SYBASE®
An SAP® Company

Installation and Configuration Guide

SAP Sybase RAP 5.0

DOCUMENT ID: DC00793-01-0500-02
LAST REVISED: December 2012

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

Contents

CHAPTER 1 Planning Your Installation

Review the latest information for this release and ensure that your environment is ready to install.

1. Review the *RAP Users Guide* to familiarize yourself with the current SAP® Sybase® RAP architecture.
2. Consult the documentation for the product components (SAP HANA (or Adaptive Server® Enterprise), Sybase IQ, Event Stream Processor, and Sybase Control Center) you are installing to ensure you:
   a) Obtain the latest information on the components.
   b) Plan your client and server installations appropriately.
   c) Have sufficient available disk space to install the RAP components. See the *Disk Space Requirements* topic.
   d) Have an operating system that meets the version level and network protocol requirements for your platform.
   e) Have a login with appropriate permissions for all the RAP installation tasks.
3. Install operating system patches, if required. If you are having trouble starting an installer, make sure you have the required operating system patches for the Java Runtime Environment (JRE).
4. If you use a SySAM license server, see the *Sybase Software Asset Management (SySAM) 2 Users Guide* for information on how to run it, and how to install the license onto the license server.
5. Verify that your network software is configured. Sybase software uses network software even if the SAP HANA (or Adaptive Server) and Sybase IQ client applications are installed on a machine that is not connected to a network.
6. To verify your network configuration, ping the host.
7. On your Adaptive Server Enterprise server machine, allocate a shared memory segment that is at least as large as the Adaptive Server total logical memory configuration parameter.
8. Ensure that the system clocks on all machines in the RAP system have the same time.

Supported Operating Systems

The primary server components of RAP run on specific platforms and operating systems.

Check the Platform Certifications *http://certification.sybase.com/ucr/search.do* Web site to see whether additional operating systems are supported.

Table 1. Supported Platforms and Operating Systems

Platform	Supported OS
Linux, 64-bit (AMD/Intel)	Red Hat 5.5, 6.0
Sun, 64-bit (SPARC, UltraSPARC T)	Solaris 10
Sun, 64-bit (AMD)	Solaris 10

Software Included with Your RAP License

A RAP license entitles you to various software components, including Adaptive Server Enterprise, Sybase IQ, PowerDesigner, and Event Stream Processor.

RAP supports SAP HANA as the preferred in-memory database. If you do not have SAP HANA installed, you can use Adaptive Server Enterprise as the in-memory database. For information on where to obtain an SAP HANA client installation, see the *SAP HANA Client Installation Guide*. For instructions on setting up the SAP HANA server as your in-memory database for RAP, see *Setting Up SAP HANA as the In-Memory Database* on page 7.

For instructions on setting up Sybase Control Center for RAP, see *Setting Up Sybase Control Center for Monitoring and Administering* on page 27.

Table 2. Installation Media

Media	Description
RAP Product DVD	Includes the RAP enablement package which contains configuration scripts and samples.
Adaptive Server Enterprise 15.7	Includes the product and client DVDs, which install the Adaptive Server Enterprise server and the components required for client connections to an Adaptive Server Enterprise server. Before installing any product, review the `readme.txt` file, which contains a brief description of products and dependencies, as well as last-minute instructions or changes.
Sybase IQ 15.4	Includes the product DVD and client CDs, which install the Sybase IQ server and the components required for client connections to a Sybase IQ server.

Media	Description
PowerDesigner CDs	Install PowerDesigner modeling environment for use with RAP. If you purchased RAP from Sybase or an authorized Sybase reseller, it comes with PowerDesigner DataArchitect Enterprise 16.1. If you ordered RAP under an SAP® contract and were directed to download from SAP Service Marketplace (SMP), it comes with PowerDesigner Information Architect 16.1. Both versions provide PowerDesigner functionality needed for RAP.
Event Stream Processor 5.1 CDs	Install Event Stream Processor for use with RAP. Includes Windows x64 and x86 versions, as well as CDs for the platform and operating system on which you are installing RAP. The Windows CDs install the PowerDesigner extensions for RAP. First, install PowerDesigner using the media provided above, and then install the PowerDesigner extensions using the Event Stream Processor custom installation mode.

Disk Space Requirements

Review the disk space requirements to plan for your installation.

Typically, a production environment has four servers: one each for Sybase IQ, SAP HANA (or Adaptive Server Enterprise), Event Stream Processor, and Sybase Control Center for RAP. Memory requirements vary depending on the size, number, and complexity of projects you run. See the individual Sybase IQ, SAP HANA (or Adaptive Server Enterprise), and Event Stream Processor documentation for memory recommendations.

Component	Disk Space
RAP Enablement Package	1.2GB
Event Stream Processor	4GB
Sybase IQ Database Server	8GB
SAP HANA Client	0.4GB
Adaptive Server Enterprise Database Server	10GB
Sybase Control Center for Sybase RAP	4GB

Obtaining Licenses

Before installing your product, choose a Sybase Software Asset Management (SySAM) license model, determine license server information, and obtain license files.

Sybase RAP is licensed through SySAM and supports both served and unserved licenses. Sybase recommends the use of served license models with RAP so that the RAP package license, which includes Event Stream Processor, Sybase IQ and Adapter Server Enterprise licenses, as well as other RAP licenses (the PowerDesigner floating license), may be shared by all RAP machines. Because the primary RAP package license includes these other license keys, you do not need to generate additional individual license keys for these components, unless you purchased add-on or optional licenses separately.

If you are migrating to RAP 5.0 from a previous version of RAP, you may need to obtain additional license files to preserve your existing entitlements. See the *Migration Information* section in the *Sybase RAP Release Bulletin* for more information on obtaining these additional licenses.

These steps summarize the actions required to install a Sybase Software Asset Management (SySAM) license. Refer to the *Sybase Software Asset Management Users Guide* for details.

1. Choose a SySAM license model.

License model	Description
Served license model	A license server manages the allocation of licenses among multiple machines. Sybase recommends the use of served license models with RAP.
Unserved license model	Licenses are obtained directly from a license file. If you are using an unserved license, save the license file to the machine on which you install the product.

2. For the served license model, decide whether to use an existing or new license server.

   The license server and your product installation do not need to share the same machine, operating system, or architecture.
3. If you chose the served license model, do one of the following:
   - Install the new license server on a machine that does not have an existing license server.
   - To install on a machine that is running a SySAM 1 license server, follow migration instructions in the *Sybase Software Asset Management User's Guide* to migrate to SySAM 2.
4. Before installing your product:
   - If you purchased your product from Sybase or an authorized Sybase reseller, go to the secure Sybase Product Download Center (SPDC) at *https://sybase.subscribenet.com*

and log in to generate license keys. The license generation process may vary slightly, depending on whether you ordered directly from Sybase or from a Sybase reseller.

- If you ordered your product under an SAP contract and were directed to download from SAP Service Marketplace (SMP), you can use SMP at *http://service.sap.com/licensekeys* to generate license keys for Sybase products that use SySAM 2-based licenses.

5. From SPDC or SMP, download the appropriate license keys for the following:

   - Base IQ VLDB Option for RAP - download the license file for the amount of terabytes of VLDB provided as a base to all RAP licenses.
   - RAP - The Trading Edition - download the RAP software and RAP package license.
   - PowerDesigner for RAP - download the software and license for PowerDesigner. One floating license is provided with RAP.
   - If applicable, download the license keys for any additional options or add-ons purchased for Sybase RAP.

Next

For more information on SySAM licensing, see the *Sybase Software Asset Management Users Guide* on the Sybase Product Manuals Web site at *http://sybooks.sybase.com*.

See also

- *Activating Your SySAM Sybase IQ License* on page 61
- *Activating Your SySAM Adaptive Server Enterprise License* on page 59

CHAPTER 2 Install and Configure RAP

Install and configure the in-memory and historical databases, Event Stream Processor, Sybase Control Center, and deploy the RAP sample project for SAP HANA (or Adaptive Server).

Each component is installed and configured individually. If you wish to copy and paste commands directly from the guide, use the HTML version of the guide to avoid integrating line breaks into the commands.

The terms RAPCache and in-memory database, and RAPStore and historical database are used interchangeably throughout the RAP documentation. RAP consists of SAP HANA as the preferred in-memory database and Sybase IQ as the historical database. If you do not have SAP HANA installed, you can use Adaptive Server Enterprise as the in-memory database.

Set Up the In-Memory Database

Set up and configure an in-memory database for RAP. The preferred in-memory database for RAP is SAP HANA but if you do not have SAP HANA, you can use Sybase Adaptive Server Enterprise.

Do not use spaces in the installation directories for any of the RAP components.

Setting Up SAP HANA as the In-Memory Database

Install and configure SAP HANA for use as the in-memory database for Sybase RAP.

Prerequisites

Install the SAP HANA Client. For more information, see the *SAP HANA Client Installation Guide*.

Task

1. Open up a new command prompt. Do not reuse a command prompt from the installation of another RAP component.
2. Install revision 1.00.37 or higher of the SAP HANA database server.
3. Copy the RAP enablement package to the SAP HANA server machine, and extract it into a directory of your choice. This is referred to as the RAP enablement installation directory throughout the rest of this guide.
4. Run the `configure.sh` shell script from the RAP enablement installation directory.

   This file prepares the SAP HANA script files for execution.
5. Use the **hdbsql** tool to create a RAP_USER in the SAP HANA database.

a) Run **hdbsql** in interactive mode as the administrative user: **hdbsql -n <server name>:<port> -u <username>** where **<username>** corresponds to an SAP HANA account with user administrative permissions and grantable import permissions.

   You are asked for this user's password when you run the first SQL command in this tool.

b) In the interactive terminal, run: **CREATE USER RAP_USER PASSWORD <new password>**

   The password you set here is only temporary. You are required to change your password after your first log onto the SAP HANA server as RAP_USER.

c) In the interactive terminal, run: **GRANT IMPORT TO RAP_USER**

d) Type **quit** to leave the interactive terminal.

6. Use the **hdbsql** tool to set up the RAP schema and sample data.

a) Run **hdbsql** in interactive mode as the user (RAP_USER) created in the previous step: **hdbsql -n <server name>:<port> -u RAP_USER** .

   You are prompted for RAP_USER's password after you run the first SQL command in this tool.

b) In the interactive terminal, run: **ALTER USER RAP_USER PASSWORD <new password>** where **<new password>** is the new and permanent password for RAP_USER.

c) Create the RAP tables in the SAP HANA database. Before loading the database schema, decide whether the RAP system will need to support multibyte character sets for string data.

   1. If the RAP system requires support for multibyte character sets, load the multibyte RAP schema. In the interactive terminal, run: **\I <RAP enablement installation directory>/databases/hana/scripts/mbcs/create_tables_mbcs.sql** where **<RAP enablement installation directory>** is replaced by the absolute path to the RAP enablement installation directory.
   2. Otherwise, load the single-byte RAP schema. In the interactive terminal, run: **\I <RAP enablement installation directory>/databases/hana/scripts/create_tables.sql** where **<RAP enablement installation directory>** is replaced by the absolute path to the RAP enablement installation directory.

d) Load the sample supporting data into the newly created tables in the SAP HANA database. The sample supporting schema is required to use the RAP sample project.

   If you installed the multibyte schema in the previous step, you can choose to load the default sample supporting data for English (single-byte characters) or for Japanese (multibyte characters). If you install the Japanese sample supporting data, use the Japanese sample input files when you run the RAP project.

Sample Supporting Data	Required Steps
English	1. In the interactive terminal, run: **\I <RAP enablement installation directory>/databases/hana/scripts/load_tables.sql**
Japanese (multibyte schemas only)	1. Change directories to `<RAP enablement installation directory>/databases/hana/data/supportingdata`. 2. To keep a copy of the English sample supporting data for future use, rename the files `INSTRUMENT.csv` to `INSTRUMENT.csv.en`, and `MARKET_INDEX.csv` to `MARKET_INDEX.csv.en`. 3. Copy the `INSTRUMENT.csv.jp` to `INSTRUMENT.csv`, and `MARKET_INDEX.csv.jp` to `MARKET_INDEX.csv`. 4. In the interactive terminal, run: **\I <RAP enablement installation directory>/databases/hana/scripts/load_tables.sql**

e) (Optional) Load sample data (STOCK_QUOTE and STOCK_TRADE) into the RAP database. In the interactive terminal, run: **\I <RAP enablement installation directory>/databases/hana/scripts/load_TAQ_tables.sql**

f) Type **quit** to leave the interactive terminal.

Setting Up Adaptive Server Enterprise as the In-Memory Database

Install and configure Adaptive Server for use as the in-memory database for RAP if you are not using SAP HANA as your in-memory database.

1. Open a new command prompt. Do not reuse a command prompt from the installation of another RAP component.
2. Install Adaptive Server Enterprise, Enterprise Edition version 15.7 ESD #1 or higher.

   The installer may display additional options other than the options below but the suggestions detailed in this step are important when configuring Adaptive Server Enterprise for RAP. If no suggestion is made for an option, you can accept a default value or see the *Adaptive Server Enterprise Installation Guide* for additional information.

   a) Choose a typical installation. If you are operating your server in a non-English server environment, you may wish to install additional Adaptive Server language packs. To do this, choose a custom installation.

   b) When prompted, choose to install a licensed copy and continue the installation without a license key. Or if you already have a licensed server set up, choose the licensed server.

   c) Select **Enterprise Edition** for the product edition, and select the license type that corresponds to your RAP license.

d) After copying files to your computer, the installer prompts you to specify which products you wish to configure. Select the Adapter Server, Backup Server, and the Sybase Control Center.

The next steps detail the configuration of each of these options.

3. Configure the ASE server.

a) Choose a server name.

The default server name is RAPASE. This name is used in all of the **isql** commands used to connect to the database. If you use a different server name, replace RAPASE in the **isql** commands with the name of your server (the **-S** option specifies the server name).

b) Choose **mixed** as the application type.

c) Choose a page size.

Sybase recommends an 8k page size. If you choose a different page size, you will need to make some modifications to SQL scripts later in the installation process.

d) Choose a language and character set.

If you wish to process multibyte characters, select UTF-8 (Unicode) for the character set and binary (binary ordering for UTF-16) as the sort order. Otherwise, accept the default of iso_1 and a sort order of bin_iso_1. The character set chosen determines which SQL scripts are required when configuring your RAP schema.

e) Select the option to optimize your server, and allocate as much shared memory to the ASE server as possible. If you are planning on creating an in-memory database for this server, ensure that at least 5 gigabytes of shared memory are allocated to the server.

f) Select the option to create sample databases.

g) These are recommended minimum database configuration options:

Property	Recommended Configuration
Master Device Size	500MB
Master Database Size	100MB
System Procedure Device Size	160MB
System Procedure Database Size	160MB
System Device Size	30MB
System Database Size	30MB
Tempdb Device Size	100MB
Tempdb Database Size	100MB

h) Do not enable the PCI property.

4. Configure the Backup Server. Sybase recommends using the server name RAPASE_BS but if you chose a different name for your database server, use that name in your backup server instead.

5. Configure Sybase Control Center.

   This installs the agent that communicates with SCC. The actual SCC server is installed using the Sybase Control Center installer.

   If you are installing other Sybase components on this machine, ensure that the RMI port for the Adaptive Server agent is unique from the RMI ports specified for other Sybase components. The default for all Sybase components is 9999.

   a) Configure the UDP Adaptor for discovery purposes.
   b) Change the passwords for SCC users (uafadmin, sccadmin, and sccuser) by selecting each one and changing their properties to reflect their new passwords.

6. Wait for the ASE installer to finish.

   Once the installation is complete, the install tool starts the ASE server and creates the specified databases. The following scripts primarily modify the properties and schema of the Adaptive Server model database. This database is copied to create the RAP database.

7. Copy the RAP enablement package to the ASE server machine, and extract it into a directory of your choice. This is referred to as the RAP enablement installation directory throughout the rest of this guide.

8. Run the `configure.sh` shell script from the RAP enablement installation directory.

   This file prepares the Adaptive Server script files for execution.

9. Source the `SYBASE.sh` or `SYBASE.csh` script file located in the RAP enablement installation directory.

10. Source the `SYBASE.sh` or `SYBASE.csh` script file located in the Adaptive Server installation directory.

11. Customize the creation of the rapTAQ database. Edit the `$RAP_HOME/databases/ase/scripts/create_TAQdatabase.sql` file.

    a) Create separate locations to store your data devices if they do not already exist. For performance reasons, consider storing your data devices on multiple disks.
    b) Change all instances of <ADD PATH> in the script to these locations.
    c) If you chose a different page size from the recommended 8k page size for RAP, create a memory pool for the log cache since by default, a log cache has an i/o size that is twice the page size. Modify the line `sp_poolconfig "rapTAQLogCache", "10M", "16K"`. Change the "`16K`" portion to be twice the page size selected.

       For example, with a 2K page size that line becomes `sp_poolconfig "rapTAQLogCache", "10M", "4K"`.

12. Configure initial database properties.

    a) Run the command **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/initial_setup.sql**.

       This script modifies database properties including locks, number of connections, transfer history, and maximum memory.

b) If errors display which refer to a procedure not being found, ignore them. The script verifies whether a procedure exists before it creates one.

13. Create and load the RAP tables with supporting data.

a) Create the tables.

1. If you elected to support multibyte characters in the Adaptive Server setup, load the multibyte schema by running **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/mbcs/create_tables_mbcs.sql** . This adds necessary memory to the model and temporary databases. Multibyte versions of the RAP tables are then created.
2. Otherwise, run the command **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/create_tables.sql**. This adds necessary memory to the databases and creates all RAP tables.

b) Load the sample supporting data into the newly created tables in the Adaptive Server database. The sample supporting data is required to use the RAP sample project.

1. If you installed the multibyte schema in the previous step, you can choose to load the default sample supporting data for English (single-byte characters) or for Japanese (multibyte characters). If you install the Japanese supporting data, use the Japanese sample input files when you run the RAP project.

Sample Supporting Data for Multibyte Schemas	Required Steps for Multibyte Schemas
English	1. Load data into the supporting tables by running **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/mbcs/load_tables_mbcs.sql**
Japanese	1. Change directories to `$RAP_HOME/databases/ase/data/supportingdata/mbcs`. 2. To keep a copy of the English sample supporting data for future use, rename the files `INSTRUMENT.bin` to `INSTRUMENT.bin.en` and `MARKET_INDEX.bin` to `MARKET_INDEX.bin.en` 3. Copy the `INSTRUMENT.bin.jp` to `INSTRUMENT.bin` and `MARKET_INDEX.bin.jp` to `MARKET_INDEX.bin`. 4. Load data into the supporting tables by running **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/mbcs/load_tables_mbcs.sql**

2. If you installed single-byte schemas, run the command **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/load_tables.sql** to load data into the supporting tables.

14. Create the three RAP SQL procedures by running: **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/create_procedures.sql**

15. Create RAP_USER.
    a) Run **isql** in the interactive mode: **isql -S RAPASE -U sa**
    b) Enter these lines in the terminal where <password> represents the password being assigned to RAP_USER:

```
sp_addlogin RAP_USER,<password>,model
go
use model
go
sp_adduser RAP_USER,rap_user
go
use master
go
sp_adduser RAP_USER,rap_user
go
```

    c) Type **quit** to leave the interactive terminal.
16. Create the rapTAQ database and bind the log cache. You can create the rapTAQ database as a reduced durability database (RDDB; default) or an in-memory database (IMDB; optional):
    a) To create the RDDB:
        1. Execute the command: **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/create_TAQdatabase.sql**

           While the create database script is running, ignore any errors relating to devices or databases not existing. These are due to the script attempting to drop any current devices and databases before creating new ones.
        2. Run the command: **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/bind_TAQcache.sql**

    This script binds the model database's `syslog` table to a recently created log cache. A separate cache improves performance by isolating log I/O operations from data I/O operations.
    b) To create an in-memory database, execute the command **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/create_memoryTAQdatabase.sql**

    This script attempts to drop the TAQ cache before it creates it. If you are running this script for the first time, you will see errors about being unable to drop the named cache. You can safely ignore these errors.
17. Set RAP_USER's default database and modify permissions on the model database by running the command: **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/modify_user.sql**
18. Shutdown the ASE server so that the changes that have been made to the model database during installation are copied into the rapTAQ database. Run: **isql -S RAPASE -U sa** and then type:

```
shutdown
go
```

19. License your ASE database using your RAP license if you have not already done so using the installer.
20. Start your ASE server by issuing these commands in a new terminal:
    a) Source the `SYBASE.sh` or `SYBASE.csh` script in the Adaptive Server installation directory.
    b) Run: **<ASE Install Location>/ASE-15_0/install/RUN_RAPASE**. If you chose your own server name earlier, this file will have a different name.
21. (Optional) Load your new database with sample trade and quote data. In the terminal where you previously sourced the `SYBASE.sh` or `SYBASE.csh` files for the RAP enablement package and Adaptive Sever:
    a) If you are using the multibyte schema, load data into the TAQ tables by running **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/mbcs/load_TAQ_tables_mbcs.sql**
    b) If you are using a single-byte schema, load data into the TAQ tables by running **isql -S RAPASE -U sa -i $RAP_HOME/databases/ase/scripts/load_TAQ_tables.sql**

Setting Up the Historical Database

Set up Sybase IQ as the historical database in RAP.

Some of the installation steps require that you use **dbisqlc**. After running a **dbisqlc** command, an error message displays stating "Invalid user ID or password." Click **OK**, then in the next window, fill in the password for the user specified, and select **OK** to run the script. This is necessary to prevent passing your password on the command line since this is undesirable for security reasons.

Do not use spaces in the installation directories for any of the RAP components.

1. Open a new command prompt. Do not reuse a command prompt from the installation of another RAP component.
2. Install Sybase IQ 15.4 ESD#1 or higher.

   The installer may display additional options other than the options below but the suggestions detailed in this step are important when configuring Sybase IQ for RAP. If no suggestion is made for an option, you can accept a default value or see the *Sybase IQ Installation Guide* for additional information.
   a) Choose a typical installation.
   b) When prompted, install a licensed copy of the server but select the option that allows you to continue the installation without a license key because the RAP license will not be accepted until further steps are performed. Or if you already have a licensed server set up, choose the licensed server.
   c) Select **Enterprise Edition** for the product edition, and select the license type corresponding to your RAP license.

d) Install all RAP components on separate machines. If other Sybase components exist on this machine, ensure that the RMI port specified in the installer is unique from the RMI ports specified for other Sybase components. The default for all Sybase components is 9999.
e) Change the Sybase Control Center password for the uafadmin user if the installer you are using gives you this option. Otherwise, follow the instructions in the *Setting Passwords or Disabling Default Login Accounts* topic in the *Sybase Control Center 3.2.7 Installation Guide* to change the password for this user once you have completed your installation.

   The SCC portion of this installer installs the agent that communicates with SCC. The actual SCC server is installed using the Sybase Control Center installer.
f) When asked whether you wish to start a Sybase Control Center Server, decline.

3. Copy the RAP enablement package to the Sybase IQ server machine, and extract it into a directory of your choice.
   If any of the RAP components from this current RAP version are installed on this machine, you can use the existing RAP enablement installation. This is referred to as the RAP enablement installation directory throughout the rest of this guide.
4. If you did not run it previously, run the `configure.sh` shell script from the RAP enablement installation directory.
   This file prepares the Sybase IQ script files for execution.
5. Source the `SYBASE.sh` or `SYBASE.csh` script file located in the RAP enablement installation directory.
6. To be able to use the Sybase IQ utilities required for setting up the database, source the `IQ.sh` or `IQ.csh` files in `<Sybase IQ Install Directory>`.
7. Create the RAP database for your Sybase IQ server. If you set up your in-memory database to support multibyte characters, set up the historical database to also support multibyte characters. Note that some of the commands below have line spacing added for the sake of clarity, and that you should remove these line spaces before running the commands.
   a) If your in-memory database was set up to support multibyte characters, create a multibyte enabled database. Create a directory to store your database files if it does not already exist. Then, replace the three `<Sybase IQ Database Location>` placeholders with this location, and run the command:

      ```
      dbisqlc -c "uid=dba;eng=myserver;dbn=utility_db"
      "CREATE DATABASE '<Sybase IQ Database Location>/RAPIQ'
      IQ SIZE 500
      COLLATION 'UTF8BIN'
       IQ PATH '<Sybase IQ Database Location>/RAPIQ.iq'
      TEMPORARY PATH '<Sybase IQ Database Location>/RAPIQ.iqtmp'
      TEMPORARY SIZE 1000;"
      ```

   b) Otherwise, set up a single-byte historical database. First, replace the three `<Sybase IQ Database Location>` placeholders with a file system location in which to store your database files. Then, run the command:

      ```
      dbisqlc -c "uid=dba;eng=myserver;dbn=utility_db"
      "CREATE DATABASE '<Sybase IQ Database Location>/RAPIQ'
      ```

```
IQ SIZE 500
IQ PATH '<Sybase IQ Database Location>/RAPIQ.iq'
TEMPORARY PATH '<Sybase IQ Database Location>/RAPIQ.iqtmp'
TEMPORARY SIZE 1000;"
```

8. Customize the configuration of the RAP database.
   a) Make a copy of the `default.cfg` file located in the `<Sybase IQ Install Directory>/IQ-15_4/scripts` directory, place it into the `<Sybase IQ Database Location>`, and rename it `RAPIQ.cfg`.

      The IQ-15_4 portion of the directory may change if you are using a newer version of Sybase IQ.
   b) Edit the `RAPIQ.cfg` file with the following changes:

      Use the default Sybase IQ server port of 2638, if possible. If you need to use a different port, update all of the **dbisqlc** commands in these installation instructions to use your new port and change the line **-x tcpip(port=2638)** to **-x tcpip(port=<new port>)**.

      1. Add these lines to name your server, allow any user to stop the database, assign a port for the server to start on, and set the main and temp cache size (MB):

         ```
         -n RAPIQ
         -gk all
         -x tcpip(port=2638)
         -iqmc 1000
         -iqtc 1000
         ```

         If a server called RAPIQ already exists on the same subnet as this machine, choose another name for this server. Change the **-n RAPIQ** line to **-n <Server name>**. Also change **eng=RAPIQ** to **eng=<Server name>** in all the **dbisqlc** connections.
      2. Increase the number of connections the Sybase IQ server allows by changing the line that says “-gm 10” to “-gm 100”.
9. Change the directory to the `<Sybase IQ Database Location>` and start the RAP database using the command: **start_iq @RAPIQ.cfg RAPIQ.db**
10. Create a dbspace in which to store the database data. Run: **dbisqlc -c "uid=DBA;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/create_dbspace.sql**
11. Create RAP_USER.
    a) Run **dbisqlc** in interactive mode: **dbisqlc -c "uid=DBA;eng=RAPIQ;links=tcpip(host=localhost:2638)"**
    b) Replace the `<RAP_USER Password>` placeholder with a password for RAP_USER, and run: **grant connect to RAP_USER IDENTIFIED BY <RAP_USER Password>**
12. Create the RAP schema in the Sybase IQ database:
    a) If you created a historical store with multibyte support, run this command: **dbisqlc -c "uid=DBA;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/mbcs/create_tables_mbcs.sql**

b) Otherwise, for single-byte only support, run this command: **dbisqlc -c "uid=DBA;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/create_tables.sql**

13. Load the sample supporting data into the newly created tables in the Sybase IQ database. The sample supporting schema is required to use the RAP sample project.

    If you installed the multibyte schema in the previous step, you can load the default sample supporting data for English (single-byte characters) or for Japanese (multibyte characters). Otherwise, install the single-byte sample data (English). If you install the Japanese sample supporting data, use the Japanese sample input files when you run the RAP project.

Sample Supporting Data	Required Steps
English	1. Run this command: **dbisqlc -c "uid=RAP_USER;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/load_tables.sql**
Japanese (multibyte schemas only)	1. Change directories to `$RAP_HOME/databases/iq/data/supportingdata`. 2. To keep a copy of the English sample supporting data for future use, rename the files • `INSTRUMENT.csv` to `INSTRUMENT.csv.en` • and `MARKET_INDEX.csv` to `MARKET_INDEX.csv.en` 3. Copy the `INSTRUMENT.csv.jp` to `INSTRUMENT.csv` and `MARKET_INDEX.csv.jp` to `MARKET_INDEX.csv`. 4. Run this command: **dbisqlc -c "uid=RAP_USER;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/load_tables.sql**

14. (Optional) Load sample data (STOCK_QUOTE and STOCK_TRADE) into the RAP database by running this command: **dbisqlc -c "uid=RAP_USER;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/load_TAQ_tables.sql**
15. Grant required privileges to RAP_USER by executing this command: **dbisqlc -c "uid=DBA;eng=RAPIQ;links=tcpip(host=localhost:2638)" $RAP_HOME/databases/iq/scripts/set_permissions.sql**
16. License your Sybase IQ server using your RAP license if you have not already done so using the installer.
17. Stop your Sybase IQ server using the command **stop_iq**.
18. Restart your Sybase IQ server by changing the directory to the `<Sybase IQ Database Location>` and running: **start_iq @RAPIQ.cfg RAPIQ.db**

Set Up Event Stream Processor for Streaming

Setting up Event Stream Processor for streaming involves installing Event Stream Processor, configuring ODBC, configuring your Event Stream Processor environment, and configuring your `service.xml` file.

Do not use spaces in the installation directories for any of the RAP components.

Installing Event Stream Processor

Install Event Stream Processor for use within a RAP system. See the *Event Stream Processor Installation Guide* for instructions on installing Event Stream Processor from the command line.

1. Open a new command prompt. Do not reuse a command prompt from the installation of another RAP component.
2. Install Event Stream Processor 5.1 or later (Event Stream Processor 5.1 Service Pack 1 is required when using SAP HANA).

   The installer may display additional options other than the options below but the suggestions detailed in this step are important when configuring Event Stream Processor for RAP. If no suggestion is made for an option, you can accept a default value or see the *Event Stream Processor Installation Guide* for additional information.

   a) Choose a custom installation. Set a custom password (for example, for the keystore, agent, or node) whenever the installer gives you this opportunity.
   b) Select to install all default packages.
   c) When prompted, choose to install a licensed copy of Event Stream Processor but continue the installation without a license key. Or if you already have a licensed server set up, choose the licensed server.
   d) During the cluster configuration portion of the installation:
      1. Provide the fully qualified host name of this machine.
      2. Choose to enable or disable SSL. If you enable SSL, modify the ESP Server URL in all of the commands throughout this guide to "esps://" instead of "esp://".
      3. Choose a custom cache port.
      4. Use native OS authentication to secure your cluster.

   The installer prompts for information regarding the Sybase Control Center (SCC) server. This SCC server is not your primary SCC server, so the port values relating to the server are not important but ensure you specify a port that is not in use. This installs the agent that communicates with SCC. The actual SCC server is installed using the Sybase Control Center installer.

   The installer completes copying files.

e) If you are installing other Sybase components on this machine, ensure that the RMI port for the Event Stream Processor agent is unique from the RMI ports specified for other Sybase components. The default for all Sybase components is 9999.
f) If the installer prompts you to start the Sybase Control Center service, decline.
g) If the installer prompts you to start the ESP Studio, decline.

Configuring ODBC for Event Stream Processor

Set up ODBC and the Sybase IQ ODBC driver to allow Event Stream Processor to connect to the historical database.

1. The Sybase IQ ODBC driver is installed using the Sybase IQ Network Client Install. Use a typical installation and install a licensed copy.
2. If you plan on using the SAP HANA project, obtain and install the unixODBC ODBC 2.3.0 driver manager on your system. However, if you do not plan to use the SAP HANA project, you can use the driver manager that is installed as part of Sybase IQ Network Client Installation.
3. Event Stream Processor requires you to call your driver manager's library `libodbc.so.1`. If your driver manager has a different library name (typically `libodbc.so`), create a symbolic link to the driver manager library using the name libodbc.so.1.
4. Create or add to an existing `.odbc.ini` file. This file must be accessible to any user that has permission to start Event Stream Processor.

   Sybase suggests placing the file inside the Event Stream Processor installation directory.

   Add an entry for the previously installed Sybase IQ server by replacing `<host name>` with the hostname of the Sybase IQ database server, and `<Sybase IQ ODBC Driver Installation Directory>` with the installation path for the Sybase IQ ODBC driver.

   If you specified a different port for you server, replace `2638` with your port.

   ```
   [RAPIQ]
   ServerName=RAPIQ
   DatabaseName=RAPIQ
   CommLinks=tcpip{host=<host name>;port=2638}
   Driver=<Sybase IQ ODBC Driver Installation Directory>/IQ-15_4/lib64/libodbc.so
   Charset=utf8
   ```

   Include **Charset=utf8** only if you have a multibyte character install.
5. If you plan to use the SAP HANA sample project with Event Stream Processor, add connection information about your SAP HANA database to the `.odbc.ini` file you previously created or modified.

   Replace `<hana host name>`, `<hana port>`, and `<path to SAP hana driver library file>`. If you have not yet installed your SAP HANA client, update `<path to SAP hana driver library file` with the correct location after you install it. Then add the following to your `.odbc.ini` file:

```
[RAPHANA]
ServerNode=<hana host name>:<hana port>
Driver=<path to SAP hana driver library file>/libodbcHDB.so
CHAR_AS_UTF8=1
```

Include **CHAR_AS_UTF8=1** only if Event Stream Processor passes multibyte characters to the SAP HANA database. If no multibyte characters need to be processed, you may omit this line.

Configuring Event Stream Processor

Configure your Event Stream Processor environment.

1. Edit the `SYBASE.sh` and `SYBASE.csh` files in the Event Stream Processor installation folder so that Event Stream Processor can find your driver manger and your driver manager can find the Sybase IQ drivers. In each file:
   a) Set the environment variable ODBCINI to the absolute path of your `.odbc.ini` file.
      - Add these lines to the bottom of the `SYBASE.sh` script:
        - **ODBCINI=<Path to ODBC INI File>/.odbc.ini**, where **<Path to ODBC INI File>** is the absolute path to the `.odbc.ini` file created or modified in the previous step.
        - **export ODBCINI**
      - Add a line to the bottom of the `SYBASE.csh` script: **setenv ODBCINI <Path to ODBC INI File>/.odbc.ini** where **<Path to ODBC INI File>** is the absolute path to the `.odbc.ini` file created or modified in the previous step.

   b) Add the absolute path to your driver manager libraries and the library symbolic link to the beginning of the LD_LIBRARY_PATH.
      - Add these lines to the bottom of the `SYBASE.sh` script:
        - (Linux and Solaris) **LD_LIBRARY_PATH="<path to lib folder in driver manager>":"<path to IQ Client install location>/IQ15_4/lib64": $LD_LIBRARY_PATH**, where **<path to lib folder in driver manager>** is the absolute path to the directory which contains the ODBC driver manager's libraries.
        - (Linux and Solaris) **export LD_LIBRARY_PATH**
        - (Solaris only)**LD_LIBRARY_PATH_64="<path to lib folder in driver manager>":"<path to IQ Client install location>/IQ15_4/lib64": $LD_LIBRARY_PATH_64**, where **<path to lib folder in driver manager>** is the absolute path to the directory which contains the driver manager's libraries.
        - (Solaris only)**export LD_LIBRARY_PATH_64**
      - Add these lines to the bottom of the `SYBASE.csh` script:
        - (Linux and Solaris) **setenv LD_LIBRARY_PATH <path to lib folder in driver manager>:<path to IQ Client install location>/IQ15_4/ lib64:"$LD_LIBRARY_PATH"** where **<path to lib folder in driver manager>** is the absolute path to the directory which contains the driver manager's libraries.

- (Solaris only) **setenv LD_LIBRARY_PATH_64 <path to lib folder in driver manager>:<path to IQ Client install location>/IQ15_4/lib64:"$LD_LIBRARY_PATH_64"**, where **<path to lib folder in driver manager>** is the absolute path to the directory which contains the driver manager's libraries.

2. If you selected native OS authentication for your cluster, see the *Event Stream Processor Administrators Guide* and follow the instructions in the *Configuring a Pluggable Authentication Module (PAM) for UNIX* topic to enable this feature on your operating system.
3. Enable persistence on your cluster to ensure that projects added to the cluster will stay added between sessions. Edit the file `$ESP_HOME/cluster/nodes/<node name>/<node name>.xml` where <node name> is the name of the node you chose in the Event Stream Processor installer. Change the tag **<Persistence enabled="false">** to **<Persistence enabled=" true">**.
4. If you plan to use the SAP HANA sample project with Event Stream Processor:
   a) Install the SAP HANA client on the machine containing Event Stream Processor.
   b) Edit the `SYBASE.sh` and the `SYBASE.csh` file in the Event Stream Processor installation folder, and add the SAP HANA driver library directory to the LD_LIBRARY_PATH.
      - For `Sybase.sh`, add the lines to the bottom of the script:

        (Linux and Solaris)

        ```
        LD_LIBRARY_PATH=$LD_LIBRARY_PATH:"<path to SAP HANA ODBC Driver Library>"
        export LD_LIBRARY_PATH
        ```

        (Solaris only)

        ```
        LD_LIBRARY_PATH_64=$LD_LIBRARY_PATH_64:"<path to SAP HANA ODBC Driver Library>"
        export LD_LIBRARY_PATH_64
        ```

      - For `Sybase.csh`, add the lines to the bottom of the script:

        (Linux and Solaris)

        ```
        setenv LD_LIBRARY_PATH $LD_LIBRARY_PATH:"<path to SAP HANA ODBC Driver Library>"
        ```

        (Solaris only)

        ```
        setenv LD_LIBRARY_PATH_64 $LD_LIBRARY_PATH_64:"<path to SAP HANA ODBC Driver Library>"
        ```

Configuring Your Service.xml File

Configure your `service.xml` file to connect to the SAP HANA (or ASE) and Sybase IQ databases from Event Stream Processor.

1. Run your installed Event Stream Processor cluster to encrypt the passwords for the database to which you wish to connect using Event Stream Processor. Encrypt any Sybase IQ, SAP HANA, or Adaptive Server Enterprise passwords that you intend to use:

a) Source the `SYBASE.sh` or `SYBASE.csh` script in the Event Stream Processor installation directory.
b) Run the cluster you created during the Event Stream Processor install to use its encryption facilities. Replace **<node name>** with the node name specified in the installer and run this command as a background process: **$ESP_HOME/bin/esp_server --cluster-node $ESP_HOME/cluster/nodes/<node name>/<node name>.xml**

   If your are running a certain version of Solaris, an error may display stating that the libraries for the executable cannot be located. To correct this issue, see the instructions in the *Known Issues for Event Stream Processor* topic in the *Sybase RAP Release Bulletin*.
c) Connect to your cluster node and encrypt all of your database passwords for use in the Event Stream Processor configuration files:
   - Connect to the cluster using the cluster admin tool. Replace **<RPC port>** with the RPC port for your ESP cluster, and **<user name>** with your cluster user name (with native OS authentication this is your OS login name). Run this command : **$ESP_HOME/bin/esp_cluster_admin --uri esp://localhost:<RPC port> --username <user name>**
   - The tool prompts you for your password and then for commands. Encrypt your password using this command: **encrypt <password>**. Replace **<password>** with the password for RAP_USER in the corresponding database. Repeat this for each database password you need to encrypt, and copy the encrypted results as you will need this information later.
   - After encrypting all database passwords, stop the Event Stream Processor node using the command: **stop node <node name>**, where **<node name>** is the name of the ESP node you created.
   - Type **quit** to exit the cluster admin tool.

2. Edit the Event Stream Processor `service.xml` file in `$ESP_HOME/bin/service.xml` to add information about the Sybase IQ database and make it available for connections within Event Stream Processor.

   Replace **<encrypted password>** with the encrypted password for your Sybase IQ database, and add the following ODBC entry for Sybase IQ:

```
<Service Name="SybaseIQService" Type="DB">
 <Parameter Name="DriverLibrary">esp_db_odbc64_lib</Parameter>
 <Parameter Name="DSN">RAPIQ</Parameter>
 <Parameter Name="User">RAP_USER</Parameter>
 <Parameter Name="Password" encrypted="true"><encrypted password>
 </Parameter>
</Service>
```

   If you are using a 32-bit driver manager, change the **DriverLibrary** value to **esp_db_odbc_lib**. Also, ensure that no other entries in the file use **esp_db_odbc64_lib**.
3. If you plan to use the SAP HANA sample project with Event Stream Processor, edit the Event Stream Processor `service.xml file` in `$ESP_HOME/bin/service.xml`

to add information about the SAP HANA database and make it available for connections within Event Stream Processor.

Replace **<encrypted password>** with the encrypted password for your SAP HANA database:

```
<Service Name="HANAService" Type="DB">
 <Parameter Name="DriverLibrary">esp_db_odbc64_lib</Parameter>
 <Parameter Name="DSN">RAPHANA</Parameter>
 <Parameter Name="User">RAP_USER</Parameter>
 <Parameter Name="Password" encrypted="true"><encrypted password></Parameter>
</Service>
```

If you are using a 32-bit driver manager, change the **DriverLibrary** value to **esp_db_odbc_lib**. Also, ensure that no other entries in the file use **esp_db_odbc64_lib**.

4. If you plan to use the Adaptive Server Enterprise sample project with this server, edit the Event Stream Processor `service.xml file` in `$ESP_HOME/bin/service.xml` to add information about the Adaptive Server database and make it available for connections within Event Stream Processor.

   Replace:

   - **<ASE Host Name>** with the host name for your ASE server
   - **<ASE Port>** with the port of your ASE server
   - **<encrypted password>** with the encrypted password for your ASE database

```
<Service Name="SybaseASEService" Type="DB">
 <Parameter Name="DriverLibrary">esp_db_ocs_lib</Parameter>
 <Parameter Name="Host"><ASE Host Name></Parameter>
 <Parameter Name="Port"><ASE Port></Parameter>
 <Parameter Name="User">RAP_USER</Parameter>
 <Parameter Name="Password" Encrypted="true"><encrypted password></Parameter>
 <Parameter Name="AppName">ASEOutputAdapter</Parameter>
</Service>
```

Deploy the RAP Sample Project to Event Stream Processor

Deploy the RAP sample projects to Event Stream Processor before you can use them.

If you are using SAP HANA as your in-memory database, deploy the RAP sample project for SAP HANA. Otherwise, if you are using Adaptive Server Enterprise for your in-memory database, deploy the RAP sample project for Adaptive Server Enterprise.

Do not use spaces in the installation directories for any of the RAP components.

Deploying the RAP Sample Project for SAP HANA

Configure and deploy the RAP sample project using SAP HANA as the in-memory database.

1. Open a new command prompt. Do not reuse a command prompt from the installation of another RAP component.

2. Copy the RAP enablement package to the Event Stream Processor machine, and extract it into a directory of your choice. If any of the RAP components from this current RAP version are installed on this machine, you can use the existing RAP enablement installation. This is referred to as the RAP enablement installation directory throughout the rest of this guide.
3. If you did not run it previously, run the `configure.sh` shell script from the RAP enablement installation directory.

   This file prepares the ESP project files for use.
4. Source the `SYBASE.sh` or `SYBASE.csh` script file located in the RAP enablement installation directory.
5. Source the `SYBASE.sh` or `SYBASE.csh` script in the Event Stream Processor installation directory.
6. If you installed the Japanese supporting data in your historical and in-memory databases, use Japanese sample data with this project.
   a) Change the directory to `$RAP_HOME/sample/inputdata`.
   b) To maintain a copy of the English sample data, rename `STOCK_QUOTE.csv` to `STOCK_QUOTE.csv.en`, and `STOCK_TADE.csv` to `STOCK_TRADE.csv.en`.
   c) Copy `STOCK_QUOTE.csv.jp` to `STOCK_QUOTE.csv`, and `STOCK_TRADE.csv.jp` to `STOCK_TRADE.csv`.
7. Compile the CCL file for the SAP HANA sample project by running: **$ESP_HOME/bin/esp_compiler -i $RAP_HOME/sample/projects/hana/rap.ccl -o $RAP_HOME/sample/projects/hana/rap.ccx**
8. Create primary and overflow directories for the Sybase IQ database adapter to use when loading data into your historical database.

   These directories are initially empty and will be used to temporarily store data during adapter operation. Place the primary and overflow directories on separate file systems and secure them as they may contain sensitive data during the loading process.
9. Edit the project configuration file (`$RAP_HOME/sample/projects/hana/rap.ccr`) to fill in run time values for the project. Replace <ADD PATH>:
   a) Set the **IQPrimaryFileLocation** parameter to the absolute path of the Sybase IQ Output adapter primary directory, which you created during the previous step.
   b) Set the **IQoverflowFilelocation** parameter to the absolute path of the Sybase IQ Output adapter overflow directory, which you created during the previous step.
10. Deploy the SAP HANA sample project to your cluster:
    a) Run the ESP cluster you created during the Event Stream Processor installation. Replace **<node name>** with the node name you specified in the installer and run this command: **$ESP_HOME/bin/esp_server --cluster-node $ESP_HOME/cluster/nodes/<node name>/<node name>.xml**
    b) In another terminal where you have also sourced the `SYBASE.sh` or `SYBASE.csh` script which is installed with Event Stream Processor, start the cluster administration tool. Replace **<RPC port>** with the RPC port for your ESP cluster, **<user name>** with

your cluster user name (with native OS authentication this is your OS login name), and run: **$ESP_HOME/bin/esp_cluster_admin --uri esp://localhost:<RPC port> --username <user name>**

c) Replace **<RAP Enablement installation directory>** with the installation directory for the RAP enablement package and run this command: **add project default/RAP <RAP Enablement installation directory>/sample/projects/hana/rap.ccx <RAP Enablement installation directory>/sample/projects/hana/rap.ccr**

d) Run the project by running this command: **start project default/RAP**

e) Type **quit** to exit the cluster administration tool.

11. To work with the SAP HANA sample project in the ESP Studio:

a) Source the `SYBASE.sh` or `SYBASE.csh` file in the Event Stream Processor installation directory.

b) Start the ESP Studio using this command: **$ESP_HOME/studio/esp_studio**

c) From the File menu, select **Open** and then **Project**.

d) Select **Browse** and select the directory `<RAP Enablement installation directory>/sample/projects/hana/` from the file dialog.

After import, any changes saved to the project are saved to a CCL file in the RAP enablement installation directory. If you make changes to your project, recompile it and redeploy.

Deploying the RAP Sample Project for Adaptive Server Enterprise

Configure and deploy the RAP sample project using Adaptive Server Enterprise as the in-memory database.

1. Open a new command prompt. Do not reuse a command prompt from the installation of another RAP component.
2. Copy the RAP enablement package to the Event Stream Processor machine, and extract it into a directory of your choice. If any of the RAP components from this current RAP version are installed on this machine, you can use the existing RAP enablement installation. This is referred to as the RAP enablement installation directory throughout the rest of this guide.
3. If you did not run it previously, run the `configure.sh` shell script from the RAP enablement installation directory.

   This file prepares the ESP project files for use.
4. Source the `SYBASE.sh` or `SYBASE.csh` script file located in the RAP enablement installation directory.
5. Source the `SYBASE.sh` or `SYBASE.csh` script in the Event Stream Processor installation directory.
6. If you installed the Japanese supporting data in your historical and in-memory databases, use Japanese sample data with this project.

a) Change the directory to `$RAP_HOME/sample/inputdata`.
b) To maintain a copy of the English sample data, rename `STOCK_QUOTE.csv` to `STOCK_QUOTE.csv.en`, and `STOCK_TADE.csv` to `STOCK_TRADE.csv.en`.
c) Copy `STOCK_QUOTE.csv.jp` to `STOCK_QUOTE.csv`, and `STOCK_TRADE.csv.jp` to `STOCK_TRADE.csv`.

7. Compile the CCL file for the Adaptive Server sample project by running: **$ESP_HOME/bin/esp_compiler -i $RAP_HOME/sample/projects/ase/rap.ccl -o $RAP_HOME/sample/projects/ase/rap.ccx**
8. Create primary and overflow directories for the Sybase IQ database adapter to use when loading data into your historical database.

   These directories are initially empty and will be used to temporarily store data during adapter operation. Place the primary and overflow directories on separate file systems and secure them as they may contain sensitive data during the loading process.
9. Edit the project configuration file (`$RAP_HOME/sample/projects/ase/rap.ccr`) to fill in run time values for the project. Replace <ADD PATH>:
   a) Set the **IQPrimaryFileLocation** parameter to the absolute path of the Sybase IQ Output adapter primary directory, which you created during the previous step.
   b) Set the **IQoverflowFileLocation** parameter to the absolute path of the Sybase IQ Output adapter overflow directory, which you created during the previous step.
10. Deploy the Adaptive Server sample project to your cluster:
    a) Run the ESP cluster you created during the Event Stream Processor installation. Replace **<node name>** with the node name you specified in the installer and run this command: **$ESP_HOME/bin/esp_server --cluster-node $ESP_HOME/cluster/nodes/<node name>/<node name>.xml**
    b) In another terminal where you have also sourced the `SYBASE.sh` script which is installed with Event Stream Processor, start the cluster administration tool. Replace **<RPC port>** with the RPC port for your ESP cluster, **<user name>** with your cluster user name (with native OS authentication this is your OS login name), and run: **$ESP_HOME/bin/esp_cluster_admin --uri esp://localhost:<RPC port> --username <user name>**
    c) Replace **<RAP Enablement installation directory>** with the installation directory for the RAP enablement package and run this command: **add project default/RAP <RAP Enablement installation directory>/sample/projects/ase/rap.ccx <RAP Enablement installation directory>/sample/projects/ase/rap.ccr**
    d) Run the project by running this command: **start project default/RAP**
    e) Type **quit** to exit the cluster administration tool.
11. To work with the RAP ASE sample project in the Studio:
    a) Start the ESP Studio using this command: **$ESP_HOME/studio/esp_studio**
    b) From the File menu, select **Open** and then **Project**.
    c) Select **Browse** and select the directory `<RAP Enablement Installation Directory>/sample/projects/ase/` from the file dialog.

After import, any changes saved to the project are saved to a CCL file in the RAP enablement installation directory. If you make changes to your project, recompile it and redeploy.

Setting Up Sybase Control Center for Monitoring and Administering

Monitor and administer the components of a RAP system from a single instance of Sybase Control Center (SCC).

Do not use spaces in the installation directories for any of the RAP components.

1. Open a new command prompt. Do not reuse a command prompt from the installation of another RAP component.
2. Go to *http://www.sybase.com/downloads*.
3. Select **Maintenance Releases and Bug Fixes (EBFs)**.
4. Select the latest version (3.2.7 or later) of Sybase Control Center for your platform.
5. Run the SCC framework installer version 3.2.7 or higher.

   The installer may display additional options other than the options below but the suggestions detailed in this step are important when configuring SCC for RAP. If no suggestion is made for an option, you can accept a default value or see the *Sybase Control Center Installation Guide* for additional information.

   a) Install a licensed copy of SCC.
   b) Select a full installation.
   c) Select a unique port that is not in use when prompted for an RMI port for the SCC agents. If you are installing multiple Sybase products on this machine, the default port of 9999 may already be in use.

      The installer copies files.
   d) Once the installer completes copying files, it prompts you to enable shared disk mode. Choose to disable shared disk mode.
   e) Do not start the Sybase Control Center service through the installer.
6. After the installation completes, install the Event Stream Processor management agent plugin (MAP) into the SCC installation.

   a) Install the ESP MAP into your SCC server by copying the `ESPMAP` directory from the `<ESP Installation Directory>/<SCC-3_2>/plugins` directory to the `plugins` directory inside this SCC installation.

      The `SCC-3_2` portion of this path may be different depending on the version of SCC that you installed.
   b) Remove the `ESPMAP` directory from your Event Stream Processor installation directory.

7. Run Sybase Control Center:
   a) Source the `SYBASE.sh` or `SYBASE.csh` script under the SCC installation directory.
   b) Run **<SCC Install Location>/SCC-3_2/bin/scc.sh**, where **<SCC Install Location>** is the installation location for SCC.

      The `SCC-3_2` portion of this path may be different depending on the version of SCC that you installed.
8. Start the Event Stream Processor, Adaptive Server, and Sybase IQ agents to allow the SCC server to manage them. Start the Adaptive Server agent only if you installed Adaptive Server as your in-memory database.

   Start a new terminal for each product before sourcing. Also note that the SCC directories listed below may differ slightly from your actual directories depending on the version of SCC installed with each product.
   a) Start the Event Stream Processor agent for SCC. From the Event Stream Processor installation folder:
      - Source `SYBASE.sh` or `SYBASE.csh`.
      - Run **SCC-3_2/bin/scc.sh**
   b) (Perform if ASE is your in-memory database) If you have shut down the ASE agent for SCC since your ASE installation, start the agent again. Otherwise, the agent is started because the ASE installer starts this agent for you. From the Adaptive Server installation folder:
      - Source `SYBASE.sh` or `SYBASE.csh` .
      - Run **SCC-3_2/bin/scc.sh**
   c) Start the Sybase IQ agent for SCC. From the Sybase IQ installation folder:
      - Source `IQ.sh` or `IQ.csh`.
      - Run **SCC-3_2/bin/scc.sh**

CHAPTER 3 Performing Postinstallation Tasks

Complete postinstallation tasks after you install the RAP components.

1. *Adjusting the Named rapTAQcache Data Cache Size in Adaptive Server Enterprise*

   Adjust the size of the named rapTAQcache data cache to use as much available memory as possible on the system running Adaptive Server Enterprise.
2. *Configuring Adaptive Server Enterprise to Use an In-Memory Database*

   If you can fit your entire database in memory, you may manually configure Adaptive Server to use an in-memory database (IMDB).
3. *Adjusting the Number of Locks in Adaptive Server Enterprise*

   Use sp_configure to ensure that the number of locks in ASE is greater than the number of locks required by the ASE Output adapter or other applications that may query the data.
4. *Verify that All Servers are Running*

   After you start Sybase IQ and SAP HANA (or Adaptive Server Enterprise), verify that your servers are running and that you can connect to them
5. *Configure a Pluggable Authentication Module (PAM) for UNIX*

   Set up Event Stream Processor to support username and password (Native OS) login using accounts on the UNIX operating system.
6. *Configuring a Scalability Environment*

   To avoid losing information within the writer nodes during synchronization, develop a process to automate the backup and restoration of information in the RAP_WORK_FILE table.
7. *Install Client Tools*

   To connect to RAP databases, you must install the appropriate tools on each client machine.
8. *Displaying a Component Version Number*

   (Optional) Display the version string for a RAP component.
9. *Installation Log Locations*

   See the individual component installation logs for information regarding your installation.

Adjusting the Named rapTAQcache Data Cache Size in Adaptive Server Enterprise

Adjust the size of the named rapTAQcache data cache to use as much available memory as possible on the system running Adaptive Server Enterprise.

Adjusting the named rapTAQcache data cache size increases Adaptive Server performance as an in-memory database. The amount of shared memory available to Adaptive Server determines the values of these configuration variables:

- **max memory**
- **default data cache**

The value for **max memory** should be the maximum amount of shared memory available on your system. The amount of shared memory available is generally controlled by the operating system configuration, but is likely to require modification before you can properly configure a usable Adaptive Server.

The amount of memory to reserve for shared memory is dependent on total memory available on your system. The maximum amount of memory should be allocated for shared memory, but enough memory should be retained for use by the various processes that will be running on the system. For example, if a system is dedicated to the Adaptive Server and ASE Output adapter, and has a total of 32GB of memory (RAM), then you should configure approximately 28-30 GB of operating system shared memory.

1. Use the standard commands and options for your operating system to set the operating system shared memory.
2. Configure Adaptive Server to use the memory.
   a) Configure **max memory** to the total amount of shared memory available on your system
   b) Configure the default data cache.

      The data cache size is dependent on the max memory configuration value. Use as much memory as can be spared, given the system on which Adaptive Server is running. If max memory is 16GB, allocate 15GB for data cache; the rest is to be used for the procedure, Adaptive Server system memory usage and so on.
3. When configuring the named rapTAQcache data cache size, also partition the cache, so that access to the cache from multiple engines does not cause contention.

   The partition count should be roughly equivalent to the number of engines, but you must specify the value in powers of two.

   For example, if you have configured 6 engines, then the cache partition count should be set to 8.

Example of the rapTAQcache Configuration Script

This SQL script configures `max memory` to 16GB, and `default data cache` with 15GB of memory, and 8 cache partitions:

```
sp_configure 'max memory', 8000000
go
sp_cacheconfig 'default data cache', '15G', 'cache_partition=8'
go
```

See also

- *Configuring Adaptive Server Enterprise to Use an In-Memory Database* on page 31

Resolving Adaptive Server Enterprise Tuning Issues for Linux

While tuning the ASE server (by adding named caches or increasing the sizes of existing caches), Linux users may find that the server cannot start.

Even when the memory allocated to the server is large enough to accommodate the additional cache size, and the operating system level parameter SHMMAX (the maximum size in bytes for a shared memory segment) is large enough, the server may not restart.

The problem may be that the operating system level parameter SHMALL, the system-wide maximum of shared memory pages, is not large enough. Increase the value of SHMALL in `/proc/sys/kernel/shmall`.

Modifying `/proc/sys/kernel/shmall` requires root permission.

Configuring Adaptive Server Enterprise to Use an In-Memory Database

If you can fit your entire database in memory, you may manually configure Adaptive Server to use an in-memory database (IMDB).

Switch from the relaxed durability disk-resident database (RDDB) to the IMDB:

1. Use **sp_cacheconfig** to create a named cache of type **inmemory_storage** that is used to host data and log cache devices.

   A single cache must be shared for data and log devices.
2. Use **DISK INIT** to create data and log cache devices on the named cache above.
3. Use **create database** command with **inmemory** keyword and with **DURABILITY = NO_RECOVERY** to create the IMDB for Adaptive Server.

   By default, the IMDB is created based on the model database.

```
-- Create 4 GB cache for in-memory storage
sp_cacheconfig rapTAQcache, "4.5G", inmemory_storage
go
```

```
-- Create 4 GB cache device for rapTAQ data
DISK INIT name = "rapTAQdata",
          physname = "rapTAQcache",
          size = "4G",
          type = "inmemory"
go

-- Create 0.5 GB cache device for rapTAQ log
DISK INIT name = "rapTAQlog",
          physname = "rapTAQcache",
          size = "0.5G",
          type = "inmemory"
go

-- Create rapTAQIMDDB database using existing rapTAQ model database
create inmemory database rapTAQIMDB
    on rapTAQdata = '4g'
    log on rapTAQlog  = '0.5g'
with DURABILITY = NO_RECOVERY
go
```

See also

- *Adjusting the Named rapTAQcache Data Cache Size in Adaptive Server Enterprise* on page 30

Adjusting the Number of Locks in Adaptive Server Enterprise

Use **sp_configure** to ensure that the number of locks in ASE is greater than the number of locks required by the ASE Output adapter or other applications that may query the data.

By default, the Adaptive Server is configured with 100,000 locks. System administrators can use `sp_configure` to change this limit. For example:

```
sp_configure 'number of locks', 200000
```

You may also need to adjust the `sp_configure` parameter, `max memory`, since each lock uses memory.

At a minimum, the number of locks needs to be greater than the number of locks required by the ASE Output adapter. Adaptive Server may also need locks for other applications which may be querying the data. The ASE Output adapter requires locks when inserting data into Adaptive Server. The number of required locks is affected by the bulk batch size parameter in the `rap.ccr` configuration file. This parameter tells the ASE Output adapter how many rows to include in a bulk insert before committing the transaction.

The minimum number of locks for each adapter is:

```
(<number of messages in unpartitioned tables> + <number of partitions
for partitioned message A> + <number of partitions for partitioned
message B> + ... ) * <bulk batch size>
```

For example, the reference application contains 14 message types. Two of these message types (stock quote and stock trade) have tables with 3 data partitions in each. The remaining 12 message types do not have partitioned tables. With a bulk batch size of 10,000, the ASE Output adapter needs a minimum of (12 + 3 + 3) * 10,000 = 180,000 locks. If you have multiple ASE Output adapters running against the same database, use the sum of the locks calculated for each.

Verify that All Servers are Running

After you start Sybase IQ and SAP HANA (or Adaptive Server Enterprise), verify that your servers are running and that you can connect to them

See also

- *Configure a Pluggable Authentication Module (PAM) for UNIX* on page 34

Verifying the Sybase IQ Sever

Use the **stop_iq** utility to verify that the Sybase IQ server is running.

1. Open a command window.
2. Change to `<Sybase IQ Installation Directory>`.
3. Source the environment variables:

   ```
   source IQ.csh
   ```

   or

   ```
   . IQ.sh
   ```
4. Run: `stop_iq`

   The Sybase IQ server displays a message that indicates the current status of your server.
5. Enter N when you are prompted to shut down the server.
6. If the Sybase IQ server is not running, change to `<Sybase IQ Database Location>` and enter:

   ```
   start_iq @RAPSTORE.cfg RAPSTORE.db
   ```

Verifying the SAP HANA Server

Verify that the SAP HANA server is running.

1. Open a command window.
2. Change to `<SAP HANA Installation Directory>`.

3. Source the environment variables:

   ```
   hdbenv.sh
   ```

   or

   ```
   hdbenv.csh
   ```

4. Run:

   ```
   sapcontrol -nr <instance number> -function GetVersionInfo
   ```

5. Check that the resulting message references your SAP HANA installation directory.

Verifying the Adaptive Server Enterprise Server

Use the **showserver** command to verify that the Adaptive Server Enterprise server is running.

1. Open a command window.
2. Change to `<Adaptive Server Enterprise Installation Directory>`.
3. Source the environment variables:

   ```
   source SYBASE.csh
   ```

   or

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

4. Run:

   ```
   <Adaptive Server Enterprise Installation Directory>/ASE-15_0/
   install/showserver
   ```

5. Check that the resulting message references your Adaptive Server Enterprise installation directory.

Configure a Pluggable Authentication Module (PAM) for UNIX

Set up Event Stream Processor to support username and password (Native OS) login using accounts on the UNIX operating system.

See the *Configure a Pluggable Authentication Module (PAM) for UNIX* topic in the *Event Stream Processor Administrators Guide* for detailed instructions.

See also

- *Verify that All Servers are Running* on page 33

Configuring a Scalability Environment

To avoid losing information within the writer nodes during synchronization, develop a process to automate the backup and restoration of information in the RAP_WORK_FILE table. See *Using Sybase IQ 15.4 Multiplex Guide* for information.

Install Client Tools

To connect to RAP databases, you must install the appropriate tools on each client machine.

Sybase IQ client tools are required for a client to connect to the Sybase IQ server. See *Sybase IQ Installation Guide>Client Installations* for your platform.

SAP HANA also requires client tools for a client to connect to the SAP HANA server. See the *SAP HANA Client Installation Guide* for more information. If you are using Adaptive Server Enterprise as your in-memory database, Adaptive Server client tools are required for a client to connect to the in-memory database. See *Adaptive Server Enterprise Installation Guide>PC-Client Installation* for your platform.

See also

- *Displaying a Component Version Number* on page 36

Managing Client Connections

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

If you are connecting to the Sybase IQ and SAP HANA (or Adaptive Server) databases using clients, the connection information is the same as the information that was provided during the installation.

Adaptive Server Enterprise

Clients, such as Open Client™, can access Adaptive Server Enterprise.

Some Adaptive Server client applications, such as PowerDesigner, use ODBC drivers to connect to Adaptive Server. Other third-party applications may also require an ODBC driver.

You can install an ODBC driver separately on client workstations on which you will run third-party or developed products. For more information about the ODBC driver, see the *Adaptive Server Configuration Guide* for your platform and the *ODBC Driver Reference Guide* on the Product Documentation Website.

Sybase IQ and SAP HANA

Sybase IQ and SAP HANA require an ODBC data source on the client computer for each database you want to access using ODBC.

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

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

```
[RAPIQ]
ServerName=RAPIQ
DatabaseName=RAPIQ
CommLinks=tcpip{host=<host name>;port=2638}
Driver=<Sybase IQ ODBC Driver Installation Directory>/IQ-15_4/lib64/
libodbc.so
Charset=utf8
```

Include Charset=utf8 only if you have a multibyte character install.

Here is an example of an entry for SAP HANA:

```
[RAPHANA]
ServerNode=<hana host name>:<hana port>
Driver=<path to SAP hana driver library file>/libodbcHDB.so
CHAR_AS_UTF8=1
```

Include **CHAR_AS_UTF8=1** only if Event Stream Processor passes multibyte characters to the SAP HANA database. If no multi-byte characters need to be processed, you may omit this line.

For information about configuring ODBC, see *Configuring ODBC for Event Stream Processor* on page 19.

Displaying a Component Version Number

(Optional) Display the version string for a RAP component.

Use this procedure to display the version string and copyright for Sybase IQ, SAP HANA (or Adaptive Server Enterprise), Event Stream Processor, and Sybase Control Center.

1. For Sybase IQ, change the directory to `<Sybase IQ Installation Directory>`.
2. Source: `source IQ.csh` or `. IQ.sh`
3. Execute the **start_iq -v** command to check the Sybase IQ version string.
4. Open a new command prompt.

5. For SAP HANA, change the directory to `<SAP HANA Installation Directory>`.
6. Source: `hdbenv.sh` or `hdbenv.csh`
7. Run: **sapcontrol -nr <instance number> -function GetVersionInfo**
8. Open a new command prompt.
9. For Adaptive Server Enterprise, change the directory to `<Adaptive Server Enterprise Installation Directory>`.
10. Source: `source SYBASE.csh` or `SYBASE.sh`.
11. Execute the **dataserver -v** command to check the Adaptive Server Enterprise version string.
12. Open a new command prompt.
13. For Event Stream Processor, change the directory to `<Event Stream Processor Installation Directory>`.
14. Source: `source SYBASE.csh` or `SYBASE.sh`.
15. Execute the **esp_server -v** command to check the Event Stream Processor version string.
16. Open a new command prompt.
17. For Sybase Control Center, change the directory to `<SCC Installation Directory>`.
18. Source: `SYBASE.sh` or `SYBASE.csh`
19. Change directory to `SCC-3_2/bin`.
20. Run: **./scc.sh -v**

See also

- *Install Client Tools* on page 35

Installation Log Locations

See the individual component installation logs for information regarding your installation.

Component	Log File
Sybase IQ	`<Sybase IQ Installation Directory>/log/IQ_Suite.log`
Adaptive Server Enterprise	`<ASE Installation Directory>/log/ASE_Suite.log`

Component	Log File
Event Stream Processor	`<Event Stream Processor Installation Directory>/log/esp_suite.log`
Sybase Control Center	`<SCC Installation Directory>/log/SCCSuite.log`

CHAPTER 4 Migrating from RAP R4.1 to RAP 5.0

There are significant architectural changes in RAP 5.0 that you need to be aware of to successfully migrate to this version.

Key Differences in RAP 5.0

There are key differences between RAP R4.1 and the current version of RAP.

In RAP 4.1, each message was defined in a template. RAP templates contained a message name, type, and a destination table name, as well as column names and their datatypes. The publishers (feed handlers) multicast data in the network while the subscribers got the data off the network and persisted it to the Adaptive Server Enterprise and Sybase IQ databases. Both publishers and subscribers used RAP templates to validate and process data.

Sybase RAP 5.0 is now a project within the Event Stream Processor environment. The publishers are now input adapters, and the subscribers are now the Sybase ASE and Sybase IQ Output adapters. RAP templates are now schemas for streams and windows in Event Stream Processor. Each message type requires a separate stream or window that has an ASE and Sybase IQ Output adapter attached to it.

Event Stream Processor has streams and windows with defined schema, and these streams and windows are attached to input and output adapters. The role of input adapters is to convert data into Event Stream Processor datatypes whereas the role of output adapters is to write out the data into the datatype used by the destination. See the Event Stream Processor documentation for additional information.

Data is brought into Event Stream Processor using a variety of input adapters and is output using various output adapters. RAP ships with a sample project that you can use to accept data from these input adapters (or you can write your own custom input adapter) and feed it into the ASE and Sybase IQ Output adapters. These adapters output data into the RAPCache (Adaptive Server Enterprise) and RAPStore (Sybase IQ) databases, and also enable you to persist the data into the databases. You can monitor and administer RAPCache, RAPStore, and the ASE and Sybase IQ Output adapters using the Sybase Control Center for Event Stream Processor.

Due to these differences, you need to perform the steps below to migrate from RAP 4.1 to 5.0:

1. Create a project in Event Stream Processor. This includes creating streams to represent your RAP messages and attaching the ASE and Sybase IQ Output adapters to these streams.
2. Upgrade the Adaptive Server and Sybase IQ databases and make the necessary schema changes.
3. Convert your existing RAP publishers to use the Event Stream Processor C SDK.

Create Projects in Event Stream Processor

You need to create at least one stream or window for each message definition in the RAP templates.

Make these streams output streams, attach the ASE and Sybase IQ Output adapters to them and have them point to the particular table to which the stream writes data. Specify the table within the table property in both adapters. The adapters write out the data coming into this output stream to the Adaptive Server and Sybase IQ databases.

The data you want to store in the databases can be delivered to Event Stream Processor using internal or external input adapters. Internal adapters run within Event Stream Processor and are managed by it while external adapters use the ESP SDK and are run outside of Event Stream Processor within a separate process. For more information on adapters, see the *Event Stream Processor Adapters Guide*.

Event Stream Processor can perform complex processing of the incoming data before it is stored in the database. For further information on how to preprocess the information before storing it, see the Event Stream Processor documentation. In RAP 4.1, each database subscriber stored database connection information in an xml file (`rapstore.xml` or `rapcache.xml`). In ESP database connectivity information is provided in the `$ESP_HOME/bin/service.xml` file. The RAP installer asks you for connection credential information for the Adaptive Server and Sybase IQ databases, and updates this file. To change it later on, update the SybaseASEService and SybaseIQService entries in the `service.xml` file.

The RAP 5.0 sample project is an example of the conversion of a RAP 4.1 system using the stock quote and stock trade message types and serves as a good starting point for your conversion. The model contains two streams that process incoming data. One stream processes data for stock trades and the other processes stock quotes.

These input streams are attached to the DSV Input adapter that reads sample stock quote and stock trade data from CSV files. The output streams for both trades and quotes are derived from these input streams by looking up a stock symbol using data read from the instrument table in the Sybase IQ database. This is an example of how to achieve the RAP 4.1 lookup functionality using Event Stream Processor's capabilities. The ASE and Sybase IQ Output adapters are attached to both of these output streams. The sample project can be found under

`<install>/RAP-5_0 /RAPCore/workspace/rap/rap.ccl`. Here is a sample from `rap.ccl`:

```
CREATE SCHEMA StockQuoteInSchema (
                INSTRUMENT_ID STRING,
                QUOTE_DATE DATE,
                QUOTE_SEQ_NBR INTEGER,
                TRADING_SYMBOL STRING,
                QUOTE_TIME BIGDATETIME,
                ASK_PRICE MONEY(4),
                ASK_SIZE INTEGER,
                BID_PRICE MONEY(4),
                BID_SIZE INTEGER
                );
CREATE INPUT STREAM StockQuoteInStream SCHEMA StockQuoteInSchema;
```

The ASE and Sybase IQ Output adapters have the same properties that the RAPCache and RAPStore subscribers previously had (`rapcache.xml` and `rapstore.xml`). If you changed those property values from the default values, make this change in the adapters as well.

For detailed information on creating and running a project in Event Stream Processor, see the *Event Stream Processor Getting Started Guide*, *Studio Users Guide*, and the `Programmers Guide`. Generate projects using the ESP Studio.

Datatype Mapping for RAP

RAP datatypes map to Event Stream Processor datatypes.

RAP Datatype	Event Stream Processor Datatype
`uint8`	`integer`
`uint16`	`integer`
`uint32`	`long`
`uint64`	`long`
`sint8`	`integer`
`sint16`	`integer`
`sint32`	`integer`
`sint64`	`long`
`decimal(p, s)`	`float` or `money(s)`
`datetime`	`timestamp`

RAP Datatype	Event Stream Processor Datatype
datetime2	bigdatetime
date	date
time	timestamp
time2	interval
string	string

If a RAP template definition contains a decimal datatype and you are planning on using the Event Stream Processor money datatype, make sure that the Event Stream Processor money(s) datatype has enough precision and scale. All Event Stream Processor money datatypes have a precision of 18 and a scale from 1 to 15.

However, Event Stream Processor does not have datatypes equivalent to certain datatypes in Adaptive Server and Sybase IQ. Here are the suggested mappings for these datatypes in Event Stream Processor:

Adaptive Server Enterprise Datatypes	Sybase IQ Datatypes	Event Stream Processor Datatypes
tinyint	tinyint	integer
smallint	smallint	integer
unsigned smallint		integer
unsigned int	unsigned int	long
smallmoney	smallmoney	money
unsigned bigint	unsigned bigint	long
decimal (p,s)	decimal (p,s)	money
numeric (p,s)	numeric (p,s)	money

During message flow, the adapters verify whether incoming values can fit into the specified database columns without losing precision and reject the rows if they cannot. The mappings that require runtime data checks are identified at adapter startup, and only these mappings are verified for each record flowing through the adapter. Due to the negative performance impact of runtime checks on each Event Stream Processor record, Sybase recommends that you only use these mappings if changes to existing schemas are not an option.

For a complete list of supported Adaptive Server and Sybase IQ datatypes and datatype mappings requiring runtime checks, see the sections for the Sybase ASE Output Adapter and the Sybase IQ Output Adapter in the *Event Stream Processor Adapters Guide*.

Database Upgrades

The second step in the migration process is migrating to the most recent versions of the Adaptive Server Enterprise and Sybase IQ databases.

RAP 4.1 used Adaptive Server Enterprise 15.5 and Sybase IQ 15.2. RAP 5.0 uses Adaptive Server Enterprise 15.7 and Sybase IQ 15.4. See the Adaptive Server Enterprise and Sybase IQ migration guides for detailed upgrade procedures.

The schema for the RAPStore table, RAP_WORK_FILE, has also changed. To update this table, run the following SQL:

```
alter table RAP_USER.RAP_WORK_FILE
DROP TRAN_ID,
DROP OP_DB_IP,
DROP OP_DB_PORT,
DROP OP_DB_USER,
DROP OP_DB_PASS,
ADD ADAPTER_NAME SYB_VERYLARGESTRING NOT NULL;
commit;
```

The Sybase IQ Output adapter also requires permissions to allow for client-side loading of files. The following commands will enable client side loading on the database, and grant permissions to perform client-side loading to the RAP_USER:

```
set option public.allow_read_client_file='on';
grant READCLIENTFILE to RAP_USER;
```

Aside from the change to the RAP_WORK_FILE table and the permissions on the Sybase IQ database, there are no additional changes to the Adaptive Server and Sybase IQ schema.

Converting Publishers to Adapters

The third step in the migration process, after connecting to an Event Stream Processor project and upgrading the databases, is upgrading your RAP publishers so that they can publish data to Event Stream Processor.

RAP publishers are standalone programs that need to be started manually. They use the RAP Publisher API to publish data. The least invasive migration route is to replace the RAP Publisher API calls with the ESP C SDK calls. In doing so, the publisher becomes an external adapter to Event Stream Processor. However, this adapter is not managed by Event Stream Processor, and therefore, cannot be started or stopped by it. If you want to manage your publisher through Event Stream Processor, see the Event Stream Processor Adapters Guide for more information. Note that when creating a managed adapter, the publisher must run on the same machine as Event Stream Processor.

Follow these general steps to move your existing publishers to use the new C SDK. For detailed information on the C SDK, see the Event Stream Processor C SDK documentation.

1. Your publisher needs certain information to connect to Event Stream Processor.

Property	Description
ESPURI	The URI needed to connect to a project. For example, esp://localhost:19011/default/RAP.
Encryption	Indicates whether encryption is enabled. See the Event Stream Processor SDK esp_uri_set_ssl API.
AuthenticationType	Indicates how the client is authenticated. Possible values are USER_PASSWORD, SERVER_RSA, and KERBEROS.
Username	User name for connecting to the ESP cluster. This is needed for authentication.
Password	The user password. Required for authentication.

2. Connect to your project:
   a) Use esp_credentials_create() API to set the authentication type.
   b) If the authentication type is USER_PASSWORD (Native OS), use esp_credentials_set_user() and esp_credentials_set_password() to set the username and password.
   c) If the authentication type is SERVER_RSA, use esp_credentials_set_keyfile to set the key file.
   d) If the authentication type is KERBEROS, use esp_credentials_set_user() to set the username. There are additional environment variables that you need to set to use this authentication type. See the C SDK documentation for more information.
   e) Use esp_uri_create_string() to set the uri provided to connect to the project.
   f) If you provided the encryption property and set it to true, use esp_uri_set_ssl() API.
   g) Connect to the project using the esp_project_get() API and then the esp_project_connect() API.
   h) Create an Event Stream Processor publisher object using esp_project_create_publisher() API. Use these publisher options for optimal performance:
      - Mode: EXPLICIT_BLOCKING
      - Access: DIRECT_ACCESS
      - Async: true
      - Buffer Size: 4096
   i) Connect to the publisher object using the esp_publisher_connect() API.
   j) Get socket MSS (Maximum Segment Size) for the ESP Server using these APIs to determine how much data the Server can accept within a packet. This helps you maximize network packet utilization.

- esp_socket_details_create()
- esp_publisher_get_socket_details()
- esp_socket_details_get_mss()

3. Get access to an output stream and publish data.

a) Get the stream object using the esp_project_get_stream() API.

b) Create a message writer object by passing in the publisher and stream objects as parameters into the esp_publisher_get_writer() API.

c) Open a new envelope using esp_message_writer_start_envelope() API. You can also use transaction by using esp_message_writer_start_transaction() API. Transactions are atomic whereas envelopes are not.

d) Start adding a new row (esp_relative_rowwriter_start_row() API) and set the column values by calling the SDK APIs. This table displays mapping between Event Stream Processor datatypes and the C SDK APIs:

Event Stream Processor Datatype	Event Stream Processor C SDK API
`integer`	esp_relative_rowwriter_set_integer
`long`	esp_relative_rowwriter_set_long
`float`	esp_relative_rowwriter_set_float
`timestamp`	esp_relative_rowwriter_set_timestamp
`bigdatetime`	esp_relative_rowwriter_set_bigdatetime
`date`	esp_relative_rowwriter_set_date
`interval`	esp_relative_rowwriter_set_interval
`string`	esp_relative_rowwriter_set_string

e) Keep track of the amount of data written. If you reach the MSS size, close the envelope using the esp_message_writer_end_block() API and publish this envelope using the esp_publisher_publish() API. At this point, start a new envelope. If you do not want to use MSS, you can always send a set number of rows in an envelope. This depends on the amount of data in a row. In this case, find the optimal block size by trying out different values of block size. For the RAP Sample Project, the block size of 4096 gave optimal performance.

CHAPTER 5 Uninstalling RAP Components

Uninstall the product.

Prerequisites

Shut down all Sybase IQ and SAP HANA (or ASE) servers, as well as any other running components.

Task

To uninstall product components, use the **rm -rf** command to remove the specific folders.

CHAPTER 6 RAP in a High Availability Environment

RAP ensures continuous availability of financial market data by providing high availability and disaster recovery through a combination of server configurations and procedures.

If you are using SAP HANA as your in-memory databases, see the SAP HANA documentation for information on how to ensure disaster recoverability and high availability of the SAP HANA server.

If you are using Adaptive Server as your in-memory database, if the in-memory cache becomes unusable, Adaptive Server Enterprise is synchronized with a secondary copy of the Adaptive Server (RAPCache) to ensure automatic, timely recovery and failover of the primary cache.

For disaster recovery of Sybase IQ, the Sybase IQ multiplex is used for automatic failover of the writer node from a designated reader node. For more information, see the Adaptive Server Enterprise 15.7 documentation, and the Sybase IQ 15.4 documentation.

Ensuring Market Data Consistency

You can configure RAP to enhance the availability of market data to trading applications.

By leveraging feed capabilities, and those provided by input adapters, RAP handles market data messages at a very high rate with a high probability of delivery. To achieve this, run two instances of Event Stream Processor on separate machines. Run the RAP ASE project on both instances. The primary instance feeds both RAPCache (Adaptive Server) and RAPStore (Sybase IQ), while the secondary instance feeds only RAPCache. The two RAPCache data stores should contain the same data during normal operations.

Generally, you do not need to configure a secondary RAPStore for the primary RAPStore (an instance of an IQ multiplex server). The RAPStore potentially contains several years' worth of market data, making a redundant configuration impractical. If the data in the RAPStore becomes stale due to loss of the data stream or subscriber, the RAPStore must be loaded from data in the RAPCache. RAP has no automatic mechanism to load the RAPStore from the RAPCache, but you can load it manually.

This figure illustrates the RAP HA configuration, which can be used to substantially enhance availability of market data to trading applications. You can establish similar configurations using third-party software.

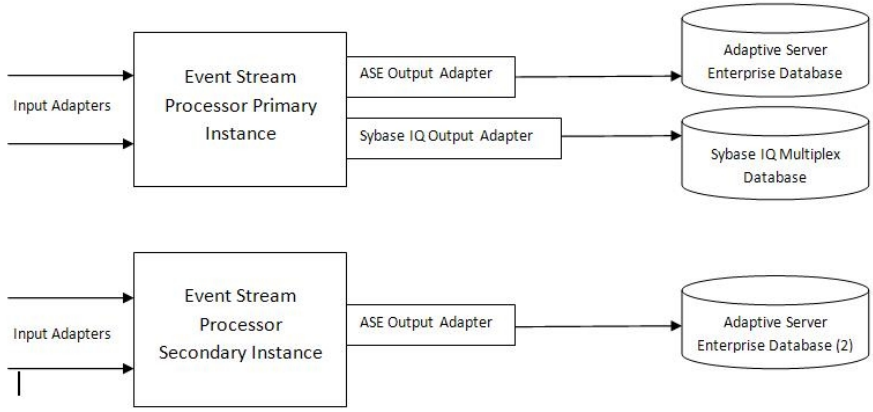

Roles of RAP Components

All high availability requirements are met through the configuration of installed RAP components.

Install and run all primary software components (Event Stream Processor, RAPCache, and RAPStore) on their own hardware and network. Install and run all secondary software components (Event Stream Processor and RAPCache) on their own hardware and network. Primary and secondary software components must run on separate hardware and networks.

RAP components that have a role in capturing and delivering market data are:

- RAPCache (Adaptive Server Enterprise) database
- RAPStore (Sybase IQ) database
- Event Stream Processor

RAPCache Database

If high data availability is required for the RAPCache (Adaptive Server Enterprise), install and configure a secondary RAPCache.

The RAPCache failover strategy requires a secondary installation of Adaptive Server Enterprise to complement the secondary configuration of Event Stream Processor.

If high availability is required for the RAPCache, a secondary RAPCache must be available. During normal operations, both RAPCaches contain the same data.

The primary instance of the RAPCache is installed and runs on a single host, while the secondary RAPCache is installed and runs on a separate host, that ideally uses a separate network.

RAPStore Database

Since the RAPStore (historical database) may save and maintain several years worth of market data, a redundant configuration is not practical.

If the data in the RAPStore becomes stale due to loss of the data stream or subscriber, you must manually load the RAPStore data from the RAPCache.

You may choose to install a complete secondary system for the RAPStore, as in the case of the RAPCache (the secondary instance of Event Stream Processor does not preclude this configuration), but the expense of a secondary RAPStore could be prohibitive, as there can be many terabytes of data.

If you configure a secondary system for the RAPStore, the RAPStore is loaded only by the primary instance of Event Stream Processor.

If there is a failure anywhere in the primary Event Stream Processor instance, data in the RAPStore becomes stale until manual intervention occurs. In this case, the RAPStore data is loaded from either the primary RAPCache or the secondary RAPCache, whichever is current.

Install a redundant RAPStore (an instance of a Sybase IQ multiplex server) on a separate machine. The redundant RAPStore does not run unless there is a hardware failure on the machine running the RAPStore. In the event of hardware failure of the RAPStore machine, you must, at the end of the trading day, launch the redundant RAPStore, load data from the current RAPCache into the RAPStore, and use the redundant RAPStore for the operations of the following day until the original hardware is repaired.

Loading the RAPStore from the RAPCache

If the data in the RAPStore becomes stale due to the loss of the primary Event Stream Processor instance, use a manual procedure to load the RAPStore from data in the RAPCache.

1. Determine the maximum key value currently in the RAPStore.

   The maximum key value is determined by date and time. Each market data message is stored in its corresponding database table with a date and timestamp value. Corresponding values for these key fields are the current date and maximum timestamp value less 10 minutes. Since the messages may arrive out of order, Sybase recommends that you allow a grace period of at least 10 minutes.
2. Create a view in the current RAPCache to select from the corresponding Adaptive Server Enterprise table where key values are greater than the key values from RAPStore. For example, if you are transferring the *STOCK_QUOTE* table, you can create a view called *STOCK_QUOTE_XFER* using:

   ```
   create view RAP_USER.STOCK_QUOTE_XFER ( INSTRUMENT_ID,
   QUOTE_DATE, QUOTE_SEQ_NBR, TRADING_SYMBOL, QUOTE_TIME, ASK_PRICE,
   ASK_SIZE, BID_PRICE, BID_SIZE ) as select from STOCK_QUOTE where
   QUOTE_TIME > 'datetime key value from previous step'
   ```
3. Bulk-copy the table out to a file from the view in current RAPCache using the Adaptive Server **bcp** utility.

Use these flags: -c -t , -r \\n:

- -c flag - tells **bcp** to output character data instead of binary data.
- -t , flag - tells **bcp** to use the comma as a field terminator.
- -r \\n flag - tells **bcp** to use a newline as a row terminator.

Here is the **bcp** command, using the example view *STOCK_QUOTE_XFER*:

```
bcp RAP_USER.STOCK_QUOTE_XFER out
stock_quote_xfer.csv -c -t , -r \\n
-URAP_USER -Prap_user
-S<RAPCache Database Server Name>
```

4. Transfer the bulk copy files from the Adaptive Server Enterprise server to the Sybase IQ server that is hosting the RAPStore database.
5. Load the data into the RAPStore on the Sybase IQ server using **dbisql** to submit a **LOAD TABLE** command. The **IGNORE CONSTRAINT UNIQUE 0** option allows IQ to throw away any duplicate key values without terminating the load operation.

   Continuing the example from above, the **LOAD TABLE** command looks similar to:

```
load table RAP_USER.STOCK_QUOTE
(
INSTRUMENT_ID    null (blanks,'NULL') ,
QUOTE_DATE    null (blanks,'NULL') ,
QUOTE_SEQ_NBR    null (blanks,'NULL') ,
TRADING_SYMBOL null (blanks,'NULL') ,
QUOTE_TIME    null (blanks,'NULL') ,
ASK_PRICE    null (blanks,'NULL') ,
ASK_SIZE    null (blanks,'NULL') ,
BID_PRICE    null (blanks,'NULL') ,
BID_SIZE  '\x0a'
)
from '<path>/stock_quote_xfer.csv'
quotes off
escapes off
preview on
ignore constraint unique 0
;
commit
;
```

Event Stream Processor

To have RAP running in a high availability environment, install and configure a primary and secondary instance of Event Stream Processor.

Both instances of Event Stream Processor remain running at all times. These instances are installed on two different machines and, ideally, use different networks. Both instances run a RAP project, but only the primary instance feeds the RAPStore. Each instance feeds a separate RAPCache.

Application Connection Failover

If the primary RAPCache (Adaptive Server Enterprise) fails, the application using it may not automatically failover to the secondary RAPCache. Other, manual means may be needed to redirect the failover.

The failure may not be with Adaptive Server itself; upstream components may have failed.

For this reason, do not rely on the normal Open Client (and ODBC or JDBC™) connection failover mechanisms that are available to Sybase client applications, unless the primary Adaptive Server shuts down. When a failure in the primary Adaptive Server is detected, redirect software that uses the primary Adaptive Server to fail over by manual means.

Minimizing Complications of Connection Failover

If connection failover occurs, choose the strategy most appropriate for your environment to minimize complications.
Choose from several possible response scenarios for connection failover:

- If the primary RAPCache is running, shut down the RAPCache server to force automatic client-side connection failover, if failover is configured for client applications.
- Change applications to be aware of both the primary and secondary Adaptive Server Enterprise installations. For existing application software, this may be impractical for the users.
- Change the `interfaces` or `sql.ini` file of the application. This may not be a practical solution if there are hundreds of applications running on the PCs or workstations.

Cleaning Up the RAPStore

Once the **LOAD TABLE** command completes for every table that must be reloaded, clean up the RAPStore before restarting the Sybase IQ Output adapter. The RAPStore loads data using files written by the Sybase IQ Output adapter.

1. Go to the directory where the files are located, which is specified by the primaryFileLocation or overflowFileLocation property in the `rap.ccr` configuration file.
2. In the *RAP_WORK_FILE* table, remove each file in the directory:
   ```
   delete from RAP_WORK_FILE where FILE_NAME='filename'
   ```
3. Delete all of the files in the directory.

CHAPTER 7 Samples Installation Directory

Sample SQL query scripts and data are installed by default into the `<RAP Enablement Installation Directory>/sample` directory.

Each database (Sybase IQ, SAP HANA, Adaptive Server) has `data` and `queries` subdirectories within their main installation directory. These directories contain data used to populate the tables for running the sample queries and the actual queries.

Table 3. /sample Subdirectory Names and Files

Directory	Description
`/projects/hana`	Sample RAP project for SAP HANA.
`/projects/ase`	Sample RAP project for Adaptive Server Enterprise.
`/inputdata`	Sample input data to load into Event Stream Processor for running the sample projects.

CHAPTER 8 Troubleshooting the Installation

If you did not install successfully, refer to the individual component installation logs to troubleshoot errors in your installation.

Installation Log Files

Refer to the individual component installation log files to troubleshoot your installation.

Component	Log File
Sybase IQ	`<Sybase IQ Installation Directory>/log/IQ_Suite.log`
Adaptive Server Enterprise	`<ASE Installation Directory>/log/ASE_Suite.log`
Event Stream Processor	`<Event Stream Processor Installation Directory>/log/esp_suite.log`
Sybase Control Center	`<SCC Installation Directory>/log/SCCSuite.log`

CHAPTER 9 Appendix

Convert your existing Adaptive Server or Sybase IQ databases to use RAP licenses.

Activating Your SySAM Adaptive Server Enterprise License

To convert your existing Adaptive Server database to use a RAP license, activate your SySAM Adaptive Server Enterprise license after you obtain the license file and install Adaptive Server on your host machine.

Prerequisites

Obtain a RAP license from SPDC or SMP.

Task

1. Change the directory to `<Adaptive Server Enterprise Installation Directory>`.
2. Source the environment variables: `source SYBASE.csh` or `. SYBASE.sh`
3. Log in to the Adaptive Server Enterprise server.
   - To check the license type, enter:

     ```
     sp_lmconfig
     go
     ```

     The valid license types for RAP are:
     - CP (CPU license for Production)
     - DT (Development and Test, not CPU dependent)
     - SF (CPU license for Stand By Production)
   - To change the license type to match your license file, if it is different than the license type you selected when installing ASE, run the appropriate command:

     ```
     sp_lmconfig ‘license type’, 'LT'
     go
     ```

     where *LT* is a valid license type.
4. Enable your Adaptive Sever Enterprise installation to use your RAP license.

   ```
   sp_lmconfig 'application type','RAP'
   go
   ```
5. Shut down the Adaptive Server Enterprise server:

   ```
   shutdown with nowait
   ```

```
go
```

6. If you are running in served mode, copy the license file to `<license_server_install_dir>/SYSAM-2_0/licenses` on the license server.

   Sybase recommends the use of the served license model with RAP.

7. If you are running in unserved mode, copy the license file to `<Adaptive Server Enterprise Installation Directory>/SYSAM-2_0/licenses` on the host machine.

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

8. Change to `<license_server_install_dir>/SYSAM-2_0/bin`. Start the license server, if neccessary:

   ```
   sysam start
   ```

   If the license server was started before you copied the license information to `<license_server_install_dir>/SYSAM-2_0/licenses` run:

   ```
   sysam reread
   ```

   SySAM reloads all licenses located in `<license_server_install_dir>/SYSAM-2_0/licenses`.

   You must have at least one served license in `<license_server_install_dir>/SYSAM-2_0/licenses`, or you see `"License Manager: Can't initialize..."` when you attempt to start the license server.

9. Update the license server options file:
   a) Change to `<license_server_install_dir>/SYSAM-2_0/licenses` on the license server.
   b) Use a text editor to add `INCLUDE SY_RAP PROJECT RAP` to the file `SYBASE.opt`.

10. Create a `sybase.lic` file in `<Adaptive Server Installation Directory>/SYSAM-2_0/licenses`.

    The `sybase.lic` file specifies the port number of the license server, so SySAM can retrieve the appropriate license. The contents of the `sybase.lic` file is:

    ```
    SERVER <license_server_host_name> ANY 27000
    USE_SERVER
    ```

    where 27000 is the default port. You can use a different port number.

11. Restart the Adaptive Server Enterprise server.

Note: To verify your license, either run the **sp_lmconfig** stored procedure in Interactive SQL (**isql**), or check the Adaptive Server error log.

In the Adaptive Server error log, a line beginning with `"Checked out license..."` indicates that the license configuration is successful. If you see a `"Sysam: FLEXnet`

`Licensing error:"` check with your Sybase Technical Support or customer service representative to resolve the issue.

See also

- *Obtaining Licenses* on page 4

Restarting the Adaptive Server Enterprise Server

After you activate your SySAM Adaptive Server Enterprise license, restart the Adaptive Server Enterprise server.

1. Open a command window.
2. Change the directory to `<Adaptive Server Enterprise Installation Directory>`.
3. Source the environment variables: `source SYBASE.csh` or `. ./SYBASE.sh`.
4. Change the directory to `<Adaptive Server Enterprise Installation Directory>/ASE-15_0/install`.
5. Start the Adaptive Server Enterprise server: **./startserver -f RUN_<server_name> &**

   where *<server_name>* is the name of the server. The ampersand (&) starts a background process that continues to run after you exit the terminal window.

Activating Your SySAM Sybase IQ License

To convert your existing Sybase IQ database to use a RAP license, activate your SySAM Sybase IQ license after you obtain the license file and install Sybase IQ on your host machine.

Prerequisites

Obtain a RAP license from SPDC or SMP.

Task

1. Change the directory to `<Sybase IQ Installation Directory>/`.
2. Source the environment variables: `source IQ.csh` or `. IQ.sh`.
3. Log in to the Sybase IQ server.
   - To check the license type, enter:

     ```
     sp_iqlmconfig
     ```

     The valid license types for RAP are:
     - CP (CPU license for Production)
     - DT (Development and Test, not CPU dependent)

    - SF (CPU license for Stand By Production)
  - To change the license type, run the appropriate command:

    ```
    sp_iqlmconfig 'Edition', 'EE'
    ```

    where EE is the Enterprise Edition, and:

    ```
    sp_iqlmconfig 'license type', 'LT'
    ```

    where *LT* is a valid license type.
4. To shut down the Sybase IQ server, run the **stop_iq** script.
5. If you are running in served mode, copy the license file to `<license_server_install_dir>/SYSAM-2_0/licenses` on the license server.

   Sybase recommends the use of the served license model with RAP.
6. If you are running in unserved mode, copy the license file saved from SPDC or SMP to `<Sybase IQ Installation Directory>/SYSAM-2_0/licenses` on the host machine.

   If you use the **ftp** utility to transfer the license file, be sure to specify ASCII mode.
7. Change to `<license_server_install_dir>/SYSAM-2_0/bin`. Start the license server, if neccessary:

   ```
   sysam start
   ```

   If the license server was started before you copied the license information to `<license_server_install_dir>/SYSAM-2_0/licenses` run:

   ```
   sysam reread
   ```

   SySAM reloads all licenses located in `<license_server_install_dir>/SYSAM-2_0/licenses`.

   You must have at least one served license in `<license_server_install_dir>/SYSAM-2_0/licenses`, or you see "`License Manager: Can't initialize...`" when you attempt to start the license server.
8. Create a `sybase.lic` file in `<Sybase IQ Installation Directory>/SYSAM-2_0/licenses`.

   The `sybase.lic` file specifies the port number of the license server, so SySAM can retrieve the appropriate license. The contents of the `sybase.lic` file is:

   ```
   SERVER <license_server_host_name> ANY 27000
   USE_SERVER
   ```

   where 27000 is the default port. You can use a different port number.
9. Restart the Sybase IQ server.

Note: To verify your license, you can run the **sp_lmconfig** stored procedure in Interactive SQL (**isql**) or check the Sybase IQ error log.

In the Sybase IQ error log, a line beginning with "`Checked out license...`" indicates that the license configuration is successful. If you see a "`Sysam: FLEXnet Licensing error:`" message, check with your Sybase technical support or customer service representative to resolve the issue.

See also

- *Obtaining Licenses* on page 4

Restarting the Sybase IQ Server

After you activate your SySAM Sybase IQ license, restart the Sybase IQ server.

1. Open a command window.
2. Change the directory to `<Sybase IQ Installation Directory>`.
3. Source the environment variables: `source IQ.csh` or `. IQ.sh`.
4. Shut down the Sybase IQ server, if it is running: `stop_iq`
5. Change the directory to `<Sybase IQ Database Location>`.
6. Start the Sybase IQ server: `start_iq @RAPSTORE.cfg RAPSTORE.db`

Index

L

M

O

P

R

S

T

U

V