products:

- Sybase IQ 12.6 serves as the market data repository
- Sybase Adaptive Server Enterprise 15.0 (ASE) provides an in-memory data cache for near real-time data access
- Sybase Risk Analytics Platform 2.0 includes a data model, configuration scripts, DDL scripts, load scripts, sample data, and sample queries
- Sybase Power Designer 11.1 includes modeling tools you can use to customize the data model

Platforms and Operating Systems

Refer to the *Release Bulletin Risk Analytics Platform* for information on available platforms that support Sybase Risk Analytics Platform.

Note
The components on the Risk Analytics Platform 2.0 product CD and the PowerDesigner 11.0 product CD are installed on a Windows system.

Adaptive Server Enterprise 15.0 and Sybase IQ 12.6 are installed on one of the supported platforms described in the "Product summary" section of the *Release Bulletin Risk Analytics Platform*.

Risk Analytics Platform and data replication

You can use Sybase Replication Server 12.6 to replicate transactions between a source database containing reference data and the Sybase IQ data repository to update the repository data at the end of the trading day or whenever an update occurs. See Appendix B, "Data Replication" for additional information.

Installation

Default locations during installation may differ for some components. For example, the default root directory may appear as *c:\sybase* or as *c:\Program Files\Sybase*, depending on the component. In this document, the installation directory is referred to as *$SYBASE* (UNIX) or *%SYBASE%* (Windows).

Before you begin

- Read the release bulletin

  Your Risk Analytics Platform package includes a printed release bulletin, which is also available on the Sybase Product Manuals Web page. The release bulletin may contain information not available from other documentation sources.

- Review the system requirements

  Knowing the requirements for each product included in this package can make installation and configuration tasks easier.

- Locate your license keys

  Some of the products distributed with the Risk Analytics Platform require a license key to install. License keys are included as separate printed documents in your product package.

- Accept the license agreement

  During installation, you must accept the Sybase license agreement to install Sybase software.

  If, for any reason, you cannot read the license agreement, go to the Sybase Software Licenses web site at http://www.sybase.com/softwarelicenses and follow the instructions for your geographical location.

  After you read the License Agreement, return to the installation screen and click **Yes** to agree to the terms and conditions of the license.

Installing package components

This section summarizes installation procedures for each component included in your package.

Risk Analytics Platform 2.0

To install Risk Analytics Platform 2.0 on a Windows system from the CD included with your Risk Analytics Platform package:

1 Insert the Risk Analytics Platform product CD into the CD drive.

If Risk Analytics Setup does not start automatically:

- Click Start, choose Run.
- Browse to select your CD drive letter, choose *setup.exe*.

2 Follow the instructions on your screen to install Risk Analytics Platform.

The Risk Analytics Platform installation directory structure

The default installation directory for Risk Analytics Platform is *c:\Program File\Sybase\Risk Analytics Platform 2.0* on Windows. There are four major subdirectories within the installation directory:

- *\Data* contains *\IQ* and *\ASE* subdirectories, which contain input and load directories for the sample data files and load scripts.
- *\Documentation* contains Risk Analytics Platform 2.0 documentation for this release in Adobe Acrobat PDF format and Microsoft Excel format.
- *\Model* contains *\IQ* and *\ASE* subdirectories for the Sybase IQ and ASE data models and scripts that support the data models.
- *\Scripts* contains *RAP_Load_Scripts*, *RAP_Queries*, and *RAP_Setup* folders. These subdirectories contains SQL scripts for the data load, sample queries, and database configuration files.

Sybase IQ 12.6

To install Sybase IQ 12.6 from the CD included with your Risk Analytics package, see the appropriate installation and configuration guide for your platform for information about installing the following software components:

- Sybase IQ
- Sybase Central Java™ Edition

- Open Client Software Developer's Kit (SDK)
- Sybase jConnect JDBC Driver
- Java Runtime Environment

Installation options

Refer to the following sections in the installation and configuration guide for the type of components you want to install:

- To install components required for operation as a network server, install the Sybase IQ Server Components on your UNIX system. See "Installing Sybase IQ Server Components" in Chapter 1, "Installing Sybase IQ."
- To install components required for connection to a network server, see Chapter 2, "Installing Sybase IQ Network Client."

Note Before installing Sybase IQ 12.6, refer to Chapter 2, "Repository Database Setup" for any additional information you may need during the installation.

Check the online Sybase Support Page at http://www.sybase.com/support for software updates (ESD or EBF) and install the latest available ESD or EBF after you install Sybase IQ 12.6 from the product CD.

Additional information

For platform specific installation issues, refer to the following documents:

- *Sybase IQ 12.6 Installation and Configuration Guide for Sun Solaris*
- *Sybase IQ 12.6 Installation and Configuration Guide for IBM RISC SYSTEM/6000 AIX*

Adaptive Server Enterprise 15.0

To install Adaptive Server Enterprise 15.0 from the CD included with your Risk Analytics Platform package, see the appropriate installation and configuration guide for your platform. Selecting `Typical Install` during the installation process will install the default ASE components necessary for Risk Analytics Platform. Assign a minimum of 500MB of space to the master device (the default is 30MB); select the 4k page size.

Note Before installing Adaptive Server Enterprise 15.0, refer to Chapter 3, "Cache Database Setup" for any additional information you may need during the installation.

Check the online Sybase Support Page at http://www.sybase.com/support for software updates (ESD or EBF) and install the latest available ESD or EBF after you install Adaptive Server Enterprise 15.0 from the product CD.

Additional information

For platform specific installation issues, refer to the following documents:

- *Adaptive Server Enterprise 15.0 Installation Guide for Sun Solaris*
- *Adaptive Server Enterprise 15.0 Installation Guide for IBM AIX*

PowerDesigner

PowerDesigner 11.1 provides full support for the Risk Analytics Platform 2.0 data model. To install PowerDesigner 11.1, you must first install PowerDesigner 11.0 from the CD included with your Risk Analytics Platform package, then download and install PowerDesigner 11.1 Maintenance Release and PowerDesigner EBF 12918: 11.1.0 ESD #5 (or later).

Installing PowerDesigner 11

1 Insert the PowerDesigner 11.0 product CD into your CD drive on a Windows system.

   If PowerDesigner Setup does not start automatically:

   - Click Start, choose Run.
   - Browse to select your CD drive letter, choose *setup.exe*.

2 Follow the instructions on your screen to install PowerDesigner. See the PowerDesigner product documentation for additional information.

Installing PowerDesigner 11.1 Maintenance Release

Use these instructions to download and install PowerDesigner 11.1 Maintenance Release.

Note You must install the 11.0 Maintenance Release *before* you install EBF 12918: 11.1.0 ESD #5 (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 PowerDesigner 11.1 Maintenance Release 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.

Installing EBF 12918: 11.1.0 ESD #5

Use these instructions to download and install PowerDesigner EBF 12918: 11.1.0 ESD #5 (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 12918: 11.1.0 ESD #5 (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.

Managing databases with Sybase Central

Sybase Central is a graphical management tool, which is installed as part of many Sybase products, including Sybase IQ and Adaptive Server Enterprise. By registering Sybase IQ and ASE as plug-ins in a single instance of Sybase Central 4.3, you can connect to ASE and Sybase IQ through the same interface.

This example shows you how to add Sybase IQ 12.6 to Sybase Central for ASE 15.

1 Start Sybase Central for ASE 15.

   - In Windows, click Start, point to Programs | Sybase | Sybase Central 4.3.
   - From a UNIX command line, enter `source $SYBASE/ SYBASE.csh`, then enter `$SYBASE/ASEP/bin/aseplugin`.

2 Click Tools, choose Plug-ins.

3 Click Register, then choose Register a plug-in by specifying a plug-in registration file.

4 Click Browse, navigate to:

   - <*Sybase IQ path*>*/java/* (UNIX)
   - <*Sybase IQ path*>*java*\ (Windows)

5 Choose *IQPlugin1260.jar*.

6 Click Finish.

CHAPTER 2

Repository Database Setup

About this Chapter

This chapter includes the steps you use to create and load the RAP repository in a Sybase IQ database.

Contents

Overview

Risk Analytics Platform 2.0 uses Sybase IQ 12.6 as a consolidated data repository. The repository database provides fast access to large amounts of historical, reference, and real-time data.

Where to find additional information about Sybase IQ
To find additional information about Sybase IQ, see *Introduction to Sybase IQ*, the *Sybase IQ System Administration Guide*, or the *Sybase IQ Performance and Tuning Guide*, especially if you are using IQ on a UNIX platform.

Baseline requirements

The system on which the RAP repository database is created requires Sybase IQ 12.6 server software and Sybase IQ 12.6 client tools. Sybase IQ 12.6 server software requires a minimum of 3.5GB of disk space for the repository database. Sybase IQ 12.6 client tools must be loaded on the machine from which you install the database. Both the Sybase IQ client and server can be on the same machine.

Notes

- The names of the database, user, and ODBC Data Source Name (DSN) are required when you set up a Sybase IQ database. The examples in this section use RAPIQ as the database name, and DBA as the user name.
- Sybase IQ 12.6 requires a minimum of 3.5GB of disk space for the sample data in the repository database. Your own data may require more or less disk space; adjust this value as appropriate.

Creating a Sybase IQ database

This section tells you how to create the RAP repository database and database server with Sybase Central and Interactive SQL Java.

Using Sybase Central to create the database

Sybase Central is a graphical management tool for Sybase products and plug-ins. Sybase Central includes a database creation utility you can use to start the server and create the database.

1 Start Sybase Central.

- On Windows, click Start, point to Programs | Sybase | Adaptive Server IQ 12.6 | Sybase Central Java Edition.
- If you use Sybase Central on the UNIX Platform, first set the DISPLAY variable. The command you use depends on your shell:

Shell...	Command...
C-shell	setenv DISPLAY <terminal host name>:0.0

Shell...	Command...
B-shell	export DISPLAY=<terminal host name>:0.0

- To start Sybase Central on UNIX, source the *ASIQ-12_6.sh* (or *.csh*) script in the *$SYBASE/ASIQ-12_6* directory, then enter this command:

```
% scjview
```

2 Click Sybase IQ | Utilities tab | Create Database. A dialog box prompts you for the server options.

3 Click Next. The next dialog box prompts you for the connection parameters.

4 Specify an appropriate user account with DBA authority and connection parameters for the server. See the table below for IQ system defaults.

Parameter	Value
Username	*DBA*
Password	*SQL*
Port Number	*2638*
Host machine name	Machine where IQ is installed and where you want to create the database.

5 Click Next. A dialog box prompts you for the name and location of the sample IQ files.

6 Specify an appropriate path and filename for the RAP repository database, for example, *<path>\RAPIQ.db*

Note The Sybase IQ server and the IQ database must have identical names, if you are using IQ as a target database for Replication Server. This is a requirement of Replication Server.

For more information, see Appendix B, “Data Replication.”

This path should exist locally on the given host. If you omit the path, Sybase Central creates the database relative to the directory where you started the server.

7 Type a path and filename for the IQ Store, which contains the IQ data, for example, *<path>\RAPIQ.iq*

8 If your RAP repository database is on an operating system file system, deselect the *This is a raw device* checkbox, then click Next.

9 On the next dialog box, enter an IQ Store Size of 2 GB (2000 MB) and accept the defaults for IQ_Page_Size and Database Page Size.

10 Click Next. A dialog box prompts you for the temporary store path and size.

11 If your database is on an operating system file system, deselect *This is a raw device*.

- Deselecting this option creates a temp file IQ uses as work space during processing. Add the full path and file name for the temporary dbspace file you want to use. This file must have an *.iqtmp* extension.
- Specify 1000 (1 GB) for File Size in MB. Accept the defaults for all remaining parameters, and click Next until you reach the Default Database Attributes page.
- On the Default Database Attributes page, click the Emulate as ASE button. Accept the remaining defaults, and click Next until you reach the summary screen.
- On the summary screen, review your choices, click Finish.

Using Interactive SQL to create the database

As an alternative to Sybase Central, you can use Interactive SQL to create a database. This example uses a CREATE DATABASE statement to create a Sybase IQ database called RAPIQ on a raw device in UNIX:

```
CREATE DATABASE 'RAPIQ'
 IQ PATH '/dev/rdsk/c2t6d0s3'
```

See the Sybase IQ documentation for details.

Note To create an IQ database with multiplex capability, you must use Sybase Central. See the Sybase IQ documentation for details.

Editing the IQ configuration file

The easiest way to start Sybase IQ is with a configuration file. You can make a copy of the *\Sybase\ASIQ-12_6\demo\asiqdemo.cfg* configuration file that is created during IQ installation, rename it, and edit it for use with your RAP repository database. The default *asiqdemo.cfg* file sets these Sybase IQ 12.6 startup options:

```
-n <hostname>_asiqdemo
-c 32m
-gc 20
-gd all
–gl all
-gm 10
-gp 4096
-ti 4400
-x tcpip(port=2638)
```

Use a text editor to change the following values:

- Set the -n parameter to the system and server of your sample IQ database. For example, if your hostname is RAP, set -n RAP_RAPIQ
- -c 32m is for 32-bit platforms. Set -c 48m for 64-bit platforms.
- Set -gm 40 for the number of connections.
- Set the port to an available port number for that server, if port number 2638 is already in use. The default port number for IQ is 2638.
- Include the host name, if you intend to start IQ from a remote server. For example,

  ```
  -x tcpip(host=IQ_server;port=1234)
  ```

- Add the following parameters:

  ```
  -iqmc 256
  -iqtc 128
  ```

  The -iqmc parameter specifies 256 MB for IQ Main buffer cache and the -iqtc parameter specifies 128 MB for IQ Temp buffer cache. Increase the values of -iqmc and -iqtc, if more memory is available.

- Rename the file, keeping the *.cfg* suffix, for example, *RAP_RAPIQ.cfg*

Note
This *.cfg* file must be in the same location as the **.db* and **.iq* files.

These should be the only changes required to use the configuration file with your RAP repository database. For more information on startup options and configuration files, see the Sybase IQ documentation.

Creating ODBC connections

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 about creating ODBC data sources, see the chapter "Configuring Sybase IQ" in the *Sybase IQ 12.6 Installation and Configuration Guide for Windows*.

UNIX

On UNIX 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 12.6 Installation and Configuration Guide* for your server platform.

Starting the RAP repository database

Although there are several ways to start a Sybase IQ server and connect to a database, these procedures give you some simple ways to get started.

Starting the database from the command line

- On the system where you created the Sybase IQ database, start the IQ server with the RAP repository database.

  For example, to start the server from the command line using the edited configuration file, change to the directory containing the *.cfg* and *.db* files and enter:

  ```
  start_asiq @RAP_RAPIQ.cfg RAP_RAPIQ.db
  ```

  Use the -n switch to name the server, either in the configuration file or on the command line when you start the server.

- Connect to the sample IQ database with the user ID DBA and password SQL.

  For example, to connect using Interactive SQL from the machine where the Sybase IQ server RAPIQ and the database RAPIQ are running, you can enter on the command line:

  ```
  dbisql -c "uid=DBA;pwd=SQL;eng=RAPIQ_RAPIQ;
  dbn=RAP_RAPIQ;dbf=RAP_RAPIQ.db;
  links=tcpip{port=2638;host=IQ_server}"
  ```

Starting the database from Sybase Central

In Windows, you can also use Sybase Central to start the server and connect to the RAP repository database.

1 Start Sybase Central.

  - On Windows, click Start, point to Programs | Sybase | Adaptive Server IQ 12.6 | Sybase Central Java Edition.
  - On UNIX, source the *ASIQ-12_6.sh* (or *.csh*) script in the *$SYBASE/ASIQ-12_6* directory, then enter this command:

    ```
    % scjview
    ```

2 Click Sybase IQ | Utilities tab | Start Server.

  A dialog box prompts you for the start-up options.

3 Select Single server and click Next.

4 On the Connection Parameters dialog, type the appropriate connection parameters. See the following table for system defaults.

Field name...	Value
Server Name	*server_name*
Host machine	*host_name*
Port number	*port_number*
Database Name	*database_name*
Username	DBA
Password	SQL

5 Click Next.

6 Specify the full path to the database file, then click Next.

7 Click Finish.

Other startup and connection options
See the *Sybase IQ Installation and Configuration Guide* for your platform for more information.

Creating database objects

At this point the IQ database is ready for you to execute the DDL script and create objects in the RAP repository database.

1 Start Interactive SQL Java.

- On Windows, click Start, point to Programs | Sybase | Adaptive Server IQ 12.6 | Interactive SQL Java.
- On UNIX, enter this command:

```
dbisql
```

2 Use the ODBC DSN you created to connect to the RAP repository database.

3 Click File, choose Open.

4 On the Open dialog, using the Look in drop-down menu, go to the directory where the DDL script *RAP_IQ_Tables.sql* is located. This script is installed in the *%SYBASE%\Model\IQ* folder. Select the *.sql* script file and open it.

5 Click the Execute SQL statement icon on the toolbar to execute the script and create the RAP repository database tables and other database objects.

Indexes in the IQ database

The RAP data model includes only those indexes that support the sample queries. Statements needed to create these indexes appear in the DDL scripts generated from the RAP IQ data model. Consequently, the indexes supplied with the model will be created automatically when you run the corresponding DDL scripts (see above).

Depending on site-specific issues such as limits on the available load time and the actual queries in the database, you may want to add or remove indexes from the RAP IQ data model. For detailed information on IQ indexes, refer to the Sybase IQ product documentation.

Loading the sample data

You work with four file types to load the sample data:

- Shell script files (*.sh*) start the database setup and data loading processes.
- Load scripts (*.tst*) load the data files.
- Data files (*.csv*) contain sample data.
- Script files (*.sql*) set up the repository and create the indexes.

Before you can run any of the scripts, you must move the sample data to the server, and edit the *.sh* and *.tst* files to correspond to the repository database you created.

Moving the sample data to the server

All sample data for the RAP repository is located in the *%Sybase%\Data\IQ\RAP_Input* directory. Although Sybase distributes sample data as comma separated value (*.csv*) files, some data has been compressed (*.zip)* to conserve disk space.

Preparing the reference sample data

The reference table sample data is loaded into all tables except the STOCK_TRADE and STOCK_QUOTE tables. You must unzip one file, then use FTP to transfer all of the *.csv* data files to the machine on which the IQ server resides.

1 Unzip the *STOCK_HISTORY.zip* file located in the *%Sybase%\Data\IQ\RAP_Input* directory.

2 Use FTP to transfer all of the *.csv* files in the *%Sybase%\Data\IQ\RAP_Input* folder from Windows to the UNIX machine on which the IQ server resides. The *.csv* files, which contain the sample data, **must** be transferred in ASCII format. The sample data in these files is loaded into the reference tables in the ASE database.

Note
The *.csv* files must be in the same location (folder or directory) as the Sybase IQ database files and the directory from which the Sybase IQ server is started.

Preparing the TAQ sample data

The TAQ sample data, which is loaded into the STOCK_TRADE and STOCK_QUOTE tables, must be unzipped and transferred using FTP.

1 Unzip the *STOCK_QUOTE.zip* and *STOCK_TRADE.zip* files located in the *%Sybase%\Data\IQ\RAP_Input* directory.

2 Use FTP to transfer in ASCII mode the data files *STOCK_QUOTE.csv* and *STOCK_TRADE.csv* from Windows to the UNIX machine on which the IQ server resides.

Note The *.csv* files, which contain the sample data, **must** be transferred in ASCII format.

Refer to the section “Loading additional TAQ sample data” in Chapter 4, “Loading TAQ Data” for information on loading this sample data into the STOCK_TRADE and STOCK_QUOTE tables in the RAP repository.

Editing the shell script files

Three shell script (*.sh*) files start the database setup and data loading processes:

Filename	Location
rap_setup.sh	*%Sybase%\Scripts\RAP_Setup\IQ*
load_unix_rap_tables.sh	*%Sybase%\Data\IQ\RAP_Load*
load_unix_IQ_TAQ_tables.sh	*%Sybase%\Data\IQ\RAP_Load*

These files contain placeholders that you must change to values which correspond to the RAP repository database. Use a text editor to open each file and make the following changes:

Change this...	To your value for this...
<USER_NAME>	User ID created for the target database
<PASSWORD>	Corresponding password
<ENG_NAME>	Server name
<DB_NAME>	Database name
<DB_FILE>	Database file name
<PORT_NUMBER>	Port number
<HOST_NAME>	Host name

Editing the load scripts

load_unix_rap_tables.tst and *load_unix_IQ_TAQ_tables.tst* (see the table for the location of these files), which load the data files, contain a placeholder <*path*> that points to the location of the data files.

Filename	Location
load_unix_rap_tables.tst	*%Sybase%\Data\IQ\RAP_Load*
load_unix_IQ_TAQ_tables.tst	*%Sybase%\Data\IQ\RAP_Load*

1 Use a text editor to open *load_unix_rap_tables.tst* and globally change <*path*> to the location of the reference tables sample data *.csv* files on the server.

2 Use a text editor to open *load_unix_IQ_TAQ_tables.tst* and globally change <*path*> to the location of the *STOCK_TRADE.csv* and *STOCK_QUOTE.csv* TAQ sample data files on the server.

In both scripts, the path must reflect the location of *.csv* files as the server sees them and the *.csv* files must be accessible to the server.

Moving the scripts to the server

Use FTP to move all script files from *%Sybase%\Scripts\RAP_Setup\IQ* and *%Sybase%\Data\IQ\RAP_Load* to the host machine where the IQ server resides.

Running the RAP repository database setup and load scripts

To run the scripts that setup the RAP repository and load the sample data into the reference tables, you must have the Sybase IQ 12.6 client tools installed on your system.

After you make the appropriate changes to the scripts, run the *rap_setup.sh* and *load_unix_rap_tables.sh* scripts (in this order) on the machine that hosts the Sybase IQ database. These scripts start the processes that complete the RAP repository database setup and load the reference table sample data.

Note
The scripts executed by *rap_setup.sh* create global temporary tables and set Sybase IQ database options.

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

Refer to the section “Loading additional TAQ sample data” in Chapter 4, “Loading TAQ Data” for information on loading sample data into the STOCK_TRADE and STOCK_QUOTE tables in the RAP repository.

CHAPTER 3

Cache Database Setup

About this Chapter

This chapter includes the steps you perform to create an in-memory cache database and load the sample data.

Contents

Overview

Sybase Adaptive Server Enterprise features an asynchronous commit option that allows data to be stored in an in-memory cache, ensuring extremely fast load and retrieval times. Using the ASE data store for current market data provides a separate memory space for traders who need up-to-the-second information for pretrade analysis, but may not need the full historical data held in the RAP market data repository.

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

Where to find more information about Adaptive Server Enterprise
To find additional information about Adaptive Server Enterprise, refer to the Adaptive Server Enterprise documentation set included in the ASE product package.

Preliminary steps

Install Adaptive Server Enterprise 15.0. Assign a minimum of 500MB of space to the master device (the default is 30MB); select the 4k page size.

Adaptive Server Enterprise requires a minimum of 5GB of disk space for the sample data in the cache database. Your own data may require more or less disk space; adjust this value as appropriate.

Refer to the ASE product documentation for additional instructions. Check the online support Web site for the latest software updates (ESD or EBF). See "Sybase EBFs and software maintenance" on page ix for more information.

ASE configuration recommendations

After installing Adaptive Server Enterprise, edit the default configuration file according to the configuration recommendations in Appendix A, "Sample ASE Configuration Settings." This appendix also contains an example of configuring memory for cache.

The sample ASE configuration file *RAP_ASE_Server.cfg* (described in Appendix A, "Sample ASE Configuration Settings") is installed in the folder *%SYBASE%\Risk Analytics Platform 2.0\Scripts\RAP_Setup\ASE*.

Creating the ASE database objects

Risk Analytics Platform 2.0 uses two temporary ASE databases:

- rapt for trade data
- rapq for quotes data

RAP includes a Data Definition Language (DDL) script that creates database objects in the model database for the rapt and rapq temporary databases. This DDL script is called *RAP_Table.sql* and is installed in the *%SYBASE%\Risk Analytics Platform 2.0\Model\ASE* folder.

Log into your ASE server and execute the DDL script *RAP_Table.sql* to create the RAP database objects in the model database.

For a description of modifications made to the DDL in the *RAP_Table.sql* file, refer to the section "RAP cache" in Chapter 4 "Generating DDL Scripts" in the *Risk Analytics User's Guide*.

Device initialization and ASE database creation

The sample ASE device creation script file *RAP_DB_Devices.sql* is installed in the *%SYBASE%\Risk Analytics Platform 2.0\Model\ASE* folder. This script contains SQL statements that set up ASE database devices, create the two temporary databases (rapt and rapq), and map segments to multiple devices.

Edit the <PATH> placeholder in the *RAP_DB_Devices.sql* script to include the pathnames to your devices. If possible, create your database devices on different physical devices. The script initializes and stripes the databases across four data devices; add or delete disk init commands in the script to match the number of devices available in your environment.

Note The *RAP_DB_Devices.sql* script creates devices to accommodate the RAP sample data. When you create devices for your own data, change the number and size of the devices to suit your data and environment. See the Adaptive Server Enterprise 15.0 *System Administration Guide: Volume 2* for more information.

After you modify *RAP_DB_Devices.sql*, log into the ASE server and execute the script to create the devices and temporary databases. The rapt and rapq databases now contain the Risk Analytics Platform tables, indexes, and other database objects that were created in the model database in the previous section.

Loading ASE sample data

You work with three file types to load the sample data:

- Shell scripts (*.sh*) load the sample data.
- Data files (*.csv)* contain the sample data.
- SQL scripts (*.sql*) contain database commands.

Before you can run the scripts, you must move the sample data to the server, and edit the *.sh* and *.sql* files to correspond to the RAP cache database you created.

Moving the sample data to the server

All sample data for the RAP cache is located in the *%Sybase%\Data\ASE\RAP_Input* directory. Although Sybase distributes sample data as comma separated value (*.csv*) files, some data has been compressed (*.zip)* to conserve disk space.

Preparing the reference sample data

Use FTP to transfer all of the *.csv* files in the *%Sybase%\Data\ASE\RAP_Input* folder from Windows to the UNIX machine on which the ASE server resides. The *.csv* files, which contain the sample data, **must** be transferred in ASCII format. Make sure all *.csv* files are accessible to the ASE client. The sample data in these files is loaded into the reference tables in the ASE database.

Preparing the TAQ sample data

The TAQ sample data, which is loaded into the STOCK_TRADE and STOCK_QUOTE tables, must be unzipped and transferred using FTP.

1 Unzip the *STOCK_QUOTE.zip* and *STOCK_TRADE.zip* files located in the *%Sybase%\Data\ASE\RAP_Input* directory.

2 Create a work directory (for example, */data/source/ASE*) on your UNIX machine where the ASE server resides.

3 Use FTP to transfer in ASCII mode the data files *STOCK_QUOTE.csv* and *STOCK_TRADE.csv* from Windows to the UNIX machine on which the ASE server resides.

Note The *.csv* files, which contain the sample data, **must** be transferred in ASCII format.

Refer to the section “Loading the TAQ sample data” in Chapter 4, “Loading TAQ Data” for information on loading this sample data into the STOCK_TRADE and STOCK_QUOTE tables in the RAP cache.

Editing the script files

All script files are located *%Sybase%\Data\ASE\RAP_Load* directory. These files contain placeholders that you must change to values that correspond to the RAP cache database.

Shell scripts

Use an editor to change the placeholders in the *load_unix_ase_rap_tables.sh* and *load_unix_ASE_TAQ_tables.sh* scripts to match the server, database, and path names:

Change this...	To your value for this...
<ASE Server>	Server name
<ASE Database>	Database name
<path>	Full path to the *.csv* files
<Database Name>	Database name

SQL scripts

Use an editor to change the <Database Name> placeholder in the *truncate_tables_ase.sql* script to match the RAP cache database name.

Notes

- The parameters <ASE Database> and <Database Name> are equal to rapt and rapq for the ASE sample databases. You must edit these files and run the load process for *each* database.
- Globally change <path> to the full path of the *%SYBASE%\Risk Analytics Platform 2.0\Data\ASE\RAP_Load* directory or to the location of the reference table *.csv* data files. The path must reflect the location of *.csv* files as the server sees them and the files must be accessible to the server.
- Because rapq and rapt are temporary databases, all data is truncated in these databases every time the server restarts. These databases are created by the *RAP_DB_Devices.sql* script. To preserve data in the reference tables, load the data permanently in the model database. This ensures that the temporary databases rapq and rapt contain the reference data every time the server starts.

Use an editor to change the <TRADE DataBase> and <QUOTE DataBase> placeholders in the *truncate_ASE_TAQ_tables_ase.sql* script to match the cache database names. This script truncates both the STOCK_TRADE table in one cache database and the STOCK_QUOTE table in the other cache database. For the ASE sample databases, <TRADE DataBase> is rapt and <QUOTE DataBase> is rapq.

Moving the ASE scripts to the server

Use FTP to move all script files from *%Sybase%\Data\ASE\RAP_Load* to the machine where the cache database resides.

Running the ASE load scripts

After you make these changes, run the *load_unix_ase_rap_tables.sh* script from a command window on the machine that hosts the Adaptive Server Enterprise database or on a client machine. This script uses the bulk copy utility bcp to load the sample data from the *.csv* files into the reference tables of the ASE sample database.

1 Run the *load_unix_ase_rap_tables.sh* script for the rapt database with both <ASE Database> in *load_unix_ase_rap_tables.sh* and <Database Name> in *truncate_tables_ase.sql* specified as *rapt*.

2 Change the value of <ASE Database> in the *load_unix_ase_rap_tables.sh* script and the value of <Database Name> in *truncate_tables_ase.sql* to *rapq*.

3 Run the *load_unix_ase_rap_tables.sh* script for the rapq database with both <ASE Database> in *load_unix_ase_rap_tables.sh* and <Database Name> in *truncate_tables_ase.sql* specified as *rapq*.

To check the status of the data loads, review the log file for the Adaptive Server. The default location of the ASE log file is *$SYBASE/$SYBASE_ASE/install/<servername>.log* on UNIX platforms.

Refer to the section “Loading the TAQ sample data” in Chapter 4, “Loading TAQ Data” for information on loading sample data into the STOCK_TRADE and STOCK_QUOTE tables in the RAP cache.

CHAPTER 4

Loading TAQ Data

About this Chapter

This chapter tells you how to set up and run the automated load process that loads the in-memory cache and the repository databases.

Contents

Overview

Risk Analytics Platform uses a series of scripts to run and monitor an automated process that loads the in-memory cache and repository databases. These data loading scripts require certain directories to function correctly. All environment variables required to execute the bcp and dbisql utilities must also be set for the scripts to operate properly.

How data loading works

Running *rap_start_pl.sh* starts a process which 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 links 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 and repository databases. 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

- Create directories on the UNIX machine where you installed the Risk Analytics Platform repository and cache databases. See “Load process directories” on page 31 for more information.
- Use FTP to transfer the *.sql* scripts from the *%Sybase%\Scripts\RAP_Load_Scripts\IQ_Load_Scripts* directory to the *$LOAD_IQ_SCRIPT_DIR* directory on the server. These scripts load the repository tables.
- Use FTP to transfer the shell (*.sh*) scripts from the *%Sybase%\Scripts\RAP_Load_Scripts* directory to the *$RAP_HOME* directory.
- Update *set_rap_env_pl.sh* with the correct path information. This script sets the environmental variables.
- Run *$RAP_HOME/set_rap_env_pl.sh* to set the environment variables.
- Set up the environment for the `bcp` and `dbisql` utilities.
- Run *$RAP_HOME/rap_start_pl.sh* to start the load process.

Load process directories

You need to create certain directories for the load process on the UNIX machine where you installed the repository and cache databases.

The following is an example directory structure for the automated load process:

```
/work3/LoadProcess
    data
        ase_in
            active
            log
        iq_in
            active
            log
        archive
        rap_in
        .status
    iq_load_scripts
```

In this example, the values of the environment variables defined in the *set_rap_env_pl.sh* script are:

Variable name	Value
$RAP_HOME	*/work3/LoadProcess*
$RAP_DATA	*/work3/LoadProcess/data/ase_in*
$IQ_DIR	*/work3/LoadProcess/data/iq_in*
$RAP_ARCH	*/work3/LoadProcess/data/archive*
$RAP_IN	*/work3/LoadProcess/data/rap_in*
$RAP_STAT	*/work3/LoadProcess/data/.status*
$LOAD_IQ_SCRIPT_DIR	*/work3/LoadProcess/iq_load_scripts*

The two directories *active* and *log* must be created under both $RAP_DATA and $IQ_DIR.

RAP load scripts

This section describes the scripts that instantiate the loading process.

Shell scripts

Use FTP to transfer the shell (*.sh*) scripts from the *%Sybase%\Scripts\RAP_Load_Scripts*folder on Windows to the *$RAP_HOME* directory on the server. These shell scripts control the loading process and load the RAP cache tables.

Script name	Description
rap_link.sh	Links new data source files in the $RAP_IN directory to the $RAP_DATA, $IQ_DIR, and $RAP_ARCH directories.
rap_start_pl.sh	Launches the loading process.
rap_stop_pl.sh	Creates a tickler file in the $RAP_STAT directory. The loading process stops, when the tickler file exists.
rap_tbl_bcp_pl.sh	Runs the bcp utility with arguments, determines completion status, then issues error messages, if errors occurred, and logs trail files. Called by *rap_tbl_mon_pl.sh*.
rap_tbl_dbisql_pl.sh	Runs the dbisql utility with arguments and determines completion status, then issues error messages, if errors occurred, and logs trail files. Called by *rap_tbl_mon_pl.sh*.
rap_tbl_mon_pl.sh	Looks for new data source files. This is a middle level process called by *rap_start_pl.sh*.
set_rap_env_pl.sh	Sets the environment variables.
rap_truncate.sh	Truncates the tables.

SQL scripts

Use FTP to transfer the *.sql* scripts from the *%Sybase%\Scripts\RAP_Load_Scripts\IQ_Load_Scripts* folder on Windows to the *$LOAD_IQ_SCRIPT_DIR* directory on the server. These scripts load the RAP repository tables.

Script name	Description
load_iq_stock_quote.sql	Loads the STOCK_QUOTE table.
load_iq_stock_trade.sql	Loads the STOCK_TRADE table.

Setting the environment variables

Setting up the RAP environment

set_rap_env_pl.sh sets environment variables for the shell. You must edit the placeholders in this file with the correct path information, then run this script to set the environmental variables.

Variable name	Description
RAP_HOME	Base directory for RAP. All other paths should be relative to this directory.
RAP_DATA	Input directory for loading files into ASE. Monitored by the loading process.
RAP_STAT	Used for placing 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 IQ. Monitored by loading process.
LOAD_IQ_SCRIPT_DIR	Directory where sql scripts used by the dbisql utility during IQ loading are stored.
TICKLER_FILE	Name of the tickler file that stops the loading script.
MAX_PTN_QUOTE	Maximum number of partitions for the STOCK_QUOTE table.
MAX_PTN_TRADE	Maximum number of partitions for the STOCK_TRADE table.
ASE_SERVER	Name of the ASE server.
ASE_INTERFACE	Path to the ASE interfaces file.
ASE_UID	The ASE user ID.
ASE_PWD	The ASE user password.
ASE_DB_QUOTE	ASE database name for storing STOCK_QUOTE table data.
ASE_DB_TRADE	ASE database name for storing STOCK_TRADE table data.
IQ_ENGINE	Name of the IQ server.
IQ_UID	The IQ user ID.
IQ_PWD	The IQ user password.
IQ_DB_QUOTE	IQ database name for storing STOCK_QUOTE table data.
IQ_DB_TRADE	IQ database name for storing STOCK_TRADE table data (should be the same as IQ_DB_QUOTE.)

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 ASE, run the *SYBASE.sh* script, which is located in the ASE installation directory.

To set up the environment for the dbisql utility used by IQ, run the *ASIQ-12_6.sh* script, which is located in the directory */ASIQ-12_6* in the IQ installation directory.

Starting and stopping the load process

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

- Run *$RAP_HOME/rap_start_pl.sh* to start the load process.
- Run *$RAP_HOME/rap_stop_pl.sh* to stop the load process.

You can also create a tickler file to stop the load process instead of using the *rap_stop_pl.sh* script. The file name should be the same as $TICKLER_FILE, specified in the *set_rap_env_pl.sh*. When the load process goes away, the tickler file will be removed.

Note Before running the *rap_start_pl.sh* script 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 RAP cache and RAP repository the first time you start the automated load process.

Loading the TAQ sample data

You can either run a script or use the automated load process to load the TAQ sample data located in the *%Sybase%\Data\ASE\RAP_Input* directory on Windows. Before choosing either method described in this section, be sure you have unzipped and transferred the sample data *.csv* files as described in “Preparing the TAQ sample data” on page 26.

Using a script to load cache only

To load the TAQ sample data only into the RAP cache, run the *load_unix_ASE_TAQ_tables.sh* script from a command window on the machine that hosts the ASE database or on a client machine. This script uses the bulk copy utility bcp to load the sample data into the 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 *$SYBASE/$SYBASE_ASE/install/<servername>.log* on UNIX platforms.

Loading automatically

Use the automated load process to load the TAQ sample data into both the RAP cache and the RAP repository.

Remove any TAQ sample data already loaded in the RAP cache or you will receive "Duplicate row" errors. See the section "Clearing the cache" on page 37 for information on how to remove the data from the RAP cache.

If the TAQ sample data is not in the RAP cache, 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) which contain 500,000 rows each. This step is optional, but is recommended.

Note The number of rows specified in the split command may need to be adjusted to an appropriate value for the physical resources on your system to improve load performance.

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) which 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 Run *$RAP_HOME/rap_start_pl.sh* to start the automated load process, if the load process is not already running.

Loading additional TAQ sample data

This procedure is optional. The TAQ sample data provided for the RAP in-memory cache is also loaded into the RAP repository IQ database by the automated load process. The RAP repository TAQ sample data, however, contains millions of additional rows of data.

If you want to load more data than is included for the RAP cache, be sure you have already unzipped and transferred the RAP repository TAQ sample data files *STOCK_QUOTE.csv* and *STOCK_TRADE.csv* located in the *%Sybase%\Data\IQ\RAP_Input* folder on Windows as described in "Preparing the TAQ sample data" on page 20.

To run the script that loads the additional TAQ sample data into the RAP repository, you must have the 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 additional TAQ sample data into the RAP repository database.

Note If you want to preserve the TAQ data already in the repository, 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 Sybase IQ database, <*database_name*>*.iqmsg*. The default location of the IQ log file is the directory where the IQ database file (<*database_name*>*.db*) is located.

Checking the load process

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

RAP cache load

If an error occurs loading the RAP cache ASE database, the following message is displayed 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 RAP cache load failure.

RAP repository load

If an error occurs loading the RAP repository 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 RAP repository load failure.

Clearing the cache

Clear the data from the RAP cache on a regular schedule, as the data is permanently stored in the RAP repository and is redundant in the cache. The script *rap_truncate.sh*, located in the installation directory *%Sybase%\Scripts\RAP_Load_Scripts*, clears the data by truncating the STOCK_TRADE and STOCK_QUOTE tables in the RAP cache.

The script *rap_truncate.sh* script can be run manually or called from your own scheduling program. This script depends on the environment variables set by the *set_rap_env_pl.sh* as described in “Setting the environment variables” on page 33.

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

APPENDIX A

Sample ASE Configuration Settings

A sample Adaptive Server Enterprise configuration file *RAP_ASE_Server.cfg* is included with Risk Analytics Platform in the installation folder *%SYBASE%\Risk Analytics Platform 2.0\Scripts\RAP_Setup\ASE*. This appendix contains information about the settings of the options in the RAP configuration file for Adaptive Server Enterprise which 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 51.

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 or 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

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 will go 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 The value of 100M for buffer pool size is relatively small, in order to avoid using a large buffer pool and to have fewer large I/Os. 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

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 only works 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. Plus 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 = 1
   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. This function is not necessary, since there are mostly inserts and few deletes.

enable semantic partitioning = 1 Enables partitioning other than round-robin (for example, range partitioning) in ASE. A separate license is required for this option.

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 a lot of 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 avoid 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 ASE.

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 Shutdown 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 the 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 Shutdown 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 the 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>
```

APPENDIX B

Data Replication

About this Chapter

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

Contents

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 repository database.

Instructions in this chapter assume that you purchased a copy separately and are familiar with the Replication Server 12.6 environment. Instructions also assume that the target database is Adaptive Server Enterprise.

Replication environment

- Sybase Adaptive Server Enterprise version 12.5.3 (ASE) serves as the source database and the Replication Server System Database (RSSD).
- Sybase IQ version 12.6 serves as the target database.

- Sybase Replication Server version 12.6 (RepServer) captures and processes transactions on the source database and transfers them to the Sybase IQ 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 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 version. The information below should help clarify each version of Sybase Central and its use:

1 To manage IQ, 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 to configure 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 will manage 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 repository 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 12.5.3, which is required for compatibility with Replication Server 12.6, is not included in the Risk Analytics Platform package and must be purchased separately. See the *Release Bulletin Risk Analytics Platform* for the specific version requirements of Adaptive Server Enterprise.

Install Adaptive Server Enterprise as described in the ASE product documentation. Sybase strongly recommends that you always check the online support Web site for the latest software updates (ESD or EBF). See "Sybase EBFs and software maintenance" on page vi 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 of 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.

The steps you use to create a script for ASE are the same as those for Sybase IQ. Follow the steps in "Generating DDL Scripts" in the *Risk Analytics User's Guide* with the following differences for Adaptive Server Enterprise:

- Change the default Risk Analytics owner RAP_USER to the name of the Adaptive Server Enterprise database owner.
- In the Change the Target DBMS dialog, choose Sybase AS Enterprise 12.5.2 as the target database.
- Generate the DDL as described in "Generating DDL Scripts" in the *Risk Analytics User's Guide*, but do *not* select Create index in the Indexes panel on the Keys & Indexes tab.

- Follow the instructions in the next section to execute the DDL script on your Adaptive Server Enterprise database. Do not execute the DDL script as described in “Generating a DDL Script” as this procedure is for Sybase IQ only.

Executing the DDL script

Now the DDL script you generated from the PDM is ready to be executed. Log into your ASE target database and execute the DDL script to create the database tables and other database objects.

Note Risk Analytics Platform does not include indexes for an Adaptive Server Enterprise 12.5.3 source database.

Replication Server Configuration

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

Configuration tasks

This section summarizes that 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 is available only on Microsoft Windows platforms.

   - If setup asks you whether you want to over-write 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 will be named *Sybase REP_ <server name>*, where *<server name>* is the name you gave to your server.

RSM configuration

This section summarizes that 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 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. You will need to 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 will be named *Sybase RSMServer__<server name>*, where *<server name>* is the named 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 database. All procedures are performed in Interactive SQL. Separate startup 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, type the following command at a system prompt: dbisql.

2 Run the following commands:

```
grant connect to rsuser identified by rsuser

grant membership in group rs_systabgroup to rsuser
```

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. Separate. Follow 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/sybcentral* 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, choose Connect.

3 Add the appropriate connection information, click OK.

4 Expand Sybase IQ | *Sybase IQ <target database>* | Procedures & Functions.

5 Right-click on a procedure, choose Properties, click Permissions | Grant.

6 On the Grant Permissions dialog, choose *rs_systabgroup*, 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 data server. `rs_lastcommit` commits a transaction; `rs_threads` detects deadlocks and perform transaction serialization between parallel DSI threads.

1 Expand Sybase IQ | *Sybase IQ <target database>* | Tables.

2 Right-click on the appropriate table, choose Properties, click Permissions | Grant.

3 On the Grant Permissions dialog, choose *rs_systabgroup*, 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, click Permissions | Grant.

3 On the Grant Permissions dialog, choose *rsuser* (maintenance user) *or rs_systabgroup*, 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 onto 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 of the data servers using Open Client connectivity. Replication Server database connections are configured as either replicate-only, or primary-and-replicate (bidirectional) in a Sybase-only replication system.

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 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 in order for replication to work properly.
- See “Troubleshooting” on page 64 for information on problems that may be encountered 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 do this, 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 that is similar to this message: `"Replication definition <table name> requires at least one column in the primary key."`

Every replication definition requires 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, choose Open. A property sheet for the new subscription appears.

3 In the Name box, type a name for the new subscription.

4 In the Replication Information Group, click Browse, 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','2005-10-18')
go
insert into MARKET_INDEX values
    (20,'Test2','Test2','Test2','2005-10-18')
go
insert into MARKET_INDEX values
    (30,'Test3','Test3','Test3','2005-10-18')
go
```

2 Log into 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 the data was replicated properly:

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

Troubleshooting

This section contains some basic trouble shooting 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.

We also recommend 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 the Sybase IQ <*target database*> database comes down and replication does not occur. The RS log reports the following error:

```
T. 2004/11/08 11:06:52. (28): Command(s) to
'asiqsample.asiqsample':
T. 2004/11/08 11:06:52. (28): 'insert into test2 (id,
fname) values (1, 'George')'
E. 2004/11/08 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. 2004/11/08 11:06:52. THREAD FATAL ERROR #5049 DSI
EXEC(115(1) asiqsample.
asiqsample) - dsiqmint.c(3041)The DSI thread for
database 'asiqsample.asiqsample'
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. 2004/11/08 11:06:52. The DSI thread for database
'asiqsample.asiqsample'
```

is shutdown.

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 quickly replicate multiple source tables as a group. The process creates replication definitions for multiple source tables quickly. Below 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 and 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.

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

IQ connection remains in a DOWN state upon creation.

Symptom

RepServer log reports the following errors:

```
T. 2004/11/05 16:36:37. (36): 'execute
rs_update_lastcommit @origin = 0, @origin_qid =
0x00000000000000000000000000000000000000000000000000000000
0000000000000000000, @secondary_qid =
0x00000000000000000000000000000000000000000000000000000000
0000000000000000000, @origin_time = '19000101
00:00:00:000''
E. 2004/11/05 16:36:37. ERROR #1028 DSI(116
asiqdemo.asiqdemo) - \dsioqid.c(2028)
   Message from server: Message: 201, State 0, Severity
16 -- 'ASA Error -615: Parameter '' not found in
procedure '???''.
E. 2004/11/05 16:36:37. ERROR #5046 DSI(116
asiqdemo.asiqdemo) - \dsioqid.c(2042)
   When executing the rs_get_lastcommit function in
database 'asiqdemo.asiqdemo', received data server
errors. See
logged data server errors for more information.
I. 2004/11/05 16:36:37. The DSI thread for database
'asiqdemo.asiqdemo' is shutdown.
```

Solution

The maintenance user (*rsuser*) must be a member of the rs_systabgroup group. To accomplish this, the IQ administrator must execute the following from an Interactive SQL or isql session:

```
grant membership in group rs_systabgroup to rsuser
```

Or the IQ administrator may use IQ's Sybase Central Java Edition. Highlight the maintenance user (*rsuser*) and select File | New | Memberships from the menu. Then select the *rs_systabgroup* group for the rsuser to join.

Enabling RepServer transact SQL tracing

Log into RepServer using the isql utility from a command prompt, and execute the following:

```
trace "ON", DSI, DSI_BUF_DUMP
```

This will trace the SQL transactions within the RepServer and write them to the RepServer log.

Enabling IQ tracing

Add the following parameters to the IQ startup command line:

```
-zr all -zo
```

Where is the path and name for the trace log file. This will turn on extensive tracing within the Sybase IQ server. Additionally, the standard Sybase IQ database log can be checked. This is located in the *.iqmsg* file.

Rounding with the FLOAT datatype

Minor rounding occurs after replicating FLOAT datatypes from ASE to Sybase IQ. See the example below:

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

Defining a new connection using RSM might fail with error

If this error occurs grant execute permission on the *rs_configure* stored procedure within the ASE RSSD database. This can be done via Sybase Central or manually via isql.

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 *d:\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 8 MB.You need to add more dbspace.

To add more dbspace, execute the following command in Interactive SQL:

```
CREATE DBSPACE asiqdemo4 AS 'd:\sybase\asiqdemo4.iq'IQ
STORESIZE 500
```

This adds 500 MB to the dbspace.

Index

L

M

O

P

R

S

T