SYBASE®

An **SAP®** Company

Sybase Control Center for Sybase IQ

3.2.3

Contents

About Sybase Control Center for Sybase IQ

Sybase® Control Center for Sybase IQ is a Web-based tool for managing and monitoring Sybase IQ single-node and multiplex servers. The two main features are administration and monitoring.

Sybase Control Center monitoring supports Sybase IQ version 15.1 ESD #2.1 or later. Sybase Control Center administration supports Sybase IQ version 15.3 or later.

The Sybase Control Center architecture allows a small number of Sybase Control Center servers to monitor all Sybase IQ servers in an enterprise using the Sybase Control Center agent. The Sybase Control Center agent is installed with each Sybase IQ server.

Sybase Control Center for Sybase IQ provides availability monitoring, historical monitoring, and real-time monitoring in a scalable Web application that is integrated with management modules for other Sybase products. It offers shared, consolidated management of heterogenous resources from any location, real-time notification of availability and performance, and intelligent tools for spotting performance and usage trends, all via a thin-client, rich Internet application delivered through your Web browser.

New Features in Sybase Control Center for Sybase IQ

Brief descriptions of new and enhanced features with links to complete information.

Table 1. New and enhanced features: Sybase Control Center for Sybase IQ

Feature	Topics
Functions and procedures – create and manage functions and procedures written in a variety of SQL dialects and external environments.	*Manage Functions* on page 294 *Manage Procedures* on page 301
Java external environments – install, update, modify, or delete Java classes and JAR files.	*Manage External Environments* on page 308
SQL query execution – run ad hoc queries or stored procedures.	*Executing a SQL Query* on page 323
Text indexes and text configuration objects – create and manage text indexes for full-text searches and custom text configuration objects for your text indexes.	*Manage Text Indexes* on page 314 *Manage Text Configuration Objects* on page 320
Data collections – five new data collections allow you to improve performance and better manage alerts by avoiding the collection of duplicate statistics.	*Sybase IQ Data Collections* on page 130

Feature	Topics
Shared-disk mode – run multiple Sybase Control Center instances from a single installation on a shared disk.	*Deploying an Instance from a Shared Disk Installation* on page 64 *Shared Disk Mode* on page 66 *Instances* on page 180
Run collections without saving data – configure data collection jobs so they update monitoring views but do not add data to the repository.	*Setting Up Statistics Collection* on page 128
E-mail configuration for alerts – specify a domain name or change the sender name for e-mail alert no-tifications.	*Configuring the E-mail Server* on page 104
Resizing wizards – make windows larger or smaller by dragging the edges or corners.	(No new topic)
Authority management – grant database authorities to a user or group and revoke database authorities from a user or group.	*Manage Authorities* on page 286
Database management – create databases, view and modify database options, and view properties of a database.	*Manage a Database* on page 229
Dbspace and DB file management – add, modify, view properties, or delete a dbspace. Add, modify, view properties, or delete a DB file.	*Manage a Dbspace* on page 235 *Manage DB Files* on page 242
Logical server management – group a subset of phys-ical hardware resources together as a logical entity that appears as a single multiplex server.	*Logical Servers* on page 267
Login policy management – manage the rules for user login.	*Manage Login Policies* on page 288
Multiplex management – change a multiplex server configuration, manage the secondary and failover nodes, and configure logical servers.	*Manage a Multiplex Server* on page 258
Server management – create, start, stop, view prop-erties, change server configuration, and generate ad-ministration scripts.	*Manage a Server* on page 218
Shared temporary dbspace management – manage shared temporary stores to communicate and store temporary on-disk structures for distributed query processing.	*Viewing Multiplex Dbspace Statistics* on page 257 *Adding a DB File* on page 243

Feature	Topics
User and group management – add, change, and delete users, or groups containing users or other groups, as a prerequisite to managing user authorities.	*Manage Users and Groups* on page 276
Administration Console – manage existing resources and create new ones. Column-based filtering lets you display only the objects you are interested in.	Various topics, including: *Browsing and Managing Resources* on page 161 *Common Display Options* on page 7
Testing scripts – test the execution of alert-triggered scripts to make sure they work as expected.	*Testing an Alert-triggered Script* on page 170 *Creating an Alert* on page 146
Automatic logout – a Sybase Control Center administrator can configure the logout timer to end users' login sessions after a specified period of idleness.	*Configuring the Automatic Logout Timer* on page 105 *Logging out of Sybase Control Center* on page 86
Multiple object selection – in the Perspective Resources view, Resource Explorer, and Administration Console, you can select and perform operations on several objects simultaneously.	Various topics, including *Unregistering a Resource* on page 175
Memory management – use environment variables to control Sybase Control Center's memory use. A new console command, **info -m**, displays memory usage data.	*Configuring Memory Usage* on page 78 *info Command* on page 201

See also

- *User Interface Overview* on page 4
- *Toolbar Icons* on page 5
- *Status Icons* on page 6
- *Common Display Options* on page 7
- *Style and Syntax Conventions* on page 9
- *Accessibility Features* on page 11
- *Sybase Control Center Accessibility Information* on page 12

User Interface Overview

This illustration labels important elements of the Sybase Control Center user interface so you can identify them when they appear in other help topics.

Figure 1: Sybase Control Center User Interface

See also

- *New Features in Sybase Control Center for Sybase IQ* on page 1
- *Toolbar Icons* on page 5
- *Status Icons* on page 6
- *Common Display Options* on page 7
- *Style and Syntax Conventions* on page 9
- *Accessibility Features* on page 11
- *Sybase Control Center Accessibility Information* on page 12

Toolbar Icons

Describes the icons in the Sybase Control Center toolbar for launching and managing views.

Table 2. Toolbar icons

Icon	Name	Description
	Show/Hide Perspective Resources View	Displays or minimizes the Perspective Resources view, which lists registered resources in this perspective.
	Launch Resource Explorer	Opens the resource explorer, which lists reachable resources (both registered and unregistered).
	Launch Heat Chart	Opens the perspective heat chart, which gives a status overview of the registered resources in this perspective.
	Close All Open Views	Closes all open and minimized views.
	Minimize All Views	Minimizes all open views.
	Restore All Minimized Views	Returns all minimized views to their original size.
	Cascade All Open Views	Arranges open views to overlap each other.
	Tile All Open Views Vertically	Arranges open views in a vertical manner.
	Tile All Open Views Horizontally	Arranges open views in a horizontal manner.

See also
- *New Features in Sybase Control Center for Sybase IQ* on page 1
- *User Interface Overview* on page 4
- *Status Icons* on page 6
- *Common Display Options* on page 7
- *Style and Syntax Conventions* on page 9
- *Accessibility Features* on page 11
- *Sybase Control Center Accessibility Information* on page 12

Status Icons

Sybase Control Center uses icons to indicate the status of resources and key performance indicators (KPIs).

Resource status icons indicate the condition of each resource in the heat chart. In addition, they are used as badges (small overlays) on server icons in both the heat chart and the Perspective Resources view. The Perspective Resources view also has a Status column that displays the same status as the badge in English text.

Table 3. Resource status icons: Perspective Resources view and heat chart

Icon	Status	Description
	Running	Resource is up and running
	Pending	State is changing—check again
	Stopped	Resource has been shut down
	Warning	Resource has encountered a potentially harmful situation
	Error	Resource has encountered a serious problem
	Unknown	Resource is unreachable—state cannot be determined

The heat chart uses KPI status icons to indicate the health of the KPIs it displays.

Table 4. KPI status icons: heat chart

Icon	Status	Description
	Normal	Value of performance indicator is within the normal range
	Warning	Value of performance indicator is in the warning range
	Critical	Value of performance indicator is in the critical range

See also
- *New Features in Sybase Control Center for Sybase IQ* on page 1
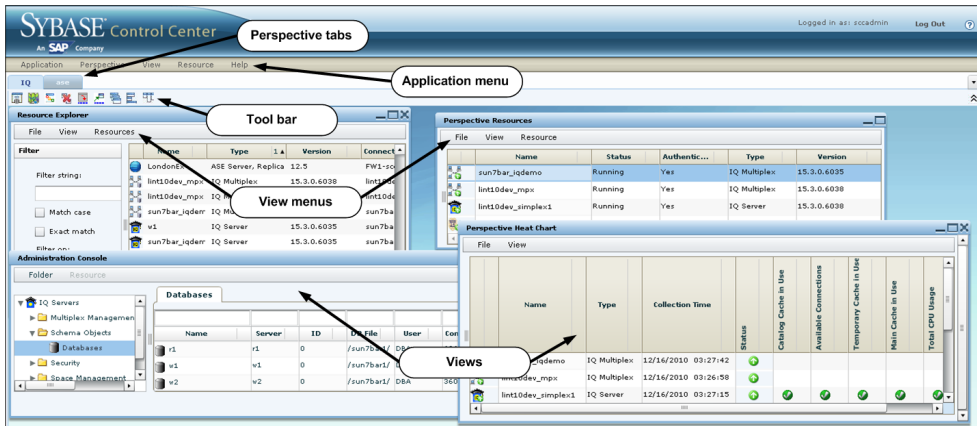- *User Interface Overview* on page 4
- *Toolbar Icons* on page 5
- *Common Display Options* on page 7
- *Style and Syntax Conventions* on page 9
- *Accessibility Features* on page 11
- *Sybase Control Center Accessibility Information* on page 12

Common Display Options

Use data display features to view resource status and to sort, search by resource name and type, and rearrange status information.

Column Options

The Administration Console, Perspective Resources view, Resource Explorer, Alert Monitor, heat chart, and other views in Sybase Control Center—including those in product modules —use a tabular grid format to display information about managed resources. You can use options provided by the grid format to sort and organize displayed data.

Table 5. Column Sorting Options

Sorting Option	Description
Simple column-based sorting	Click a column name to sort the table based on that column in ascending or descending order. The arrow in the column's sorting tab (to the right of the column name) points up when data is sorted in ascending order or down when data is sorted in descending order.
Reversing the order of a column-based sort	Click a column's sorting tab to reverse its sort from ascending to descending order or vice versa.
Nested sorting based on multiple columns	Click the column name for the primary sort. For subsidiary sorts, click the column's sorting tab. Choose the columns for subsidiary sorts in the order you want to apply them. After you click a sorting tab, it displays its sorting level (1 for the primary sort, 2 for the secondary sort, and so on).
Rearranging columns	Move columns by dragging and dropping them.

The figure below shows a table of servers sorted first by resource type; within type by software version; and within version by server name. The Type and Name columns sort in ascending order and the Version column sorts in descending order.

Figure 2: Resources sorted by type, version, and name

	Name 3 ▲	Type 1 ▲	Version 2 ▼
	mira8	ASE Server	15.0.2
	mira9	ASE Server	15.0.2
	LondonDR	ASE Server, Replication Only	12.5
	LondonEx	ASE Server, Replication Only	12.5
	NYEx	ASE Server, Replication Only	12.5
	lint10dev_mpx	IQ Multiplex	15.3.0.6038
	lint10dev_mpx	IQ Multiplex	15.3.0.6038
	sun7bar_iqdemo	IQ Multiplex	15.3.0.6035
	lamd6supt_r2	IQ Server	15.3.0.6038
	lint10dev_cn	IQ Server	15.3.0.6038
	lint10dev_r1	IQ Server	15.3.0.6038

Filter by Column
The Administration Console provides a filtering field at the top of each column. Enter a filtering term to narrow the range of objects displayed. For example:

• Enter the name of a resource at the top of the Name column to display only that server, database, group, or other named object. The display reacts as you enter each character, so you might not need to enter the entire name.
• Enter a version number at the top of the Version column to display only resources running that software version.

You can filter on multiple columns; for example, in a listing of servers, use the Status column to display only running servers, then use the Version column to display running servers using the desired software version. Delete the filtering terms to return to the original display. Filtering terms are not case sensitive.

Full Screen Mode
To increase the screen area available in Sybase Control Center for views and perspectives,

click the ⌃ icon at the upper-right corner of the perspective area. Click the icon again to return to the original screen configuration.

Tip: To increase the screen area available to SCC, press **F11** to switch Internet Explorer or Firefox to full screen mode. Press **F11** again to return to the original browser configuration.

Maximize a Section of a View

Some areas within views have a square minimize/maximize icon (🔲) in the upper-right corner. Click the icon to expand that area to fill the entire view. Click the icon again to restore the area to its former size.

View Menu

The Perspective Resources view, the Resource Explorer, the Alert Monitor, and the heat chart each have a View menu. From the View menu, you can:

- Display the filtering tool for searches. (In the heat chart, the Filter option also displays the column selection tool.)
- Toggle between an icon view and a detail view of your resources (Perspective Resources view only)
- Refresh the display (Resource Explorer only)

Note: For these tasks, use the View menu in the view window, not the application-level View menu.

See also
- *New Features in Sybase Control Center for Sybase IQ* on page 1
- *User Interface Overview* on page 4
- *Toolbar Icons* on page 5
- *Status Icons* on page 6
- *Style and Syntax Conventions* on page 9
- *Accessibility Features* on page 11
- *Sybase Control Center Accessibility Information* on page 12

Style and Syntax Conventions

A reference to the fonts and special characters used to express command syntax and to represent elements of system output and user input.

Table 6. Style Conventions

Key	Definition
`monospaced(fixed-width)`	- SQL and program code - Commands to be entered exactly as shown - File names - Directory names

Key	Definition
italic monospaced	In SQL or program code snippets, placeholders for user-specified values (see example below).
italic	• File and variable names • Cross-references to other topics or documents • In text, placeholders for user-specified values (see example below) • Glossary terms in text
bold sans serif	• Command, function, stored procedure, utility, class, and method names • Glossary entries (in the Glossary) • Menu option paths • In numbered task or procedure steps, user-interface (UI) elements that you click, such as buttons, check boxes, icons, and so on

A placeholder represents a system- or environment-specific value that you supply. For example:

```
installation directory\start.bat
```

where *installation directory* is where the application is installed.

Table 7. Syntax Conventions

Key	Definition
{ }	Curly braces indicate that you must choose at least one of the enclosed options. Do not type the braces when you enter the command.
[ ]	Brackets mean that choosing one or more of the enclosed options is optional. Do not type the brackets when you enter the command.
( )	Parentheses are to be typed as part of the command.
|	The vertical bar means you can select only one of the options shown.
,	The comma means you can choose as many of the options shown as you like, separating your choices with commas that you type as part of the command.
...	An ellipsis (three dots) means you may repeat the last unit as many times as you need. Do not include ellipses in the command.

See also

Accessibility Features

Accessibility ensures access to electronic information for all users, including those with disabilities.

Documentation for Sybase products is available in an HTML version that is designed for accessibility.

Vision impaired users can navigate through the online document with an adaptive technology such as a screen reader, or view it with a screen enlarger.

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

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

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

You may find additional information about accessibility features in the product documentation.

See also

Sybase Control Center Accessibility Information

Sybase Control Center uses the Adobe Flex application.

For the most current information about Adobe Flex keyboard shortcuts, see *http://livedocs.adobe.com/flex/3/html/help.html?content=accessible_5.html*.

Note: To use Sybase Control Center with JAWS for Windows screen reading software effectively, download and install the appropriate Adobe scripts. See *http://www.adobe.com/accessibility/products/flex/jaws.html*.

See also

Get Started

Set up Sybase® Control Center.

Quick Start for an Evaluation

(Optional) Get started using Sybase Control Center quickly if you do not need the full set of security features. This simplified process is suitable for a small-scale, temporary evaluation or proof-of-concept project, or for checking your installation.

Prerequisites

Install Sybase Control Center.

Task

Use these tasks to start Sybase Control Center, log in, register and authenticate a server, and monitor that server.

Note: After completing the tasks below and confirming that SCC is working, set up SCC for a production environment if you intend to continue using it.

1. *Registering the ODBC Driver in Windows*

   In Windows, run scc.bat with administrative privileges to register the ODBC driver.

2. *Launching Sybase Control Center*

   Use the scc command to start Sybase Control Center.

3. *Getting Started After Installing*

   Perform postinstallation testing and configuration.

4. *Configuring Sybase IQ for Monitoring*

   To enable one or more Sybase IQ users to authenticate a Sybase IQ server with Sybase Control Center, create and populate the SCC_MONITOR group.

5. *Registering a Sybase IQ Server*

   Make Sybase Control Center aware of a Sybase IQ resource (for example, a server that can be monitored) and its connection information by registering the resource.

6. *Authenticating a Login Account for a Managed Resource*

   Specify the login account Sybase Control Center will use when it connects to your server or agent to collect monitoring data or manage the resource.

7. *Displaying Resource Availability*

   Use the heat chart to view availability information on the servers in the current perspective.

8. *Viewing Overview Statistics*

Display high-level statistics for the selected Sybase IQ server.

9. *Viewing Multiplex Overview Statistics*

Display the overall health of the Sybase IQ multiplex environment.

See also
- *Get Started in a Production Environment* on page 23

Registering the ODBC Driver in Windows

In Windows, run **scc.bat** with administrative privileges to register the ODBC driver.

When Sybase Control Center starts for the first time on a Windows machine, it registers its ODBC driver. Because the automatic registration of the ODBC driver edits the registry settings, you must execute **scc.bat** using elevated administrative privileges. If you launch for the first time without adequate privileges, Sybase Control Center generates an error and fails to start.

In Windows Vista, Windows 2008, and Windows 7, you must use the **Run as administrator** setting to launch Sybase Control Center even if you already have administrative privileges. This process is described below.

In other versions of Windows, you must be logged in as an administrator to start Sybase Control Center for the first time. You need not follow the steps below.

1. In Windows Vista, Windows 2008, or Windows 7, open the Command Prompt window with administrative privileges:

- Select **Start > All Programs > Accessories.** Right-click **Command Prompt** and select **Run as administrator**.
- Alternatively, enter **cmd** in the Start Menu search box and press **Shift**+**Ctrl**+**Enter**.

2. Run **scc.bat**.

Launching Sybase Control Center

Use the **scc** command to start Sybase Control Center.

Prerequisites

Install Adobe Flash Player in the browser you will use for Sybase Control Center.

Task

1. Start Sybase Control Center.

- Windows – navigate to `<install_location>\SCC-3_2\bin` and double-click **scc.bat**.
- UNIX – execute **scc.sh**.

Messages on the progress of the launch appear in a command window. When Sybase Control Center is running, the command window becomes the Sybase Control Center console; you can issue commands to get status information on SCC and its ports, plug-ins, and services.

2. Open a Web browser and enter `https://<hostname>:8283/scc`.

Getting Started After Installing

Perform postinstallation testing and configuration.

Prerequisites

Start Sybase Control Center.

Task

1. Install Adobe Flash Player 10.1 or later in the Web browser you will use to connect to Sybase Control Center.

   Flash Player is a free plug-in. You can download the latest version from *http://get.adobe.com/flashplayer/*.

   If Flash Player is already installed but you are not sure which version you have, go to the Adobe test site at *http://adobe.com/shockwave/welcome*. Click the link that says **Test your Adobe Flash Player installation**. The version information box on the next page that appears displays your Flash Player version.

2. To connect to Sybase Control Center, direct your browser to:
   ```
   https://<scc_server_hostname>:8283/scc
   ```

   Note: If you changed the default HTTPS port during installation, use the new port number instead of 8283.

3. If you see an error about the security certificate, add Sybase Control Center to your browser's trusted sites zone (Internet Explorer) or add a security exception (Firefox).

4. Log in. Sybase Control Center has two default login accounts:

   - sccadmin – for initial configuration and setting up permanent authentication.
   - sccuser – for testing.

   Neither of these accounts requires a password.

   Note: The sccadmin and sccuser accounts and the simple login module on which they are based are not intended for use in a production environment. Sybase recommends that you pass authentication responsibility to your operating system or to LDAP, as described in the *Sybase Control Center > Get Started > Setting Up Security* section of the online help.

Sybase further recommends that you disable sccadmin and sccuser as soon as you have set up and tested authentication, and that you set passwords on the accounts if you do not plan to set up and test authentication right away.

5. (Optional) Configure passwords or disable sccadmin and sccuser—see the *Sybase Control Center Installation Guide* for instructions.

Configuring Sybase IQ for Monitoring

To enable one or more Sybase IQ users to authenticate a Sybase IQ server with Sybase Control Center, create and populate the SCC_MONITOR group.

Prerequisites
You are logged in to Sybase IQ as a user with DBA authority.

Task
To monitor a resource with SCC, you must authenticate the resource as a Sybase IQ user with DBA authority or membership in the SCC_MONITOR group.

Note: If you are performing a quick start, you need only authenticate your resource with Sybase Control Center using a Sybase IQ account with DBA authority (such as DBA). You can skip the steps below until you do a complete production set-up of SCC.

1. Using Interactive SQL or another SQL command tool, execute **scc_iq_monitor_privileges_setup.sql**, located in the directory `SCC-3_2/plugins/IQMAP`.
The script creates the SCC_MONITOR group and grants a set of permissions.

2. Assign one or more Sybase IQ users or groups to the SCC_MONITOR group. You can do this by either of these methods:

- Using Interactive SQL or another SQL command tool, execute **grant membership in group SCC_MONITOR to** <user/group>
- In a user interface tool such as Sybase Control Center or Sybase Central, add the user to the SCC_MONITOR group

Registering a Sybase IQ Server

Make Sybase Control Center aware of a Sybase IQ resource (for example, a server that can be monitored) and its connection information by registering the resource.

Prerequisites
Ensure that the Sybase IQ server does not have multiple databases. Sybase Control Center for Sybase IQ supports a maximum of one database per server.

Task

1. In the Resource Explorer, select **Resources > Register.**
2. Specify:

Table 8. New resource type details

Field	Description
Resource Name	(Required) Name of the resource to register. Enter the actual name of the server, including uppercase and lowercase letters. If the name registered in Sybase Control Center does not exactly match the server name, some monitoring functions may not work.
Resource Type	Select a resource type: • **IQ Multiplex** – register a Sybase IQ multiplex server. • **IQ Server** – register a Sybase IQ simplex server. Recommended. Can handle both simplex and multiplex servers.
Description	A brief description to help you identify the resource.

3. Click **Next**.
4. Specify the connection information for your resource:

Table 9. New resource connection details

Field	Description
Host Name	(Required) Host name of the Sybase IQ server.
IQ Port Number	(Required) Port number on server host.
Database	Name of the database.
Character Set	Character set to use for the connection.
Language	Language to use for the connection.

5. (Optional) Enter a user name and password that SCC can use to authenticate with this resource to retrieve its software version. The credentials are used only for this purpose, then discarded.

   If you prefer not to authenticate now, click **I do not want to supply authentication information**.

   This step enables SCC to display the correct version information for the server before the server is formally authenticated (later in the configuration process).

6. Click **Next**.

7. (Optional) Click **Add this resource to the current perspective**. You must add a resource to a perspective (not necessarily the current perspective) before you can manage or monitor it.

8. (Optional) Click **Open the resource explorer to view this new resource**. (This option is not present when the Resource Explorer is open.)
   The resource is added to the Resource Explorer even if you choose not to view it.

9. Click **Finish**.

Authenticating a Login Account for a Managed Resource

Specify the login account Sybase Control Center will use when it connects to your server or agent to collect monitoring data or manage the resource.

Perform this task for each resource registered with Sybase Control Center.

Note: You can also authenticate a server during administrative tasks like creating an alert or a collection job.

1. Connect a browser to Sybase Control Center and log in.

2. If the Perspective Resources view is not open, click the **Show/Hide Perspective Resources View** icon in the toolbar.

3. In the Perspective Resources view, select your resource and select **Resource > Authenticate** from the view menu.

4. Select **Use my current SCC login** or **Specify different credentials**.

5. If you chose **Specify different credentials**, enter the login and password for Sybase Control Center to use to connect to your resource.

6. If the selected server is a Replication Server, also enter the RSSD user name and password.

7. Click **OK** to save and exit the dialog.

Displaying Resource Availability

Use the heat chart to view availability information on the servers in the current perspective.

1. From the application menu bar, select **View > Open > Heat Chart.**

2. (Optional) To display tools for filtering (narrowing the list of resources in the heat chart) or changing the columns, select **View > Filter** from the Perspective Heat Chart menu bar. The Filter and Column tools appear in the left pane.

3. (Optional) To use filtering, select **View > Filter** from the view's menu bar and enter a search term in the **Filter string** field.
   The search term can be any string that appears in the tabular portion of the heat chart, such as the name, or part of the name, of a server or a resource type (ASE Server, for example).

4. (Optional) Select a filtering setting:

- **Match case** – search for resources whose displayed data includes the search term, including uppercase and lowercase letters; or
- **Exact match** – search for resources whose displayed data includes an item identical to the search term.

5. (Optional) Select a column from the **Filter on** list to restrict your search to that column.

6. (Optional) Click **Columns** to customize your heat chart.

7. (Optional) Unselect any column that should not appear in your heat chart.

8. (Optional) Click the sorting arrow in the column headers to sort the column values in either ascending or descending order.

9. (Optional) Click the resource's row and pull down the menu to the right of the resource name to view options for the selected resource.

10. (Optional) To resize the Filter and Columns tools pane, move your mouse over the border between the tools pane and the resource table. When the mouse cursor changes to a resize icon, click and drag the border to the left or the right.

11. (Optional) To hide the Filter and Columns tools, unselect **View > Filter.**

Viewing Overview Statistics

Display high-level statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.

2. In the left pane of the IQ Node Level Monitor view, select **Overview**.

3. Select the tab for the required information.

Note: Click a column header to sort the data by that column.

To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Tab	Description
Server	**State** – current status of the server. Valid states include: • Unknown • Stopped • Running **Host** – host name where the server is running. **Port** – port number where the server is running. **Server name** – name of the server. **Database** – name of the Sybase IQ database. **Server type** – type of server the database is on. Server types include Single Server, Coordinator, Reader, and Writer. **Server version** – version of the Sybase IQ server. **Platform** – operating system running on the server host.
Activities	**CPU total usage (%)** – total CPU usage percentage, including both system and user usage. **Active connections** – total number of active connections, including user and internode communication connections. **Connections available** – number of connections available for users and internode communication connections. **Active requests** – number of active requests on the server. **IQ threads in use** – number of threads being used by the Sybase IQ server. **Active transactions** – number of active transactions. **Number of committed transactions** – number of committed transactions. **Oldest transaction (minutes)** – elapsed age, in minutes, of the oldest transaction.
Caches	**Catalog cache reads (per second)** – number of catalog cache page lookups per second. **Main cache size (MB)** – size of the main cache, in megabytes. **Temp cache size (MB)** – size of the temporary cache, in megabytes. **Remaining heap size (MB)** – size of the remaining heap allocation, in megabytes.

Tab	Description
Version us-age	**Number of committed versions** – the number of table versions in the server.
	Total version space used (MB) – total space consumed by all the table versions.
	Oldest version ID – the oldest version identifier on the server.
	Number of active versions – total number of active write table versions on the server.
	Total active version space created (MB) – amount of data created by active write transactions.
	Total active version space to be destroyed (MB) – amount of data destroyed by active write transactions. If these transactions commit, the destroyed data becomes an old version and is eventually dropped. If the transactions roll back, the created data is released.
Details	**Server full version** – version of the IQ server software, including the date and time.
	Platform version – version of the operating system installed on the server host.
Alerts	Any alerts for the selected server. While the montior is open, alerts are displayed as they are created.
CPU histo-ry chart	Percentage of total CPU usage over a period of time.
IQ memo-ry chart	Allocation of the IQ memory between the main cache, temporary cache, and remaining heap.
Disk usage chart	Available and used space for the main and temporary stores.

Viewing Multiplex Overview Statistics

Display the overall health of the Sybase IQ multiplex environment.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Overview**.
3. Select the tab for the required information.

Note: Click a column header to sort the data by that column.

To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Multiplex tab > Servers	**Server** – name of the server. **Host** – host name where the server is running. **Port** – port number where the server is running. **State** – current state of the server. Valid states include: • Unknown • Stopped • Running **Role** – role the server plays in the multiplex configuration. Roles include: • Coordinator • Reader • Writer **Status** – current status of the server in the multiplex. Valid states include: • Included • Excluded
Multiplex tab > CPU History	Percentage of total CPU usage over a period of time for each server. The legend below the chart identifies the colored line associated with each server.
Multiplex tab > IQ Memory	Allocation of the IQ memory between the main cache and temporary cache for each server in the multiplex.
Disk Usage tab	Available and used space for the main store and temporary store on each server in the multiplex.
Version Usage tab > Statistics	**Number of committed versions** – the number of table versions in the servers. **Total version space used (MB)** – total space consumed by all the table versions. **Oldest version ID** – the oldest table version identifier on the server. **Number of active versions** – total number of active write table versions on the servers. **Total active version space created (MB)** – amount of data created by active write transactions. **Total active version space to be destroyed (MB)** – amount of data destroyed by active write transactions. If these transactions commit, the destroyed data becomes an old version and is eventually dropped. If the transactions roll back, the created data is released.

Area	Description
Version Usage tab > Multiplex Version Usage	**Version ID** – the table version identifier.
	Server name – the name of the server where the table version exists.
	Connection ID – the connection ID using this table version.
	WasReported – indicates whether the server has received usage information for this table version.
	MinKBRelease – the minimum amount of space returned once this table version is no longer in use.
	MaxKBRelease – the maximum amount of space returned once this table version is no longer in use.

Get Started in a Production Environment

Perform a complete set-up of Sybase Control Center, including configuration of user authentication and other one-time set-up tasks.

Prerequisites

Install Sybase Control Center and complete the follow-up tasks described in the *Sybase Control Center Installation Guide*.

1. *Deploying an Instance from a Shared Disk Installation*

   (Optional) Create a Sybase Control Center server or agent from an installation on a shared disk.

2. *Starting and Stopping Sybase Control Center in Windows*

   There are several ways to start and stop Sybase Control Center or the SCC agent. You can start manually, which is useful for testing and troubleshooting, or set the service to start automatically and to restart in case of failure.

3. *Starting and Stopping Sybase Control Center in UNIX*

   You can start Sybase Control Center or the SCC agent manually, which is useful for testing and troubleshooting, or you can set up a service to start automatically and to restart in case of failure.

4. *Configuring Memory Usage*

   (Optional) Determine whether you need to configure how much memory Sybase Control Center uses, and if so which configuration method to use.

5. *Logging in to Sybase Control Center*

   Enter the Sybase Control Center Web console.

6. *Setting Up Security*

Configure login authentication and map roles.

7. *Configuring the E-mail Server*

(Optional) Specify the e-mail server for Sybase Control Center to use to send e-mail alert notifications.

8. *Configuring the Automatic Logout Timer*

(Optional) Set Sybase Control Center to end login sessions when users are inactive for too long.

9. *User Authorization*

The authorization mechanism in Sybase Control Center employs login accounts and task-based roles.

10. *Configure*

Configure Sybase Control Center for Sybase IQ.

Deploying an Instance from a Shared Disk Installation

(Optional) Create a Sybase Control Center server or agent from an installation on a shared disk.

Prerequisites

- Install Sybase Control Center on a shared disk.
- Enable shared-disk mode.

Task

1. Log in to the host on which you plan to run the SCC server or agent.

> **Note:** You can create an instance on one host and run it on another host, but doing so interferes with the predeployment checks run by **sccinstance**. Such a deployment might generate errors (port conflicts, for example). If you are confident that the errors are caused by problems that will not be present on the host where you plan to run the instance, use the **-force** option to create the instance.

2. Change to `SCC-3_2/bin`.

3. Create the instance as an SCC agent if you plan to run a managed server on this host. Create the instance as an SCC server if you plan to manage other Sybase servers from this host.

To create an SCC agent called Boston-agent and configure it to run as a Windows service:

```
sccinstance -create -agent -instance Boston-agent -service
```

To create an SCC server called Boston and configure it to run as a Windows service:

```
sccinstance -create -server -instance Boston -service
```

4. If other SCC instances will run on this host, change the port assignments for the new instance. Change the instance names and port values in the sample commands to suit your environment, but take care to specify ports that are not in use by another SCC instance or any other application or server.

This command changes the port assignments for an SCC agent called myagent:

```
sccinstance -refresh -instance myagent -portconfig
rmi=8888,jiniHttp=9093,jiniRmi=9096,tds=9997
```

This command changes the port assignments for an SCC server called myserver:

```
sccinstance -refresh -server -instance myserver -portconfig
rmi=8889,db=3640,
http=7072,https=7073,jiniHttp=9094,jiniRmi=9097,msg=2002,tds=9996
```

5. (Optional) List the instances deployed from this installation:

```
sccinstance -list
```

6. (Optional) If you are setting up an instance in UNIX, configure it to run as a service. (See *Starting and Stopping Sybase Control Center in UNIX*).

Next

When you manage and maintain instances, keep in mind that the directory structure for instances is different from that of singleton installations. In file paths in SCC help, replace `SCC-3_2` or `<scc-install-directory>` with `SCC-3_2/instances/<instance-name>`.

For example, the path to the log directory, `SCC-3_2/log`, becomes this for an instance called kalamazoo:

```
SCC-3_2/instances/kalamazoo/log
```

See also
* *Starting and Stopping Sybase Control Center in Windows* on page 31

Enabling and Disabling Shared-Disk Mode
Turn on or turn off shared-disk mode, which allows you to run multiple Sybase Control Center agents and servers from a single installation on a shared disk.

Prerequisites
Install Sybase Control Center on a shared disk. See the *Sybase Control Center Installation Guide*.

Task
Shared-disk mode affects the entire installation; do not enable or disable individual instances.

Disabling shared-disk mode leaves the instances' file systems intact under `<SCC-install-directory>/instances`, but the instances cannot run. If you reenable, the instances are able to run again.

1. Change to `SCC-3_2/bin`.
2. Enable or disable shared disk mode.

   To enable shared disk mode:

   ```
   sccinstance -enable
   ```

   To disable shared disk mode:

   ```
   sccinstance -disable
   ```

Shared-Disk Mode

Shared-disk mode lets you run multiple Sybase Control Center instances—SCC servers, SCC agents, or a mixture of the two—from a single installation of the product.

The shared-disk capability enables SCC servers or agents on the installation host or on remote hosts to access and execute from the same installation. This feature is especially useful if you plan to use SCC to manage Adaptive Server clusters or Sybase IQ multiplexes.

After installing SCC on a shared disk, use the **sccinstance** command to enable shared-disk mode and deploy instances. **sccinstance** copies the files needed for the instance into a new directory structure. The path takes the form `<SCC-install-directory>/instances/<instance-name>` (for example, `SCC-3_2/instances/SCCserver-1`).

You can specify a name for each instance. If you do not supply a name, the instance name defaults to the host name.

An instance runs on the host on which you start it. When shared-disk mode is enabled, SCC servers and agents run out of the `SCC-3_2/instances` subdirectories, not from the base file system.

In shared-disk mode, changes made to configuration files in the base file system (everything under `SCC-3_2` except the `SCC-3_2/instances` branch) are copied to any instance deployed thereafter. Previously deployed instances are not affected.

Use **sccinstance** to deploy, remove, refresh, or convert an instance; to configure an instance's ports; and to configure a Windows instance to run as a service. Perform other tasks, including configuring a UNIX instance to run as a service, and all other configuration, using the tools and procedures documented for all installations. Use tools provided by the UI wherever possible. When you must edit a file to change the configuration of an instance (for role mapping, for example), edit the copy of the file stored under `<SCC-install-directory>/instances/<instance-name>`.

<u>**sccinstance Command**</u>

Use **sccinstance.bat** (Windows) or **sccinstance** (UNIX) to deploy an instance of Sybase Control Center from a shared-disk installation or to manage existing instances.

You can run multiple instances of Sybase Control Center, including SCC servers, SCC agents, or a mixture of the two, from a single installation on a shared disk.

Syntax

```
sccinstance[.bat]
[-agent]
[-c | -create]
[-d | -debug]
[-disable]
[-enable]
[-f | -force]
[-h | -help]
[-i | -instance [instance-name]]
[-l | -list]
[-plugins {plugin-ID,plugin-ID,...}]
[-portconfig {port-name=port-number,port-name=port-number, ...}]
[-refresh]
[-r | -remove]
[-s | -server]
[-service]
[-silent]
```

Parameters

- **-agent** – use with **-create** or **-refresh** to create or refresh an SCC agent. In a **-create** or **-refresh** command, **-agent** is the default, so you can omit it.
- **-create** – deploy a new instance. Use alone or with **-agent** to create an agent instance, or with **-server** to create a server instance.
- **-d | debug** – display debugging messages with the output of this command.
- **-disable** – turn off shared-disk mode for this installation. Generates an error if any instance is running.
- **-enable** – turn on shared-disk mode for this installation. Shared-disk mode is required if you intend to run more than one server or agent from a single installation of SCC.
- **-f | -force** – execute **sccinstance** even if there are potential conflicts (such as port clashes or a running SCC process).
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-instance** – specify an instance. Use with **-create**, **-remove**, or **-refresh**, or use alone to display the instance's status. You can omit **-instance** when you are addressing the only SCC instance or the only instance of the specified type (server or agent) on the current host.
- **-l | -list** – display a list of all instances deployed from this SCC installation.

- **-plugins {*plugin-ID,plugin-ID,...*}** – specify one or more product module plug-ins for this instance. An alternative to **-agent** and **-server**, **-plugins** is primarily for use by the SCC installation program. Use with **-create** or **-refresh**. Use commas to separate plug-in names.
- **-portconfig {*port-name=port-number*, *port-name=port-number*, ...}** – assign ports to services for this instance. Use only with **-create** or **-refresh**. For the *port-name* value, use a port name from the table below. If you plan to run more than one SCC instance on a host machine, you must reassign all the ports for every instance after the first.

Port information:

Port Name	Description	Service Names	Property Names	Default Port
db	Database port Present on SCC server	SccSADataserver Messaging Alert Scheduler	com.sybase.asa.server.port messaging.db.port alert.database.port org.quartz.data-Source.ASA.URL	3638
http	Web HTTP port Present on SCC server	EmbeddedWebCon-tainer	http.port	8282
https	Web HTTPS (secure HTTP) port Present on SCC server	EmbeddedWebCon-tainer	https.port	8283
jiniHttp	JINI HTTP server Present on SCC server and SCC agent	Jini	httpPort	9092
jiniR-mid	JINI remote method in-vocation daemon Present on SCC server and SCC agent	Jini	rmidPort	9095
msg	Messaging port Present on SCC server	Messaging	messaging.port	2000
rmi	RMI port Present on SCC server and SCC agent	RMI	port	9999

Port Name	Description	Service Names	Property Names	Default Port
tds	Tabular Data Stream™ port (used to communicate with other Sybase products) Present on SCC server and SCC agent	Tds	tdsPort	9998

- **-refresh** – recopy all the files that make up this instance (Windows) or all this instance's services and plug-ins (UNIX). Refreshing preserves any service or plug-in configuration in the deployed instance.

  You can also use **-refresh** to convert a server to an agent or an agent to a server (see the examples). Files are removed or added to change the function of the instance. Use alone or with **-agent** to refresh an agent instance, or with **-server** to refresh a server instance. Generates an error if the instance is running.

- **-r | -remove** – delete an instance. Use alone or with **-instance**. Generates an error if the instance is running. You cannot restore a removed instance.

- **-s | -server** – use with **-create** or **-refresh** to create or refresh an SCC server, including any product modules available.

- **-service** – use with **-create** or **-remove** to create or remove a Windows service for this instance. You must be logged in to Windows as an administrator to use this option.

- **-silent** – suppress the output of **sccinstance**.

Examples

- **Deploy an SCC server instance** – enables shared-disk mode, deploys a server called Boston with a Windows service, and starts the Windows service:

```
sccinstance -enable
sccinstance -create -server -instance Boston -service
net start "Sybase Control Center 3.2.3 (Boston)"
```

Note: To create the service, you must log in to Windows as an administrator.

- **Deploy an SCC agent instance** – deploys an SCC agent on this host and configures a Windows service for it. The **-agent** option, because it is the default, is not required—the command does exactly the same thing without it.

```
sccinstance -create -agent -service
```

or

```
sccinstance -create -service
```

- **Deploy a server instance and reassign ports** – deploys the server on this host and configures nondefault RMI, HTTP, and HTTPS ports.

```
sccinstance -create -server -portconfig
rmi=8888,http=7070,https=7071
```

- **Refresh a server instance or convert an agent to a server** – refreshes the server on this host. If the instance on this host is an SCC agent, refreshing it as an SCC server converts it into a server.

```
sccinstance -refresh -server
```

- **Refresh an agent instance or convert a server to an agent** – refreshes the instance named kalamazoo. If kalamazoo is a server, refreshing it as an SCC agent converts it into an agent.

```
sccinstance -refresh -agent -instance kalamazoo
```

- **Remove a server instance** – removes the instance named porcupine if it is not running:

```
sccinstance -remove -instance porcupine
```

- **Display status** – displays the status of the instance on this host:

```
sccinstance
```

- **List all instances** – displays a list of all SCC server and agent instances deployed from this SCC installation:

```
sccinstance -list
```

- **Scenario: Remove an instance by force** – suppose you have inadvertently deployed two SCC agent instances on the same host:

```
$ sccinstance -list
2 SCC instances deployed:
SCC instance node1 deployed in agent mode for host node1 RMI port
9999
SCC instance node2 deployed in agent mode for host node2 RMI port
9999
```

Both instances use the same RMI port. You must either reassign ports for one instance or remove it. But you get an error if you try remove an instance when another instance is running on the same host:

```
$ sccinstance -instance node2 -remove
[ERROR] Command execution failed.
[ERROR] SCC instance node2 could not be removed because it is
running. Shut
down the SCC before removing the instance.
```

Use the **-force** option to override the error and force the removal of the second agent instance:

```
$ sccinstance -instance node2 -remove -force
Removing SCC instance node2 ...
SCC instance node2 was successfully removed.
```

Permissions

sccinstance permission defaults to all users, except as noted for certain parameters.

Starting and Stopping Sybase Control Center in Windows

There are several ways to start and stop Sybase Control Center or the SCC agent. You can start manually, which is useful for testing and troubleshooting, or set the service to start automatically and to restart in case of failure.

This topic applies to both Sybase Control Center (the server) and the Sybase Control Center agent that runs on each product server managed by SCC. It applies to both singleton installations and instances of SCC agents and servers running from a shared disk.

If you run Sybase Control Center or the SCC agent manually, you must issue a command every time you start or shut down. If you run as a service (which is recommended), you can configure the service to start and restart automatically. These are the options:

- Use the **scc.bat** command to start Sybase Control Center or the SCC agent manually. The command gives you access to the Sybase Control Center console, which you can use to shut down and to display information about services, ports, system properties, and environment variables. You can also use **scc.bat** to change the logging level for troubleshooting purposes. Using **scc.bat** prevents you from taking advantage of the automatic start and restart features available to services.
- Use the Services list under the Windows Control Panel to start, stop, and configure the Sybase Control Center service for an SCC server or agent.
- Use the **net start** and **net stop** commands. This is another way to run Sybase Control Center or the SCC agent as a service.

  Note: To start an SCC agent or server as a service:
  - In a singleton installation, you must have selected **Yes** in the installer to install the agent or server as a service.
  - In a shared disk installation, the agent or server must have been deployed using the **-service** option of the **sccinstance** command.

In a singleton installation, the installer lets you start Sybase Control Center or the SCC agent as a service and configures the service to restart automatically. Before starting, check the Windows Services list for a Sybase Control Center service.

Here are the steps for each starting and stopping option:

- **Start Sybase Control Center or the SCC agent:**
  a) (Skip this step for the SCC agent.) If you are starting Sybase Control Center for the first time in Windows Vista, Windows 2008, or Windows 7, set the **Run as Administrator** option on the command prompt so that Sybase Control Center can register its ODBC driver. (This is necessary even if you are logged in as an administrator.)
  b) Enter the **scc** command.

     For a singleton installation:
     ```
     %SYBASE%\SCC-3_2\bin\scc.bat
     ```

For an instance:

```
%SYBASE%\SCC-3_2\bin\scc.bat -instance <instance-name>
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

- **Stop Sybase Control Center or the SCC agent:**
  a) Enter the **scc --stop** command.

  For a singleton installation:

  ```
  %SYBASE%\SCC-3_2\bin\scc.bat --stop
  ```

  For an instance:

  ```
  %SYBASE%\SCC-3_2\bin\scc.bat --stop -instance <instance-
  name>
  ```

  You can omit the **-instance** option if the instance's name is the same as its host name (the default).

  Note: You can also enter **shutdown** at the `scc-console>` prompt.

- **Start or stop from the Windows Control Panel; configure automatic start and restart:**
  a) Open the Windows Control Panel.
  b) Select **Administrative Tools > Services.**
  c) Locate "Sybase Control Center" in the Services list. It may be followed by a release number; if the service is for an instance, it is also followed by the instance name. Service names do not distinguish between agents and servers. If the service is running, the Status column displays "Started."
  d) To start or stop the service, right-click the **Sybase Control Center** entry in the Services list and choose **Start** or **Stop**.
  e) To configure automatic starting, double-click the service.
  f) To set the service to automatically start when the machine starts, change the **Startup type** to Automatic.
  g) To restart the service in case of failure, choose the **Recovery** tab and change the First, Second, and Subsequent failures to Restart Service.
  h) Click **Apply** to save the modifications and close the dialog.
- **Start or stop the Sybase Control Center service (controlling either Sybase Control Center or the SCC agent) from the Windows command line:**
  a) To start the service, enter the **net start** command.

  For a singleton installation:

  ```
  net start "sybase control center 3.2.4"

  The Sybase Control Center 3.2.4 service is starting......
  ```

```
The Sybase Control Center 3.2.4 service was started
successfully.
```

For an instance, include the instance name in parentheses:

```
net start "sybase control center 3.2.4 (Boston-1)"

The Sybase Control Center 3.2.4 (Boston-1) service is
starting......
The Sybase Control Center 3.2.4 (Boston-1) service was
started successfully.
```

b) To stop the service, enter the **net stop** command.

For a singleton installation:

```
net stop "sybase control center 3.2.4"

The Sybase Control Center 3.2.4 service is stopping.....
The Sybase Control Center 3.2.4 service was stopped
successfully.
```

For an instance, include the instance name in parentheses:

```
net stop "sybase control center 3.2.4 (Boston-1)"

The Sybase Control Center 3.2.4 (Boston-1) service is
stopping.....
The Sybase Control Center 3.2.4 (Boston-1) service was
stopped successfully.
```

See also
* *Deploying an Instance from a Shared Disk Installation* on page 24

Starting and Stopping Sybase Control Center in UNIX

You can start Sybase Control Center or the SCC agent manually, which is useful for testing and troubleshooting, or you can set up a service to start automatically and to restart in case of failure.

This topic applies to both Sybase Control Center (the server) and the Sybase Control Center agent that runs on each product server managed by SCC. It applies to both singleton installations and instances of SCC agents and servers running from a shared disk.

If you start Sybase Control Center or the SCC agent manually, you must issue a command every time you start or shut down. If you run as a service (which is recommended), you can configure the service to start and restart automatically. These are the options:

* Use the **scc.sh** script to start Sybase Control Center or the SCC agent manually. You can either:

- Run **scc.sh** in the foreground to get access to the Sybase Control Center console, which you can use to shut down and to display information about services, ports, system properties, and environment variables.
- Run **scc.sh** in the background to suppress the console.

You can use **scc.sh** to run Sybase Control Center at a nondefault logging level for troubleshooting. When you start manually with **scc.sh**, you cannot take advantage of the automatic start and restart features available to services.

- Use the **sccd** script to configure a service that starts Sybase Control Center or the SCC agent automatically.

Here are the steps for each starting and stopping option:

- **Before you start Sybase Control Center or the SCC agent for the first time, set environment variables.** Do this only once.

  a) Change to the Sybase directory (the parent of the Sybase Control Center installation directory).

  b) Execute one of the following to set environment variables.

     Bourne shell:
     ```
     . SYBASE.sh
     ```

     C shell:
     ```
     source SYBASE.csh
     ```

- **Run Sybase Control Center or the SCC agent in the foreground.**

  Running in the foreground is a method of manually starting; you must issue comands to stop and restart Sybase Control Center or the SCC agent.

  a) To start Sybase Control Center or the SCC agent and drop into the console when the start-up sequence is finished, enter the **scc** command.

     For a singleton installation:
     ```
     $SYBASE/SCC-3_2/bin/scc.sh
     ```

     For an instance:
     ```
     $SYBASE/SCC-3_2/bin/scc.sh -instance <instance-name>
     ```

     You can omit the **-instance** option if the instance's name is the same as its host name (the default).

- **Run Sybase Control Center or the SCC agent in the background.**

  You can use **nohup**, **&**, and **>** to run Sybase Control Center or the SCC agent in the background, redirect output and system error to a file, and suppress the SCC console. Running in the background is a method of manually starting; you must issue comands to stop and restart Sybase Control Center or the SCC agent.

  a) Execute a command similar to the sample below that matches your shell. Both sample commands direct output to the file scc-console.out. If the output file already

exists, you might need to use additional shell operators to append to or truncate the file.

Bourne shell (sh) or Bash

For a singleton installation:

```
nohup ./scc.sh 2>&1 > scc-console.out &
```

For an instance:

```
nohup ./scc.sh -instance <instance-name> 2>&1 > scc-console-
your-instance.out &
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

C shell

For a singleton installation:

```
nohup ./scc.sh >& scc-console.out &
```

For an instance:

```
nohup ./scc.sh -instance <instance-name> >& scc-console.out &
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

- **Shut down Sybase Control Center or the SCC agent.**
  - a) To shut down from the scc-console> prompt, enter:

    ```
    shutdown
    ```

    Warning! Do not enter **shutdown** at a UNIX prompt; it shuts down the operating system.

    To shut down from the UNIX command line, enter the **scc --stop** command.

    For a singleton installation:

    ```
    $SYBASE/SCC-3_2/bin/scc.sh --stop
    ```

    For an instance:

    ```
    $SYBASE/SCC-3_2/bin/scc.sh --stop -instance <instance-
    name>
    ```

    You can omit the **-instance** option if the instance's name is the same as its host name (the default).

- **Configure Sybase Control Center or the SCC agent to run as a service.**

  A UNIX service is a daemon process that starts automatically after the machine is started and runs in the background. UNIX installations of Sybase Control Center include a shell script, **sccd**, which you can use to configure the Sybase Control Center service. (Some UNIX platforms supply tools that make service configuration easier; Linux **chkconfig** is an example.)

> **Note:** Sybase recommends that if you are not familiar with setting up services in UNIX, you delegate this task to a system administrator or consult the system administration documentation for your UNIX platform.

a) Copy `$SYBASE/SCC-3_2/bin/sccd` into this directory:

- AIX (SCC agent only): `/etc/rc.d/init.d`
- HP-UX (SCC agent only): `/sbin/init.d`
- All other platforms: `/etc/init.d`

b) Open `sccd` and make these changes:

- Change the line that sets the SYBASE variable to the location of your Sybase installation (that is, the parent of `SCC-3_2`, the Sybase Control Center installation directory). By default, this directory is called `Sybase`.
- If you are not using shared-disk mode, or you are using shared-disk mode to run a single instance whose name is the same as the host name, skip to step *5.c* on page 36 or step *5.d* on page 36.
- If you are using shared-disk mode to run a single instance whose name is not the host name, or to run multiple instances on the same host, add the instance name to the script name. Change:

```
SCRIPT_NAME=scc.sh
```

to:

```
SCRIPT_NAME="scc.sh -instance <instance-name>"
```

- If you are using shared-disk mode to run multiple instances on the same host, append the instance name to the name of the ouput log file. Change:

```
./${SCRIPT_NAME} --start 2>&1 >> ${SCC_HOME}/log/scc-
service.out &
```

to:

```
./${SCRIPT_NAME} --start 2>&1 >> ${SCC_HOME}/log/scc-
service_<instance-name>.out &
```

- If you are using shared-disk mode to run multiple instances on the same host, save a copy of the `sccd` script for each instance, giving each copy a unique name. In each copy, add the instance name to the script name and append the instance name to the output log file name as described above. Perform the remaining steps in this procedure for each copy of `sccd`.

c) In Linux, configure the service to run in run levels 2, 3, 4, and 5:

```
/usr/sbin/chkconfig --add sccd
/usr/sbin/chkconfig --level 2345 sccd
```

You can test the `sccd` script with `/usr/sbin/service sccd status`. (The **service** command accepts these options: **start** | **stop** | **status** | restart.)

d) On non-Linux platforms, locate this directory:

- AIX (SCC agent only): `/etc/rc.d/rc<X>.d`

- HP-UX (SCC agent only): `/sbin/rc<X>.d`
- Solaris: `/etc/rc<X>.d`

where <X> is the run level (for example, 3). Make two soft links in the directory for your platform and set the links to point to:

- AIX (SCC agent only):
  `/etc/rc.d/init.d/sccd: S90sccd` and
  `/etc/rc.d/init.d/sccd: K10sccd`
- HP-UX (SCC agent only):
  `/sbin/init.d/sccd: S90sccd` and
  `/sbin/init.d/sccd: K10sccd`
- Solaris:
  `/etc/init.d/sccd: S90sccd` and
  `/etc/init.d/sccd: K10sccd`

The `S90sccd` link starts the service and the `K10sccd` link stops the service. The two-digit numbers in the links indicate the start and stop priorities of the service.

e) Use the `S90sccd` and `K10sccd` links to test starting and stopping the service. The links are called automatically when the machine is started or shut down.

Configuring Memory Usage

(Optional) Determine whether you need to configure how much memory Sybase Control Center uses, and if so which configuration method to use.

It is not usually necessary to configure memory usage for Sybase Control Center. This table lists memory options you can set and circumstances under which you should consider changing them.

Modify this value	When	Guidelines
Maximum memory • `jvmopt=-Xmx` – if you are running SCC as a Windows service • *SCC_MEM_MAX* – if you are running SCC as a UNIX service • *SCC_MEM_MAX* – if you are starting SCC from the command line	• You need to prevent Sybase Control Center from using more than a given amount of memory • SCC fails to start and may display an error: `Could not create the Java Virtual machine.` • An OutOfMemory error says SCC is out of heap space • A warning message about system memory appears during the start process • The machine where SCC is installed has less than 2GB of memory. (Starting SCC on a machine with less than 2GB of memory triggers the startup warning message about system memory.)	On machines with less than 2GB of memory, set maximum memory to 256MB or more. Default value: none. (On machines with 2GB or more of memory, maximum memory is set dynamically and is effectively limited only by the amount of system memory available.)
Permanent memory • `jvmopt=-XX:MaxPerm-Size` – if you are running SCC as a Windows service • *SCC_MEM_PERM* – if you are running SCC as a UNIX service • *SCC_MEM_PERM* – if you are starting SCC from the command line	An OutOfMemory error says SCC is out of permanent generation space	Increase by 32MB increments. If you reach a value equal to twice the default and still see the OutOfMemory error, contact Sybase technical support. Default value: 128MB

You can change memory options in two ways:

• For Sybase Control Center started from the command line – execute commands to set one or more environment variables before executing the **scc** command to start Sybase Control Center. When you use this method, your changes to the memory options last only as long as the current login session. This method is useful for testing new option values.

- For the Sybase Control Center service – modify a file used by the SCC service. When you use this method, your changes to the memory options persist—Sybase Control Center uses them every time it starts as a service.

See also
- *Logging in to Sybase Control Center* on page 40

Changing a Memory Option on the Command Line

Before you start Sybase Control Center from the command line, you can issue a command to change the value of a memory option temporarily.

Changes made using this method last only as long as the current login session. This method is useful for testing new option values.

1. If Sybase Control Center is running, shut it down.
2. Set the environment variable. Specify a size in megabytes but do not indicate the units in the command.

   Windows example:
   ```
   > set SCC_MEM_MAX=512
   ```

   UNIX example:
   ```
   bash$ export SCC_MEM_MAX=512
   ```
3. Use the **scc** command to start Sybase Control Center.

See also
- *Changing a Memory Option for an SCC Windows Service* on page 39
- *Changing a Memory Option for an SCC UNIX Service* on page 40

Changing a Memory Option for an SCC Windows Service

Add a **jvmopt** command to the `scc.properties` file to change a memory option (-Xmx or -XX:MaxPermSize) for a Sybase Control Center Windows service.

When you use this method to set memory options, your changes are permanent—Sybase Control Center uses them every time it starts as a service.

1. If Sybase Control Center is running, shut it down.
2. Open the SCC properties file:
   ```
   <SCC-install-directory>\SCC-3_2\bin\scc.properties
   ```
3. Add (or modify, if it already exists) a **jvmopt** line specifying the memory size in Java format. Use m for megabytes or g for gigabytes.

   For example:
   ```
   jvmopt=-Xmx512m
   ```

4. Save the file and start the Sybase Control Center Windows service.

See also
- *Changing a Memory Option on the Command Line* on page 39
- *Changing a Memory Option for an SCC UNIX Service* on page 40

Changing a Memory Option for an SCC UNIX Service
To change a memory setting for a Sybase Control Center UNIX service, add the appropriate environment variable (*SCC_MEM_MAX* or *SCC_MEM_PERM*) to the sccd script.

When you use this method to set memory options, your changes are permanent—Sybase Control Center uses them every time it starts as a service.

1. If Sybase Control Center is running, shut it down.

2. Open the sccd file: `/etc/init.d/sccd`

3. Add the environment variable at the top of the file (after the comments). Specify a size in megabytes but do not indicate the units in the command.

For example:

```
SCC_MEM_MAX=512
```

4. Save the file and start the Sybase Control Center UNIX service.

See also
- *Changing a Memory Option on the Command Line* on page 39
- *Changing a Memory Option for an SCC Windows Service* on page 39

Logging in to Sybase Control Center

Enter the Sybase Control Center Web console.

Prerequisites
Install Adobe Flash Player in the browser you will use for SCC. See the *Sybase Control Center Installation Guide*.

Task
Sybase Control Center typically authenticates users through the operating system or an LDAP directory service. Consult your SCC administrator if you are not sure which login account to use for SCC.

Note: When logging in to a newly installed Sybase Control Center for which secure authentication has not been configured, use the sccadmin account (with no password, by default). For more information, see the *Sybase Control Center Installation Guide*.

1. Connect to the Sybase Control Center server. In your Web browser, enter: `https://scc-hostname:8283/scc`.

2. Enter your user name and password, and click **Login**.

   Tip: If you use a Windows account to log in to SCC, enter your user name in the format username@domain. Omit top-level domain extensions such as .com or .net—for example, enter fred@sybase, not fred@sybase.com.

   See also
   • *Configuring Memory Usage* on page 37

Setting Up Security

Configure login authentication and map roles.

Read about security and follow these procedures before you configure Sybase Control Center product modules.

Note: These security topics are intended for use in a production environment. If you are evaluating or testing SCC, see the *Installation Guide* for instructions on getting started quickly.

1. *Security*

   Sybase Control Center can authenticate user logins through an LDAP server, through the operating system, or both.

2. *Configuring Authentication for Windows*

   Authentication through the Windows operating system is enabled by default, but it requires some configuration. First, set Sybase Control Center to create an account when a Windows user logs in to Sybase Control Center.

3. *Configuring a Pluggable Authentication Module (PAM) for UNIX*

   Set up Sybase Control Center to support username and password login using accounts on the UNIX operating system. Optionally, have Sybase Control Center create an account when a UNIX user first logs in to Sybase Control Center.

4. *Configuring an LDAP Authentication Module*

   Configure an LDAP authentication module for Sybase Control Center by editing the security properties file to point to the correct LDAP server.

5. *Mapping Sybase Control Center Roles to LDAP or OS Groups*

   To grant Sybase Control Center privileges to users who are authenticated through LDAP or the operating system, associate roles used in Sybase Control Center with groups in LDAP or the operating system.

6. *Encrypting a Password*

   Use the passencrypt utility to encrypt passwords and other values that must be kept secure while stored in text files.

7. *Configuring Ports*

(Optional) Use the scc --port command to assign Sybase Control Center services to new ports.

See also
* *Configuring the E-mail Server* on page 58

Security

Sybase Control Center can authenticate user logins through an LDAP server, through the operating system, or both.

* Sybase Control Center can be configured to authenticate through any LDAP server that supports the inetOrgPerson (RFC 2798) schema.
* When Sybase Control Center authenticates through the operating system, it uses the operating system of the Sybase Control Center server machine (not the client).

Although you can create native user accounts in Sybase Control Center, Sybase does not recommend this approach to authentication. It is simpler and safer to configure Sybase Control Center to authenticate using existing LDAP, Windows, or UNIX login accounts.

Sybase strongly recommends that you use a common authentication provider for all Sybase products, including Sybase Control Center. A common authentication provider ensures that single sign-on works for users of Sybase Control Center and its managed servers.

Sybase Control Center requires each authenticated login account to have a predefined role. When a login is authenticated, roles for the login are retrieved by the security module and are mapped to Sybase Control Center predefined roles. Authorization is resolved through the mappings between the security module native roles and Sybase Control Center roles. You can enable mappings by creating a "sybase" group in your operating system or LDAP server and adding all Sybase Control Center users, or by modifying the Sybase Control Center `roles-map.xml` file to configure the mapping of native roles to Sybase Control Center roles. The security module authenticates the logins and authorizes access to managed resources.

Sybase Control Center provides a set of predefined login modules for authentication. All login modules are defined in the `<install_location>/SCC-3_2/conf/csi.properties` file. The syntax is defined by the Sybase Common Security Infrastructure (CSI) framework. You can configure the different login modules to customize security strength. The login modules are:

* Simple Login – defines a user name, password, and a list of roles. The default user name is "sccadmin" with a blank password and a native role of "sccAdminRole". You can create additional accounts by adding simple login modules to `csi.properties`. However, Sybase does not recommend the use of simple login modules for authentication in production environments.

> **Note:** Add a password for the sccadmin account as soon as possible after you install
> Sybase Control Center. See the *Sybase Control Center Installation Guide* for instructions.

- NT Proxy Login – delegates authentication to the underlying Windows operating system.
  When you log in to Sybase Control Center through an NT Proxy Login module, enter your
  user name in the format *username@nt-domain-name*. For example, user@sybase.
  Windows authentication is enabled by default, but it requires some configuration.
- UNIX Proxy Login – delegates authentication to the underlying UNIX or Linux operating
  system using Pluggable Authentication Modules (PAM). When you log in to Sybase
  Control Center through a UNIX PAM, enter your UNIX user name and password. UNIX
  authentication is enabled by default, but it requires some configuration.
- LDAP Login – delegates authentication to an LDAP server you specify. When you log in to
  Sybase Control Center through an LDAP server, enter your LDAP user name and
  password. LDAP authentication is not enabled by default; you must configure the login
  module.

Configuring Authentication for Windows

Authentication through the Windows operating system is enabled by default, but it requires
some configuration. First, set Sybase Control Center to create an account when a Windows
user logs in to Sybase Control Center.

This task is optional. However, if you choose not to create Sybase Control Center accounts
automatically as described here, you must enter them manually. Sybase Control Center needs
the accounts for purposes of setting authorization (user privileges).

1. Log in to Sybase Control Center using an account with administrative privileges. (The
   login account or its group must have sccAdminRole.)
2. Select **Application > Administration > Security.**
3. Check the box labeled **Automatically add SCC login records for authenticated
   logins**.
4. Check the box labeled **Automatically grant sccUserRole to newly created logins**.
5. Click **OK** to close the Security dialog.

Next

There are two next steps:

- If you opted not to automatically create Sybase Control Center login accounts, enter each
  account into Sybase Control Center manually.
- Whether you add accounts automatically or manually, you must grant privileges to any
  login accounts that require more than basic user access. You can grant privileges by
  assigning Sybase Control Center roles directly to the login accounts, or by assigning the
  login accounts to groups and mapping Sybase Control Center roles to the groups. The
  group approach is generally more efficient.

Configuring a Pluggable Authentication Module (PAM) for UNIX

Set up Sybase Control Center to support username and password login using accounts on the UNIX operating system. Optionally, have Sybase Control Center create an account when a UNIX user first logs in to Sybase Control Center.

1. Using a login account with root privileges, configure the pluggable authentication module for your platform:

Platform	Action
Solaris	Append the contents of the `<SCC-install-dir>/utility/su-nos/pam.conf` file (provided with Sybase Control Center) to the `/etc/pam.conf` file on your Solaris platform.
Linux	Copy the `<SCC-install-dir>/utility/linux/sybase-ua` file (provided with Sybase Control Center) to the `/etc/pam.d` directory on your Linux platform.
	Note: The `sybase-ua` file provided with Sybase Control Center is not compatible with the most recent SUSE Linux versions. For SUSE 11 and later, see the example at the end of this topic.

Note: In the table above, the portion of the path that indicates the operating system might differ slightly from what is shown.

2. If the host UNIX system is not using a directory lookup for authentication (yp or NIS, for example) and authentication is carried out against the local `/etc/passwd` file, change the permissions on `/etc/shadow` to provide read access to the login account that executes SCC.

3. (Skip if you configured a PAM before starting Sybase Control Center) Restart Sybase Control Center.

4. (Optional) If you want Sybase Control Center to create an account when a UNIX user logs in to Sybase Control Center, execute these steps. If you choose not to create Sybase Control Center accounts automatically, you must enter them manually. Sybase Control Center needs the accounts for purposes of setting authorization (user privileges).

   a) Log in to Sybase Control Center using an account with administrative privileges (sccAdminRole).
   b) Select **Application > Administration > Security.**
   c) Check the box labeled **Automatically add SCC login records for authenticated logins**.
   d) Click **OK** to close the Security dialog.

Example: PAM for SUSE Linux 11 and later

For SUSE 11 and later, do not use the `sybase-ua` file provided with Sybase Control Center. Instead, in your `/etc/pam.d` directory, create a `sybase-ua` file that contains:

```
# sybase-ua PAM Configuration (SUSE style)
auth       include      common-auth
account    include      common-account
password   include      common-password
session    include      common-session
```

Next

There are two next steps:

- If you opted not to automatically create Sybase Control Center login accounts, enter each account into Sybase Control Center manually.
- Whether you add accounts automatically or manually, you must also grant privileges to the login accounts. You can grant privileges by assigning Sybase Control Center roles directly to the login accounts, or by assigning the login accounts to groups and mapping Sybase Control Center roles to the groups. The group approach is generally more efficient.

See also

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 100
- *Adding a Login Account to the System* on page 110

Configuring an LDAP Authentication Module

Configure an LDAP authentication module for Sybase Control Center by editing the security properties file to point to the correct LDAP server.

1. Open the `<SCC-install-dir>\conf\csi.properties` file.
2. Uncomment the LDAP module in the properties file by removing the # symbol at the beginning of each line (or, if necessary, add an LDAP module to the file). The sample module below specifies the LDAP server that will provide user authentication.

   The sample module shows the properties used for an OpenDS LDAP server. See the example at the end for values that work for ActiveDirectory. Configuration properties you can use in the LDAP module are described in a subtopic.

   Each line of the LDAP server module of the properties file must begin with "CSI.loginModule." followed by a module number. (The module number in this sample is 7.) The module number you assign must be unique in the properties file, and you must use the same module number in every line of the module.

```
CSI.loginModule.
7.options.AuthenticationSearchBase=ou=users,dc=example,dc=com
CSI.loginModule.7.options.BindDN=cn=Directory Manager
CSI.loginModule.7.options.BindPassword=secret
CSI.loginModule.7.options.DefaultSearchBase=dc=example,dc=com
CSI.loginModule.7.options.ProviderURL=ldap://localhost:10389
CSI.loginModule.
7.options.RoleSearchBase=ou=groups,dc=example,dc=com
CSI.loginModule.7.options.ServerType=openldap
CSI.loginModule.7.options.moduleName=LDAP Login Module
CSI.loginModule.7.controlFlag=sufficient
```

```
CSI.loginModule.
7.provider=com.sybase.ua.services.security.ldap.LDAPLoginModule
```

Note: Change the values of bolded lines only.

3. Save the file.

4. If your LDAP server's SSL certificate is signed by a nonstandard certificate authority (for example, if it is a self-signed certificate), use the **keytool** utility to configure your JVM or JDK to trust the certificate. Execute a command similar to this:

```
keytool -import -keystore <sybase-dir>/shared/JRE-6_0_6/bin/
keytool/lib/security/cacerts -file
<your cert file and path> -alias ldapcert -storepass changeit
```

LDAP configuration values for ActiveDirectory

For an ActiveDirectory server, use these values for configuration properties in your LDAP login module:

```
ServerType: msad2K
DefaultSearchBase: dc=<domainname>,dc=<tld> or o=<company
name>,c=<country code>
                E.g. dc=sybase,dc=com or o=Sybase,c=us
ProviderUrl: ldaps://<hostname>:<port>
                E.g.:  ldaps://myserver:636
AuthenticationFilter: (&(userPrincipalName={uid})
(objectclass=user))
BindDN: <User with read capability for all users>
BindPassword: <Password for BindDN user>
RoleFilter: (|(objectclass=groupofnames) (objectclass=group))
controlFlag: sufficient
```

Next
There are two additional steps:

- Set up roles and passwords for LDAP
- Map Sybase Control Center role to LDAP groups

See also

Setting Up Roles and Passwords
Set the initial user roles and passwords required for Sybase Control Center to authenticate through an LDAP server.

Prerequisites
Configure an LDAP authentication module.

Task

1. Open the `<SCC-install-dir>\conf\roles-map.xml` file and add an LDAP login module.

   Insert an LDAP login module similar to this at the end of the security-modules portion of the file, just before `</security-modules>`:

   ```
       <module name="LDAP Login Module">
         <role-mapping modRole="sybase"
   uafRole="uaAnonymous,uaPluginAdmin,sccUserRole" />
         <role-mapping modRole="administrators"
   uafRole="uaAnonymous,sccAdminRole" />
       </module>
   ```

2. Ensure that the roles defined in the LDAP repository match the roles defined in `roles-map.xml`.

3. In the `<SCC-install-dir>\conf\csi.properties` file, set the BindPassword and ProviderURL properties with values used in your deployment.

   Sybase recommends that you encrypt sensitive values before saving them in `csi.properties`.

Next
Map Sybase Control Center roles to LDAP groups.

See also
*   *LDAP Configuration Properties* on page 48

LDAP Configuration Properties

Use these properties in your `csi.properties` file to control your LDAP service.

Property	Default Value	Description
ServerType	None	Optional. The type of LDAP server you are con-necting to: • `sunone5` -- SunOne 5.x OR iPlanet 5.x • `msad2k` -- Microsoft ActiveDirectory, Win-dows 2000 • `nsds4` -- Netscape Directory Server 4.x • `openldap` -- OpenLDAP Directory Server 2.x The value you choose establishes default values for these other authentication properties: • RoleFilter • UserRoleMembership • RoleMemberAttributes • AuthenticationFilter • DigestMD5Authentication • UseUserAccountControl
ProviderURL	`ldap://local-host:389`	The URL used to connect to the LDAP server. Use the default value if the server is: • Located on the same machine as your product that is enabled with the common security in-frastructure. • Configured to use the default port (389). Otherwise, use this syntax for setting the value: `ldap://<hostname>:<port>`

Property	Default Value	Description
DefaultSearchBase	None	The LDAP search base that is used if no other search base is specified for authentication, roles, attribution and self registration: 1. `dc=<domainname>,dc=<tld>` For example, a machine in sybase.com domain would have a search base of dc=sybase,dc=com. 2. `o=<company name>,c=<country code>` For example, this might be o=Sybase,c=us for a machine within the Sybase organization.
SecurityProtocol	None	The protocol to be used when connecting to the LDAP server. To use an encrypted protocol, use "ssl" instead "ldaps" in the url. **Note:** ActiveDirectory requires the SSL protocol when setting the value for the password attribute. This occurs when creating a user or updating the password of an existing user.
AuthenticationMethod	simple	The authentication method to use for all authentication requests into LDAP. Legal values are generally the same as those of the java.naming.security.authentication JNDI property. Choose one of: • simple — For clear-text password authentication. • DIGEST-MD5 — For more secure hashed password authentication. This method requires that the server use plain text password storage and only works with JRE 1.4 or later.

Property	Default Value	Description
AuthenticationFilter	For most LDAP servers: `(&(uid={uid}) (object- class=person))` or For Active Directory email lookups: `(&(userPrinci- palName={uid}) (object- class=user)) [ActiveDirec- tory]` For Active Directory Windows username lookups: `(&(sAMAc- count- Name={uid}) (object- class=user))`	The filter to use when looking up the user. When performing a username based lookup, this filter is used to determine the LDAP entry that matches the supplied username. The string "{uid}" in the filter is replaced with the supplied username.
AuthenticationScope	onelevel	The authentication search scope. The supported values for this are: • `onelevel` • `subtree` If you do not specify a value or if you specify an invalid value, the default value is used.
AuthenticationSearchBase	none	The search base used to authenticate users. If this value is not specified, the LDAP DefaultSearch-Base is used.

Property	Default Value	Description
BindDN	none	The user DN to bind against when building the initial LDAP connection.
		In many cases, this user may need read permissions on all user records. If you do not set a value, anonymous binding is used. Anonymous binding works on most servers without additional configuration.
		However, the LDAP attributer may also use this DN to create the users in the LDAP server. When the self-registration feature is used, this user may also need the requisite permissions to create a user record. This behavior can occur if you do not set useUserCredentialsToBind to true. In this case, the LDAP attributer uses this DN to update the user attributes.
BindPassword	none	BindPassword is the password for BindDN, which is used to authenticate any user. BindDN and BindPassword are used to separate the LDAP connection into units.
		The AuthenticationMethod property determines the bind method used for this initial connection.

		If you use an encrypted the password using the CSI encryption utility, append .e to the property name. For example:
		
		`CSI.loginModule.7.options.`
		`BindPassword.e=1-AAAAEgQQOLL+LpX`
		`JO8fO9T4SrQYRC9lRT1w5ePfdczQTDs`
		`P8iACk9mDAbm3F3p5a1wXWKK8+NdJuk`
		`nc7w2nw5aGJlyG3xQ==`

RoleSearchBase	none	The search base used to retrieve lists of roles. If this value is not specified, the LDAP Default-SearchBase is used.

Property	Default Value	Description
RoleFilter	For SunONE/iPlanet: `(&(object-class=ldapsu-bentry)(ob-jectclass=nsro-ledefinition))` For Netscape Directory Server: `(object-class=groupof-names)(object-class=groupofu-niquenames))` For ActiveDirectory: `(object-class=groupof-names)(object-class=group))`	The role search filter. This filter should, when combined with the role search base and role scope, return a complete list of roles within the LDAP server. There are several default values depending on the chosen server type. If the server type is not chosen or this property is not initialized, no roles are available.
RoleMemberAttributes	For Netscape Directory Server: member,unique-member	The role's member attributes defines a comma-delimited list of attributes that roles may have that define a list of DN's of people who are in the role. These values are cross referenced with the active user to determine the user's role list. One example of the use of this property is when using LDAP groups as placeholders for roles. This property only has a default value when the Netscape server type is chosen.
RoleNameAttribute	cn	The attribute for retrieved roles that is the common name of the role. If this value is "dn" it is interpreted specially as the entire dn of the role as the role name.
RoleScope	onelevel	The role search scope. The supported values for this are: • `onelevel` • `subtree` If you do not specify a value or if you specify an invalid value, the default value is used.

Property	Default Value	Description
UserRoleMembershipAttributes	For iPlanet/SunONE: nsRoleDN For ActiveDirectory: memberOf For all others: none	The user's role membership attributes property is used to define an attribute that a user has that contains the DN's of all of the roles as user is a member of. These comma-delimited values are then cross-referenced with the roles retrieved in the role search base and search filter to come up with a list of user's roles.
UserFreeformRoleMembershipAttributes	None	The "freeform" role membership attribute list. Users who have attributes in this comma-delimited list are automatically granted access to roles whose names are equal to the attribute value. For example, if the value of this property is "department" and user's LDAP record has the following values for the department attribute, { "sales", "consulting" }, then the user will be granted roles whose names are "sales" and "consulting".
Referral	ignore	The behavior when a referral is encountered. The valid values are those dictated by LdapContext, for example, "follow", "ignore", "throw".
DigestMD5AuthenticationFormat	DN For OpenLDAP: Username	The DIGEST-MD5 bind authentication identity format.
UseUserAccountControlAttribute	For most LDAP servers: false For ActiveDirectory: true	The UserAccountControl attribute to be used for detecting disabled user accounts, account expirations, password expirations and so on. ActiveDirectory also uses this attribute to store the above information.
controlFlag	optional	Indicates whether authentication with this login module is sufficient to allow the user to log in, or whether the user must also be authenticated with another login module. Rarely set to anything other than "sufficient" for any login module. **Note:** controlFlag is a generic login module option rather than an LDAP configuration property.

See also
* *Setting Up Roles and Passwords* on page 46

Mapping Sybase Control Center Roles to LDAP or OS Groups

To grant Sybase Control Center privileges to users who are authenticated through LDAP or the operating system, associate roles used in Sybase Control Center with groups in LDAP or the operating system.

You can configure Sybase Control Center to enable users to authenticate through their local operating system or through an LDAP server. To make this type of authentication work, SCC roles must be mapped to groups that exist in the system providing authentication (LDAP or the operating system) or in the login module.

By default, SCC assumes there is a "sybase" group in the authenticating system and maps the LDAP or OS "sybase" group to SCC roles to provide basic privileges. The table lists additional default mappings of LDAP and OS groups to SCC roles.

Login Module	OS Group	Sybase Control Center Roles
UNIX Proxy	root	uaAnonymous, uaAgentAdmin, uaOSAdmin
	sybase	uaAnonymous, uaPluginAdmin, sccUserRole
	user	uaAnonymous, uaUser
	guest	uaAnonymous, uaGuest
NT Proxy	Administrators	uaAnonymous, uaAgentAdmin, uaOSAdmin
	sybase	uaAnonymous, uaPluginAdmin, sccUserRole
	Users	uaAnonymous, uaUser
	Guests	uaAnonymous, uaGuest
LDAP	sybase	uaAnonymous, uaPluginAdmin, sccUserRole

There are two ways to accomplish the mapping:

- (Recommended) Add a "sybase" group to the operating system or LDAP server Sybase Control Center is using to authenticate users, and add all users who need to access Sybase Control Center to the "sybase" group.
- Configure Sybase Control Center to use an existing group in LDAP or the operating system by editing the `roles-map.xml` file. This option is described here.

1. If Sybase Control Center is running, shut it down.
2. In a text editor, open:

   `<SCC-install-directory>/conf/roles-map.xml`
3. Locate the appropriate login module: UNIX or NT (for Windows).

4. Copy the line that maps the "sybase" group and paste it into the module just above the original sybase line.

5. Change "sybase" to the name of the group in your operating system to which Sybase Control Center users belong.

   For example, if the group is SCCusers, the new line should look like this:

   ```
   <role-mapping modRole="SCCusers"
   uafRole="uaAnonymous,uaPluginAdmin,sccUserRole" />
   ```

6. Save the file and exit.

7. Start Sybase Control Center.

See also
- *Configuring an LDAP Authentication Module* on page 45

Encrypting a Password
Use the **passencrypt** utility to encrypt passwords and other values that must be kept secure while stored in text files.

You can safely store an encrypted password in a properties file. Enter the password in clear text (unencrypted) when you execute **passencrypt** and when you use the password to log in.

passencrypt, which is located in the Sybase Control Center bin directory, uses the DES encryption algorithm.

1. Open a command window and change to the bin directory:

   Windows: `cd <SCC-install-directory>\bin`

   UNIX: `cd <SCC-install-directory>/bin`

2. To encrypt a password, enter **passencrypt**. Enter your new password at the resulting prompt.
   passencrypt encrypts the password you enter (which does not appear on the screen) and displays the password in encrypted form.

3. Copy the encrypted password.

4. Paste the encrypted password where needed.

Configuring Ports
(Optional) Use the **scc --port** command to assign Sybase Control Center services to new ports.

Prerequisites
Check for port conflicts between Sybase Control Center and other software running on the same host.

Task

Sybase Control Center cannot function properly if other services use its ports. If you discover a conflict with any port listed in the right column below, you can either reconfigure the other service's port or reconfigure Sybase Control Center as described here.

Port Name	Description	Service Names	Property Names	Default Port
db	Database port Present on SCC server	SccSADataserver Messaging Alert Scheduler	com.sybase.asa.server.port messaging.db.port alert.database.port org.quartz.data-Source.ASA.URL	3638
http	Web HTTP port Present on SCC server	EmbeddedWebContainer	http.port	8282
https	Web HTTPS (secure HTTP) port Present on SCC server	EmbeddedWebContainer	https.port	8283
jiniHttp	JINI HTTP server Present on SCC server and SCC agent	Jini	httpPort	9092
jiniR-mid	JINI remote method invocation daemon Present on SCC server and SCC agent	Jini	rmidPort	9095
msg	Messaging port Present on SCC server	Messaging	messaging.port	2000
rmi	RMI port Present on SCC server and SCC agent	RMI	port	9999

Port Name	Description	Service Names	Property Names	Default Port
tds	Tabular Data Stream™ port (used to communicate with other Sybase products) Present on SCC server and SCC agent	Tds	tdsPort	9998

1. Shut down Sybase Control Center.
2. Execute **scc --info ports** to display a list of Sybase Control Center services, their properties, and their assigned ports.
3. To reassign a port, enter a command in one of these formats:
   ```
   scc --port port-name=port-number
   scc --port service-name:property-name=port-number
   ```
   Use the first, simpler format unless you want to configure the database services to use different ports. (By default, they all use the same port.)
4. Start Sybase Control Center.
5. Execute **scc --info ports** again to confirm that the port has been reassigned.

Examples

Set all four database services (data server, messaging, database alert, and scheduler) to the same port, 3639. (The database is SQL Anywhere, used by the Sybase Control Center internal repository.)
```
scc --port db=3639
```

Set only the database messaging service to port 3639.
```
scc --port Messaging:messaging.db.port=3639
```

Set the HTTP port to 9292.
```
scc --port http=9292
```

Set the Jini RMI daemon to port 9696.
```
scc --port jiniRmid=9696
```

Set the main Sybase Control Center messaging service to port 2001.
```
scc --port msg=2001
```

Set the RMI port to 9991.
```
scc --port rmi=9991
```

Set the Tabular Data Stream port to 9997.

```
scc --port tds=9997
```

Note: scc commands that include a port-setting option (**-p** or **--port**) do not start Sybase Control Center. To start SCC, execute a separate **scc** command.

Configuring the E-mail Server

(Optional) Specify the e-mail server for Sybase Control Center to use to send e-mail alert notifications.

Prerequisites

Launch Sybase Control Center and log in using an account with administrative privileges. (The login account or its group must have sccAdminRole.)

Task

1. From the menu bar, select **Application > Administration.**
2. Select **General Settings**.
3. Click the **E-mail** tab.
4. Enter the name of the e-mail server through which Sybase Control Center will send alert notifications.
5. Change the default e-mail server port only in consultation with your e-mail administrator.
6. (Optional) Click **Customize e-mail settings** to display options for setting the domain name and e-mail sender for alert e-mail notifications.
7. (Optional) Enter your domain name (for example, mycompany.com).

   Most e-mail servers do not require SCC to provide an explicit domain name. Try providing a domain name here if your first attempt to configure e-mail alerts fails.
8. (Optional) Change the default e-mail sender name.

   This name appears in the "From" field of SCC e-mail alert messages. Do not use spaces; use hyphens or underscore characters instead.

   Tip: If you have multiple SCC servers, configure their sender names so you can tell which SCC an alert is coming from. For example, SybaseControlCenter_Boston or SCC_test11.
9. (Optional) If you entered anything in the **E-mail Domain name** or **E-mail sender name** fields, click **Apply** to make the test e-mail option reappear.
10. (Optional) To dispatch a test message, enter an e-mail address in the **Test e-mail address** field and click **Send**.
    If the test e-mail is received, you have properly configured the server for e-mail alert notifications.
11. Click **OK** (to apply the change and close the properties dialog) or **Apply** (to apply the change and leave the dialog open).

Next

(Optional) Configure automatic logout.

See also

• *Setting Up Security* on page 41

Configuring the Automatic Logout Timer

(Optional) Set Sybase Control Center to end login sessions when users are inactive for too long.

Prerequisites

Launch Sybase Control Center and log in using an account with administrative privileges. (The login account or its group must have sccAdminRole.)

Task

1. From the menu bar, select **Application > Administration.**
2. Select **General Settings**.
3. Click the **Auto-Logout** tab.
4. Enter the number of minutes after which an idle user will be automatically logged out.

   Enter 0 or leave the box empty to disable automatic logout.
5. Click **OK** (to apply the change and close the properties dialog) or **Apply** (to apply the change and leave the dialog open).

User Authorization

The authorization mechanism in Sybase Control Center employs login accounts and task-based roles.

Access to Sybase Control Center is controlled by login accounts. You grant permissions to a login account by assigning predefined roles that control tasks the user can perform in Sybase Control Center, such as administration and monitoring of particular types of Sybase servers. The roles can be assigned directly to login accounts or to groups; a login account inherits the roles of any group to which it belongs. Component product modules assign some roles automatically.

Sybase Control Center classifies roles as follows:

• System roles – define how a user can interact with Sybase Control Center.
• Product roles – define how a user can interact with a particular managed resource in Sybase Control Center, for example the Replication Server named RepBoston01.

Note: The tools described here are for managing SCC-enabled login accounts; you cannot use them to manage accounts and groups that are native to your managed resource.

See also
*   *Configure* on page 62

Assigning a Role to a Login or a Group

Use the security configuration options to add one or more roles to a Sybase Control Center login account or to a group. Roles enable users to perform tasks such as monitoring servers or administering Sybase Control Center.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task. To assign a monitoring role for a server, first register the server.

Task

Assign the sccAdminRole to any login account that will perform administrative tasks in Sybase Control Center.

1. From the application menu bar, select **Application > Administration.**
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Click **Logins** or **Groups**.
4. In the table, select the login account or group to which you want to assign a role.
5. Click the **Roles** tab.
6. In the **Available roles for resource** list, select the role, then click **Add**. For example, to grant administrative privileges, add the SCC Service:sccAdminRole. To grant monitoring privileges, add the MonitorRole for the desired server and server type.

    Note: Sybase Control Center product modules assign certain roles automatically, so you might not need to add a MonitorRole.

    If a role appears in the **Has following roles** list, this account or group has already been configured with that role.
7. Click **OK**.

See also
*   *Adding a Group* on page 60
*   *Adding a Login Account to a Group* on page 61
*   *Logins, Roles, and Groups* on page 62

Adding a Group

Use the security configuration options to create a new group.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

Groups can make roles easier to manage. Rather than assigning roles to individual users, assign roles to groups and add users to the groups or remove them as needed.

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Select **Groups**.
4. Click **Create Group**.
5. Enter a group name and a description.
6. Click **Finish**.

See also
- *Assigning a Role to a Login or a Group* on page 60
- *Adding a Login Account to a Group* on page 61
- *Logins, Roles, and Groups* on page 62

Adding a Login Account to a Group
Use the security configuration options to add one or more login accounts to a group.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Click **Groups**.
4. Select the group to which you want to assign an account.
5. Click the **Membership** tab.
6. Select the account, then click **Add**.
7. Click **OK**.

See also
- *Assigning a Role to a Login or a Group* on page 60
- *Adding a Group* on page 60
- *Logins, Roles, and Groups* on page 62

Logins, Roles, and Groups

Sybase Control Center includes predefined login accounts and roles.

In Sybase Control Center, a login account identifies a user who can connect to the application. An account may have roles that specify the tasks the user is allowed to perform.

Sybase Control Center is designed to delegate user authentication to the operating system or to an LDAP directory service. Delegation requires some configuration, however, so Sybase Control Center comes with two predefined login accounts. Sybase recommends using the predefined accounts only for installing and setting up Sybase Control Center. These accounts are not intended for use in a production environment.

Table 10. Predefined accounts

Login name	Description
sccadmin	Can use all the administration features in Sybase Control Center
sccuser	Test account with no special privileges

A role is a predefined profile that can be assigned to a login account or a group. Roles control the access rights for login accounts. Sybase Contol Center comes with predefined roles that are intended for use in production environments.

Table 11. Predefined roles

Role	Description
sccUserRole	Provides nonadministrative access to Sybase Control Center. Required for every user.
sccAdminRole	Provides administrative privileges for managing Sybase Control Center.

Monitoring privileges for SCC product modules are assigned automatically.

A group is made up of one or more login accounts; all the accounts in a group have the roles granted to the group. In Sybase Control Center you can create groups to suit your business requirements.

See also

Configure

Configure Sybase Control Center for Sybase IQ.

1. *Configuring Sybase IQ for Monitoring*

   To enable one or more Sybase IQ users to authenticate a Sybase IQ server with Sybase Control Center, create and populate the SCC_MONITOR group.

2. *Configuring Sybase IQ for Administration*

   To perform administration tasks, you must have the correct authority or group membership, and you may need to register the server's Sybase Control Center agent.

3. *Registering a Sybase IQ Server*

   Make Sybase Control Center aware of a Sybase IQ resource (for example, a server that can be monitored) and its connection information by registering the resource.

4. *Importing Resources for Batch Registration*

   (Optional) Import and register multiple servers from an interfaces or sql.ini file.

5. *Registering and Authenticating a Sybase Control Center Agent*

   Register and authenticate the Sybase Control Center agent for a managed server.

6. *Creating a Perspective*

   Create a perspective in which you can add and manage resources.

7. *Adding a Resource to a Perspective*

   Add one or more resources to the current perspective.

8. *Authenticating a Login Account for a Managed Resource*

   Specify the login account Sybase Control Center will use when it connects to your server or agent to collect monitoring data or manage the resource.

9. *Changing Update Frequency for Statistics and Charts*

   You can control the rate at which data on monitor screens and charts is refreshed, the amount of time covered by charts, and the multiplex nodes included in charts.

10. *Setting Up Statistics Collection*

    Use the Properties view of your managed resource to create a data collection job and add a schedule to the job.

11. *Creating an Alert*

    Use the Add Alert wizard to create an alert instance for your resource.

12. *Optional Configuration Steps*

    Perform additional configuration, including user authorization, alerts, data collection scheduling, backups, and setting purging options for the repository.

See also
- *User Authorization* on page 59
- *Logins, Roles, and Groups* on page 112
- *Setting Up Security* on page 87
- *Assigning a Role to a Login or a Group* on page 106

Deploying an Instance from a Shared Disk Installation

(Optional) Create a Sybase Control Center server or agent from an installation on a shared disk.

Prerequisites

- Install Sybase Control Center on a shared disk.
- Enable shared-disk mode.

Task

1. Log in to the host on which you plan to run the SCC server or agent.

   Note: You can create an instance on one host and run it on another host, but doing so interferes with the predeployment checks run by **sccinstance**. Such a deployment might generate errors (port conflicts, for example). If you are confident that the errors are caused by problems that will not be present on the host where you plan to run the instance, use the **-force** option to create the instance.

2. Change to SCC-3_2/bin.

3. Create the instance as an SCC agent if you plan to run a managed server on this host. Create the instance as an SCC server if you plan to manage other Sybase servers from this host.

   To create an SCC agent called Boston-agent and configure it to run as a Windows service:

   ```
   sccinstance -create -agent -instance Boston-agent -service
   ```

   To create an SCC server called Boston and configure it to run as a Windows service:

   ```
   sccinstance -create -server -instance Boston -service
   ```

4. If other SCC instances will run on this host, change the port assignments for the new instance. Change the instance names and port values in the sample commands to suit your environment, but take care to specify ports that are not in use by another SCC instance or any other application or server.

   This command changes the port assignments for an SCC agent called myagent:
   ```
   sccinstance -refresh -instance myagent -portconfig
   rmi=8888,jiniHttp=9093,jiniRmi=9096,tds=9997
   ```

   This command changes the port assignments for an SCC server called myserver:
   ```
   sccinstance -refresh -server -instance myserver -portconfig
   rmi=8889,db=3640,
   http=7072,https=7073,jiniHttp=9094,jiniRmi=9097,msg=2002,tds=9996
   ```

5. (Optional) List the instances deployed from this installation:

```
sccinstance -list
```

6. (Optional) If you are setting up an instance in UNIX, configure it to run as a service. (See *Starting and Stopping Sybase Control Center in UNIX*).

Next

When you manage and maintain instances, keep in mind that the directory structure for instances is different from that of singleton installations. In file paths in SCC help, replace `SCC-3_2` or `<scc-install-directory>` with `SCC-3_2/instances/ <instance-name>`.

For example, the path to the log directory, `SCC-3_2/log`, becomes this for an instance called kalamazoo:

```
SCC-3_2/instances/kalamazoo/log
```

Enabling and Disabling Shared-Disk Mode

Turn on or turn off shared-disk mode, which allows you to run multiple Sybase Control Center agents and servers from a single installation on a shared disk.

Prerequisites

Install Sybase Control Center on a shared disk. See the *Sybase Control Center Installation Guide*.

Task

Shared-disk mode affects the entire installation; do not enable or disable individual instances.

Disabling shared-disk mode leaves the instances' file systems intact under `<SCC-install-directory>/instances`, but the instances cannot run. If you reenable, the instances are able to run again.

1. Change to `SCC-3_2/bin`.

2. Enable or disable shared disk mode.

To enable shared disk mode:

```
sccinstance -enable
```

To disable shared disk mode:

```
sccinstance -disable
```

See also
- *Shared-Disk Mode* on page 66
- *sccinstance Command* on page 66

Shared-Disk Mode

Shared-disk mode lets you run multiple Sybase Control Center instances—SCC servers, SCC agents, or a mixture of the two—from a single installation of the product.

The shared-disk capability enables SCC servers or agents on the installation host or on remote hosts to access and execute from the same installation. This feature is especially useful if you plan to use SCC to manage Adaptive Server clusters or Sybase IQ multiplexes.

After installing SCC on a shared disk, use the **sccinstance** command to enable shared-disk mode and deploy instances. **sccinstance** copies the files needed for the instance into a new directory structure. The path takes the form `<SCC-install-directory>/instances/<instance-name>` (for example, `SCC-3_2/instances/SCCserver-1`).

You can specify a name for each instance. If you do not supply a name, the instance name defaults to the host name.

An instance runs on the host on which you start it. When shared-disk mode is enabled, SCC servers and agents run out of the `SCC-3_2/instances` subdirectories, not from the base file system.

In shared-disk mode, changes made to configuration files in the base file system (everything under `SCC-3_2` except the `SCC-3_2/instances` branch) are copied to any instance deployed thereafter. Previously deployed instances are not affected.

Use **sccinstance** to deploy, remove, refresh, or convert an instance; to configure an instance's ports; and to configure a Windows instance to run as a service. Perform other tasks, including configuring a UNIX instance to run as a service, and all other configuration, using the tools and procedures documented for all installations. Use tools provided by the UI wherever possible. When you must edit a file to change the configuration of an instance (for role mapping, for example), edit the copy of the file stored under `<SCC-install-directory>/instances/<instance-name>`.

See also

- *Enabling and Disabling Shared-Disk Mode* on page 65
- *sccinstance Command* on page 66

sccinstance Command

Use **sccinstance.bat** (Windows) or **sccinstance** (UNIX) to deploy an instance of Sybase Control Center from a shared-disk installation or to manage existing instances.

You can run multiple instances of Sybase Control Center, including SCC servers, SCC agents, or a mixture of the two, from a single installation on a shared disk.

Syntax

```
sccinstance[.bat]
[-agent]
[-c | -create]
[-d | -debug]
[-disable]
[-enable]
[-f | -force]
[-h | -help]
[-i | -instance [instance-name]]
[-l | -list]
[-plugins {plugin-ID,plugin-ID,...}]
[-portconfig {port-name=port-number,port-name=port-number, ...}]
[-refresh]
[-r | -remove]
[-s | -server]
[-service]
[-silent]
```

Parameters

- **-agent** – use with **-create** or **-refresh** to create or refresh an SCC agent. In a **-create** or **-refresh** command, **-agent** is the default, so you can omit it.
- **-create** – deploy a new instance. Use alone or with **-agent** to create an agent instance, or with **-server** to create a server instance.
- **-d | debug** – display debugging messages with the output of this command.
- **-disable** – turn off shared-disk mode for this installation. Generates an error if any instance is running.
- **-enable** – turn on shared-disk mode for this installation. Shared-disk mode is required if you intend to run more than one server or agent from a single installation of SCC.
- **-f | -force** – execute **sccinstance** even if there are potential conflicts (such as port clashes or a running SCC process).
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-instance** – specify an instance. Use with **-create**, **-remove**, or **-refresh**, or use alone to display the instance's status. You can omit **-instance** when you are addressing the only SCC instance or the only instance of the specified type (server or agent) on the current host.
- **-l | -list** – display a list of all instances deployed from this SCC installation.
- **-plugins {*plugin-ID,plugin-ID,...*}** – specify one or more product module plug-ins for this instance. An alternative to **-agent** and **-server**, **-plugins** is primarily for use by the SCC installation program. Use with **-create** or **-refresh**. Use commas to separate plug-in names.
- **-portconfig {*port-name=port-number*, *port-name=port-number*, *...*}** – assign ports to services for this instance. Use only with **-create** or **-refresh**. For the *port-name* value, use a port name from the table below. If you plan to run more than one SCC instance on a host machine, you must reassign all the ports for every instance after the first.

  Port information:

Port Name	Description	Service Names	Property Names	Default Port
db	Database port Present on SCC server	SccSADataserver Messaging Alert Scheduler	com.sybase.asa.server.port messaging.db.port alert.database.port org.quartz.data-Source.ASA.URL	3638
http	Web HTTP port Present on SCC server	EmbeddedWebContainer	http.port	8282
https	Web HTTPS (secure HTTP) port Present on SCC server	EmbeddedWebContainer	https.port	8283
jiniHttp	JINI HTTP server Present on SCC server and SCC agent	Jini	httpPort	9092
jiniR-mid	JINI remote method invocation daemon Present on SCC server and SCC agent	Jini	rmidPort	9095
msg	Messaging port Present on SCC server	Messaging	messaging.port	2000
rmi	RMI port Present on SCC server and SCC agent	RMI	port	9999
tds	Tabular Data Stream™ port (used to communicate with other Sybase products) Present on SCC server and SCC agent	Tds	tdsPort	9998

- **-refresh** – recopy all the files that make up this instance (Windows) or all this instance's services and plug-ins (UNIX). Refreshing preserves any service or plug-in configuration in the deployed instance.

  You can also use **-refresh** to convert a server to an agent or an agent to a server (see the examples). Files are removed or added to change the function of the instance. Use alone or

with **-agent** to refresh an agent instance, or with **-server** to refresh a server instance. Generates an error if the instance is running.

- **-r | -remove** – delete an instance. Use alone or with **-instance**. Generates an error if the instance is running. You cannot restore a removed instance.
- **-s | -server** – use with **-create** or **-refresh** to create or refresh an SCC server, including any product modules available.
- **-service** – use with **-create** or **-remove** to create or remove a Windows service for this instance. You must be logged in to Windows as an administrator to use this option.
- **-silent** – suppress the output of **sccinstance**.

Examples

- **Deploy an SCC server instance** – enables shared-disk mode, deploys a server called Boston with a Windows service, and starts the Windows service:

```
sccinstance -enable
sccinstance -create -server -instance Boston -service
net start "Sybase Control Center 3.2.3 (Boston)"
```

Note: To create the service, you must log in to Windows as an administrator.

- **Deploy an SCC agent instance** – deploys an SCC agent on this host and configures a Windows service for it. The **-agent** option, because it is the default, is not required—the command does exactly the same thing without it.

```
sccinstance -create -agent -service
```

or

```
sccinstance -create -service
```

- **Deploy a server instance and reassign ports** – deploys the server on this host and configures nondefault RMI, HTTP, and HTTPS ports.

```
sccinstance -create -server -portconfig
rmi=8888,http=7070,https=7071
```

- **Refresh a server instance or convert an agent to a server** – refreshes the server on this host. If the instance on this host is an SCC agent, refreshing it as an SCC server converts it into a server.

```
sccinstance -refresh -server
```

- **Refresh an agent instance or convert a server to an agent** – refreshes the instance named kalamazoo. If kalamazoo is a server, refreshing it as an SCC agent converts it into an agent.

```
sccinstance -refresh -agent -instance kalamazoo
```

- **Remove a server instance** – removes the instance named porcupine if it is not running:

```
sccinstance -remove -instance porcupine
```

- **Display status** – displays the status of the instance on this host:

```
sccinstance
```

- **List all instances** – displays a list of all SCC server and agent instances deployed from this SCC installation:

```
sccinstance -list
```

- **Scenario: Remove an instance by force** – suppose you have inadvertently deployed two SCC agent instances on the same host:

```
$ sccinstance -list
2 SCC instances deployed:
SCC instance node1 deployed in agent mode for host node1 RMI port
9999
SCC instance node2 deployed in agent mode for host node2 RMI port
9999
```

Both instances use the same RMI port. You must either reassign ports for one instance or remove it. But you get an error if you try remove an instance when another instance is running on the same host:

```
$ sccinstance -instance node2 -remove
[ERROR] Command execution failed.
[ERROR] SCC instance node2 could not be removed because it is
running. Shut
down the SCC before removing the instance.
```

Use the **-force** option to override the error and force the removal of the second agent instance:

```
$ sccinstance -instance node2 -remove -force
Removing SCC instance node2 ...
SCC instance node2 was successfully removed.
```

Permissions

sccinstance permission defaults to all users, except as noted for certain parameters.

See also
- *Enabling and Disabling Shared-Disk Mode* on page 65
- *Shared-Disk Mode* on page 66

Launching Sybase Control Center

Use the **scc** command to start Sybase Control Center.

Prerequisites

Install Adobe Flash Player in the browser you will use for Sybase Control Center.

Task

1. Start Sybase Control Center.
   - Windows – navigate to `<install_location>\SCC-3_2\bin` and double-click **scc.bat**.
   - UNIX – execute **scc.sh**.

   Messages on the progress of the launch appear in a command window. When Sybase Control Center is running, the command window becomes the Sybase Control Center console; you can issue commands to get status information on SCC and its ports, plug-ins, and services.

2. Open a Web browser and enter `https://<hostname>:8283/scc`.

See also
- *Sybase Control Center Console* on page 200

Registering the ODBC Driver in Windows

In Windows, run **scc.bat** with administrative privileges to register the ODBC driver.

When Sybase Control Center starts for the first time on a Windows machine, it registers its ODBC driver. Because the automatic registration of the ODBC driver edits the registry settings, you must execute **scc.bat** using elevated administrative privileges. If you launch for the first time without adequate privileges, Sybase Control Center generates an error and fails to start.

In Windows Vista, Windows 2008, and Windows 7, you must use the **Run as administrator** setting to launch Sybase Control Center even if you already have administrative privileges. This process is described below.

In other versions of Windows, you must be logged in as an administrator to start Sybase Control Center for the first time. You need not follow the steps below.

1. In Windows Vista, Windows 2008, or Windows 7, open the Command Prompt window with administrative privileges:
   - Select **Start > All Programs > Accessories.** Right-click **Command Prompt** and select **Run as administrator**.
   - Alternatively, enter **cmd** in the Start Menu search box and press **Shift**+**Ctrl**+**Enter**.
2. Run **scc.bat**.

See also
- *Starting and Stopping Sybase Control Center in Windows* on page 72
- *Starting and Stopping Sybase Control Center in UNIX* on page 74
- *Configuring Memory Usage* on page 78
- *scc Command* on page 81

Starting and Stopping Sybase Control Center in Windows

There are several ways to start and stop Sybase Control Center or the SCC agent. You can start manually, which is useful for testing and troubleshooting, or set the service to start automatically and to restart in case of failure.

This topic applies to both Sybase Control Center (the server) and the Sybase Control Center agent that runs on each product server managed by SCC. It applies to both singleton installations and instances of SCC agents and servers running from a shared disk.

If you run Sybase Control Center or the SCC agent manually, you must issue a command every time you start or shut down. If you run as a service (which is recommended), you can configure the service to start and restart automatically. These are the options:

- Use the **scc.bat** command to start Sybase Control Center or the SCC agent manually. The command gives you access to the Sybase Control Center console, which you can use to shut down and to display information about services, ports, system properties, and environment variables. You can also use **scc.bat** to change the logging level for troubleshooting purposes. Using **scc.bat** prevents you from taking advantage of the automatic start and restart features available to services.
- Use the Services list under the Windows Control Panel to start, stop, and configure the Sybase Control Center service for an SCC server or agent.
- Use the **net start** and **net stop** commands. This is another way to run Sybase Control Center or the SCC agent as a service.

  Note: To start an SCC agent or server as a service:
  - In a singleton installation, you must have selected **Yes** in the installer to install the agent or server as a service.
  - In a shared disk installation, the agent or server must have been deployed using the **-service** option of the **sccinstance** command.

In a singleton installation, the installer lets you start Sybase Control Center or the SCC agent as a service and configures the service to restart automatically. Before starting, check the Windows Services list for a Sybase Control Center service.

Here are the steps for each starting and stopping option:

- **Start Sybase Control Center or the SCC agent:**
  a) (Skip this step for the SCC agent.) If you are starting Sybase Control Center for the first time in Windows Vista, Windows 2008, or Windows 7, set the **Run as Administrator** option on the command prompt so that Sybase Control Center can register its ODBC driver. (This is necessary even if you are logged in as an administrator.)
  b) Enter the **scc** command.

     For a singleton installation:
     ```
     %SYBASE%\SCC-3_2\bin\scc.bat
     ```

For an instance:

```
%SYBASE%\SCC-3_2\bin\scc.bat -instance <instance-name>
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

• **Stop Sybase Control Center or the SCC agent:**

a) Enter the **scc --stop** command.

For a singleton installation:

```
%SYBASE%\SCC-3_2\bin\scc.bat --stop
```

For an instance:

```
%SYBASE%\SCC-3_2\bin\scc.bat --stop -instance <instance-
name>
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

Note: You can also enter **shutdown** at the `scc-console>` prompt.

• **Start or stop from the Windows Control Panel; configure automatic start and restart:**

a) Open the Windows Control Panel.

b) Select **Administrative Tools > Services.**

c) Locate "Sybase Control Center" in the Services list. It may be followed by a release number; if the service is for an instance, it is also followed by the instance name. Service names do not distinguish between agents and servers. If the service is running, the Status column displays "Started."

d) To start or stop the service, right-click the **Sybase Control Center** entry in the Services list and choose **Start** or **Stop**.

e) To configure automatic starting, double-click the service.

f) To set the service to automatically start when the machine starts, change the **Startup type** to Automatic.

g) To restart the service in case of failure, choose the **Recovery** tab and change the First, Second, and Subsequent failures to Restart Service.

h) Click **Apply** to save the modifications and close the dialog.

• **Start or stop the Sybase Control Center service (controlling either Sybase Control Center or the SCC agent) from the Windows command line:**

a) To start the service, enter the **net start** command.

For a singleton installation:

```
net start "sybase control center 3.2.4"

The Sybase Control Center 3.2.4 service is starting......
```

```
The Sybase Control Center 3.2.4 service was started
successfully.
```

For an instance, include the instance name in parentheses:

```
net start "sybase control center 3.2.4 (Boston-1)"

The Sybase Control Center 3.2.4 (Boston-1) service is
starting......
The Sybase Control Center 3.2.4 (Boston-1) service was
started successfully.
```

b) To stop the service, enter the **net stop** command.

For a singleton installation:

```
net stop "sybase control center 3.2.4"

The Sybase Control Center 3.2.4 service is stopping.....
The Sybase Control Center 3.2.4 service was stopped
successfully.
```

For an instance, include the instance name in parentheses:

```
net stop "sybase control center 3.2.4 (Boston-1)"

The Sybase Control Center 3.2.4 (Boston-1) service is
stopping.....
The Sybase Control Center 3.2.4 (Boston-1) service was
stopped successfully.
```

See also
- *Registering the ODBC Driver in Windows* on page 71
- *Starting and Stopping Sybase Control Center in UNIX* on page 74
- *Configuring Memory Usage* on page 78
- *scc Command* on page 81

Starting and Stopping Sybase Control Center in UNIX

You can start Sybase Control Center or the SCC agent manually, which is useful for testing and troubleshooting, or you can set up a service to start automatically and to restart in case of failure.

This topic applies to both Sybase Control Center (the server) and the Sybase Control Center agent that runs on each product server managed by SCC. It applies to both singleton installations and instances of SCC agents and servers running from a shared disk.

If you start Sybase Control Center or the SCC agent manually, you must issue a command every time you start or shut down. If you run as a service (which is recommended), you can configure the service to start and restart automatically. These are the options:

- Use the **scc.sh** script to start Sybase Control Center or the SCC agent manually. You can either:
  - Run **scc.sh** in the foreground to get access to the Sybase Control Center console, which you can use to shut down and to display information about services, ports, system properties, and environment variables.
  - Run **scc.sh** in the background to suppress the console.

  You can use **scc.sh** to run Sybase Control Center at a nondefault logging level for troubleshooting. When you start manually with **scc.sh**, you cannot take advantage of the automatic start and restart features available to services.
- Use the **sccd** script to configure a service that starts Sybase Control Center or the SCC agent automatically.

Here are the steps for each starting and stopping option:

- **Before you start Sybase Control Center or the SCC agent for the first time, set environment variables.** Do this only once.
  a) Change to the Sybase directory (the parent of the Sybase Control Center installation directory).
  b) Execute one of the following to set environment variables.

  Bourne shell:
  ```
  . SYBASE.sh
  ```

  C shell:
  ```
  source SYBASE.csh
  ```
- **Run Sybase Control Center or the SCC agent in the foreground.**

  Running in the foreground is a method of manually starting; you must issue comands to stop and restart Sybase Control Center or the SCC agent.

  a) To start Sybase Control Center or the SCC agent and drop into the console when the start-up sequence is finished, enter the **scc** command.

  For a singleton installation:
  ```
  $SYBASE/SCC-3_2/bin/scc.sh
  ```

  For an instance:
  ```
  $SYBASE/SCC-3_2/bin/scc.sh -instance <instance-name>
  ```

  You can omit the **-instance** option if the instance's name is the same as its host name (the default).
- **Run Sybase Control Center or the SCC agent in the background.**

  You can use **nohup**, **&**, and **>** to run Sybase Control Center or the SCC agent in the background, redirect output and system error to a file, and suppress the SCC console. Running in the background is a method of manually starting; you must issue comands to stop and restart Sybase Control Center or the SCC agent.

a) Execute a command similar to the sample below that matches your shell. Both sample commands direct output to the file scc-console.out. If the output file already exists, you might need to use additional shell operators to append to or truncate the file.

Bourne shell (sh) or Bash

For a singleton installation:

```
nohup ./scc.sh 2>&1 > scc-console.out &
```

For an instance:

```
nohup ./scc.sh -instance <instance-name> 2>&1 > scc-console-
your-instance.out &
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

C shell

For a singleton installation:

```
nohup ./scc.sh >& scc-console.out &
```

For an instance:

```
nohup ./scc.sh -instance <instance-name> >& scc-console.out &
```

You can omit the **-instance** option if the instance's name is the same as its host name (the default).

- **Shut down Sybase Control Center or the SCC agent.**

  a) To shut down from the scc-console> prompt, enter:

  ```
  shutdown
  ```

  Warning! Do not enter **shutdown** at a UNIX prompt; it shuts down the operating system.

  To shut down from the UNIX command line, enter the **scc --stop** command.

  For a singleton installation:

  ```
  $SYBASE/SCC-3_2/bin/scc.sh --stop
  ```

  For an instance:

  ```
  $SYBASE/SCC-3_2/bin/scc.sh --stop -instance <instance-
  name>
  ```

  You can omit the **-instance** option if the instance's name is the same as its host name (the default).

- **Configure Sybase Control Center or the SCC agent to run as a service.**

  A UNIX service is a daemon process that starts automatically after the machine is started and runs in the background. UNIX installations of Sybase Control Center include a shell script, **sccd**, which you can use to configure the Sybase Control Center service. (Some

UNIX platforms supply tools that make service configuration easier; Linux **chkconfig** is an example.)

Note: Sybase recommends that if you are not familiar with setting up services in UNIX, you delegate this task to a system administrator or consult the system administration documentation for your UNIX platform.

a) Copy `$SYBASE/SCC-3_2/bin/sccd` into this directory:

- AIX (SCC agent only): `/etc/rc.d/init.d`
- HP-UX (SCC agent only): `/sbin/init.d`
- All other platforms: `/etc/init.d`

b) Open `sccd` and make these changes:

- Change the line that sets the SYBASE variable to the location of your Sybase installation (that is, the parent of `SCC-3_2`, the Sybase Control Center installation directory). By default, this directory is called `Sybase`.
- If you are not using shared-disk mode, or you are using shared-disk mode to run a single instance whose name is the same as the host name, skip to step *5.c* on page 77 or step *5.d* on page 78.
- If you are using shared-disk mode to run a single instance whose name is not the host name, or to run multiple instances on the same host, add the instance name to the script name. Change:

```
SCRIPT_NAME=scc.sh
```

to:

```
SCRIPT_NAME="scc.sh -instance <instance-name>"
```

- If you are using shared-disk mode to run multiple instances on the same host, append the instance name to the name of the ouput log file. Change:

```
./${SCRIPT_NAME} --start 2>&1 >> ${SCC_HOME}/log/scc-
service.out &
```

to:

```
./${SCRIPT_NAME} --start 2>&1 >> ${SCC_HOME}/log/scc-
service_<instance-name>.out &
```

- If you are using shared-disk mode to run multiple instances on the same host, save a copy of the `sccd` script for each instance, giving each copy a unique name. In each copy, add the instance name to the script name and append the instance name to the output log file name as described above. Perform the remaining steps in this procedure for each copy of `sccd`.

c) In Linux, configure the service to run in run levels 2, 3, 4, and 5:

```
/usr/sbin/chkconfig --add sccd
/usr/sbin/chkconfig --level 2345 sccd
```

You can test the `sccd` script with `/usr/sbin/service sccd status`. (The **service** command accepts these options: **start** | **stop** | **status** | restart.)

   d) On non-Linux platforms, locate this directory:
      * AIX (SCC agent only): `/etc/rc.d/rc<X>.d`
      * HP-UX (SCC agent only): `/sbin/rc<X>.d`
      * Solaris: `/etc/rc<X>.d`

   where <X> is the run level (for example, 3). Make two soft links in the directory for your platform and set the links to point to:

   * AIX (SCC agent only):
      `/etc/rc.d/init.d/sccd: S90sccd` and
      `/etc/rc.d/init.d/sccd: K10sccd`
   * HP-UX (SCC agent only):
      `/sbin/init.d/sccd: S90sccd` and
      `/sbin/init.d/sccd: K10sccd`
   * Solaris:
      `/etc/init.d/sccd: S90sccd` and
      `/etc/init.d/sccd: K10sccd`

   The `S90sccd` link starts the service and the `K10sccd` link stops the service. The two-digit numbers in the links indicate the start and stop priorities of the service.

   e) Use the `S90sccd` and `K10sccd` links to test starting and stopping the service. The links are called automatically when the machine is started or shut down.

See also
* *Registering the ODBC Driver in Windows* on page 71
* *Starting and Stopping Sybase Control Center in Windows* on page 72
* *Configuring Memory Usage* on page 78
* *scc Command* on page 81

Configuring Memory Usage

(Optional) Determine whether you need to configure how much memory Sybase Control Center uses, and if so which configuration method to use.

It is not usually necessary to configure memory usage for Sybase Control Center. This table lists memory options you can set and circumstances under which you should consider changing them.

Modify this value	When	Guidelines
Maximum memory • `jvmopt=-Xmx` – if you are running SCC as a Windows service • *SCC_MEM_MAX* – if you are running SCC as a UNIX service • *SCC_MEM_MAX* – if you are starting SCC from the command line	• You need to prevent Sybase Control Center from using more than a given amount of memory • SCC fails to start and may display an error: `Could not create the Java Virtual machine.` • An OutOfMemory error says SCC is out of heap space • A warning message about system memory appears during the start process • The machine where SCC is installed has less than 2GB of memory. (Starting SCC on a machine with less than 2GB of memory triggers the startup warning message about system memory.)	On machines with less than 2GB of memory, set maximum memory to 256MB or more. Default value: none. (On machines with 2GB or more of memory, maximum memory is set dynamically and is effectively limited only by the amount of system memory available.)
Permanent memory • `jvmopt=-XX:MaxPerm-Size` – if you are running SCC as a Windows service • *SCC_MEM_PERM* – if you are running SCC as a UNIX service • *SCC_MEM_PERM* – if you are starting SCC from the command line	An OutOfMemory error says SCC is out of permanent generation space	Increase by 32MB increments. If you reach a value equal to twice the default and still see the OutOfMemory error, contact Sybase technical support. Default value: 128MB

You can change memory options in two ways:

• For Sybase Control Center started from the command line – execute commands to set one or more environment variables before executing the **scc** command to start Sybase Control Center. When you use this method, your changes to the memory options last only as long as the current login session. This method is useful for testing new option values.

- For the Sybase Control Center service – modify a file used by the SCC service. When you use this method, your changes to the memory options persist—Sybase Control Center uses them every time it starts as a service.

See also
- *Registering the ODBC Driver in Windows* on page 71
- *Starting and Stopping Sybase Control Center in Windows* on page 72
- *Starting and Stopping Sybase Control Center in UNIX* on page 74
- *scc Command* on page 81

Changing a Memory Option on the Command Line

Before you start Sybase Control Center from the command line, you can issue a command to change the value of a memory option temporarily.

Changes made using this method last only as long as the current login session. This method is useful for testing new option values.

1. If Sybase Control Center is running, shut it down.
2. Set the environment variable. Specify a size in megabytes but do not indicate the units in the command.

   Windows example:
   ```
   > set SCC_MEM_MAX=512
   ```

   UNIX example:
   ```
   bash$ export SCC_MEM_MAX=512
   ```
3. Use the **scc** command to start Sybase Control Center.

See also
- *Changing a Memory Option for an SCC Windows Service* on page 80
- *Changing a Memory Option for an SCC UNIX Service* on page 81
- *Starting and Stopping Sybase Control Center in Windows* on page 72
- *Starting and Stopping Sybase Control Center in UNIX* on page 74
- *scc Command* on page 81

Changing a Memory Option for an SCC Windows Service

Add a **jvmopt** command to the `scc.properties` file to change a memory option (-Xmx or -XX:MaxPermSize) for a Sybase Control Center Windows service.

When you use this method to set memory options, your changes are permanent—Sybase Control Center uses them every time it starts as a service.

1. If Sybase Control Center is running, shut it down.

2. Open the SCC properties file:

```
<SCC-install-directory>\SCC-3_2\bin\scc.properties
```

3. Add (or modify, if it already exists) a **jvmopt** line specifying the memory size in Java format. Use m for megabytes or g for gigabytes.

For example:

```
jvmopt=-Xmx512m
```

4. Save the file and start the Sybase Control Center Windows service.

See also
* *Changing a Memory Option on the Command Line* on page 80
* *Changing a Memory Option for an SCC UNIX Service* on page 81
* *Starting and Stopping Sybase Control Center in Windows* on page 72

Changing a Memory Option for an SCC UNIX Service

To change a memory setting for a Sybase Control Center UNIX service, add the appropriate environment variable (*SCC_MEM_MAX* or *SCC_MEM_PERM*) to the sccd script.

When you use this method to set memory options, your changes are permanent—Sybase Control Center uses them every time it starts as a service.

1. If Sybase Control Center is running, shut it down.

2. Open the sccd file: /etc/init.d/sccd

3. Add the environment variable at the top of the file (after the comments). Specify a size in megabytes but do not indicate the units in the command.

For example:

```
SCC_MEM_MAX=512
```

4. Save the file and start the Sybase Control Center UNIX service.

See also
* *Changing a Memory Option on the Command Line* on page 80
* *Changing a Memory Option for an SCC Windows Service* on page 80
* *Starting and Stopping Sybase Control Center in UNIX* on page 74

scc Command

Use **scc.bat** (Windows) or **scc.sh** (UNIX) to start and stop Sybase Control Center agents and servers and to perform administrative tasks like configuring ports and enabling and disabling services.

Syntax

```
scc[.bat | .sh] [-a | --address RMI-service-address]
[-b | --bitwidth]
[--dbpassword]
```

```
[-disable | --disable service-name,service-name...]
[-enable | --enable service-name,service-name...]
[-h | --help]
[-I | --info [information-category]]
[-instance [instance-name]]
[-m | --message message-level]
[-password | --password password]

[-p | --port {port-name=port-number |
    service-name:property-name=port-number}]
[{-start | --start} | {-stop | --stop}]
[-status | --status]
[-user | --user login-name]
[-v | -version | --version]
```

Parameters

- **-a | --address** *RMI-service-address* – the address for the RMI service to use; must be an IP address on this machine or the name of this machine (which is the default).
- **-b | --bitwidth** – returns a string identifying the bit width (32 or 64) of the underlying platform; Sybase Control Center uses this option to determine which libraries to use for its internal database. If you use this option, the **scc** command does not start Sybase Control Center.
- **--dbpassword** – changes the password of the default dba account provided for the repository database. It prompts you for the new password, validates it, and starts the Sybase Control Center server. This option does not work if you start Sybase Control Center in the background—the server fails to start if there is no console.
- **-disable | --disable** *service-name,service-name...* – disable the specified Sybase Control Center services. This option does not work while Sybase Control Center is running or as part of a command that starts SCC. To use it, shut down SCC, execute `scc --disable`, then restart. See under --ports for service names; separate each service from the next with a comma.
- **-enable | --enable** *service-name,service-name...* – enable the specified Sybase Control Center services. See under --ports for service names; separate each service from the next with a comma. When you use this option, **scc** does not start Sybase Control Center—use a separate command to start SCC.
- **-h | --help** – display help and usage information for the **scc** command. If you use this option, **scc** does not start Sybase Control Center.
- **-I | --info [***information-category***]** – display the specified categories of information about Sybase Control Center. Separate each category from the next with a comma. The information categories are:

  - all – returns all the information provided by the sys, ports, and services categories. Default option.
  - sys – returns general information about this instance of Sybase Control Center, including the version, the home (installation) directory, the host machine's name and

82

IP address, the RMI port number, the messaging level, and details about the platform and Java installation.

- ports – lists all the ports on which the Sybase Control Center agent and its services listen, indicates whether each port is in use, and shows the service running on each port.
- services – lists all the services known to the Sybase Control Center agent, indicates whether each service is enabled, and lists other services on which each service depends.
- sysprop – lists all the Java system properties known the Java VM and their values.
- env – lists the complete Java VM process environment.

- **-instance [*instance-name*]** – use with other options (**-start** and **-stop**, for example) to specify a Sybase Control Center instance in a shared disk deployment. If you do not enter a name for the instance, it defaults to the host name.
- **-m | --message** *message-level* – set the amount of detail recorded in system logs; also known as the logging level. Valid values are OFF, FATAL, ERROR, WARN, INFO, DEBUG, and ALL. WARN is the default.
- **-password | --password** – specify the password of the user account Sybase Control Center will use to stop servers or query them for status. Use this option with **--user**. When you enter a command with **--user** but without **--password**, the console prompts you to enter a password.
- **-p | --port {*port-name=port-number* | *service-name:property-name=port-number*}** – configure the specified service to run on the specified port. Changing ports is useful if you discover a port conflict between Sybase Control Center and other software on the same system. When you use this option, **scc** does not start Sybase Control Center—use a separate command to start SCC.

Valid port names, service names and property names are:

Port Name	Description	Service Names	Property Names	Default Port
db	Database port Present on SCC server	SccSADataserver Messaging Alert Scheduler	com.sybase.asa.server.port messaging.db.port alert.database.port org.quartz.data-Source.ASA.URL	3638
http	Web HTTP port Present on SCC server	EmbeddedWebContainer	http.port	8282
https	Web HTTPS (secure HTTP) port Present on SCC server	EmbeddedWebContainer	https.port	8283

Port Name	Description	Service Names	Property Names	Default Port
jiniHttp	JINI HTTP server Present on SCC server and SCC agent	Jini	httpPort	9092
jiniR-mid	JINI remote method invocation daemon Present on SCC server and SCC agent	Jini	rmidPort	9095
msg	Messaging port Present on SCC server	Messaging	messaging.port	2000
rmi	RMI port Present on SCC server and SCC agent	RMI	port	9999
tds	Tabular Data Stream™ port (used to communicate with other Sybase products) Present on SCC server and SCC agent	Tds	tdsPort	9998

You can also execute scc --info ports to display service names and associated property names; they appear in the first two columns of the output.

- **-start | --start** – start the Sybase Control Center server. This is the default option—if you execute **scc** with no options, it starts SCC. This option cannot be combined in the same command with options that set ports or enable or disable services; use a separate **scc** command to start SCC.
- **-status | --status** – display a status message indicating whether the Sybase Control Center server is running.
- **-stop | --stop** – shut down the Sybase Control Center server if it is running.
- **-user | --user [*login-name*]** – specify the user account Sybase Control Center will use to stop managed servers or query them for status. Use this option with **--password**. If you do not enter a login name, the console prompts you to enter one.
- **-v | -version | --version** – display the version of Sybase Control Center software running on this server. If you use this option, **scc** does not start Sybase Control Center.

Examples

- **Set the RMI port** – each of these commands sets the RMI port to 9999 (the default). The first command illustrates the port name syntax; the second illustrates the service name:property name syntax.

```
scc --port rmi=9999
scc --port RMI:port=9999
```

- **Set the RMI port and start SCC** – these commands set the RMI port to 9996, then start SCC. Two commands (separated by a semicolon here) are needed because **scc** does not start Sybase Control Center when it includes any of the port-setting options.

```
scc -p rmi=9996; scc
```

- **Set all database ports** – this command sets all four of the SQL Anywhere database ports (data server, messaging, database alert, and scheduler) to 3638. (SQL Anywhere is the Sybase Control Center internal repository.)

```
scc --port db=3638
```

- **Set the TDS port** – this command sets the TDS port to 9998 (the default):

```
scc --port Tds:tdsPort=9998
```

- **Enable a service and start SCC** – the first **scc** command enables the TDS service; the second starts SCC. (The two commands are separated by a semicolon.) The second command is needed because **scc** does not start Sybase Control Center when it includes the **-enable** option.

```
scc -enable Tds; scc
```

- **Start an SCC instance** – this command starts the SCC instance called kalamazoo. **-start** is optional because it is the default.

```
scc -start -instance kalamazoo
```

Permissions

scc permission defaults to all users. No permission is required to use it.

See also
- *Registering the ODBC Driver in Windows* on page 71
- *Starting and Stopping Sybase Control Center in Windows* on page 72
- *Starting and Stopping Sybase Control Center in UNIX* on page 74
- *Configuring Memory Usage* on page 78
- *Configuring Ports* on page 102
- *Logging or Message Levels* on page 198

Logging in to Sybase Control Center

Enter the Sybase Control Center Web console.

Prerequisites
Install Adobe Flash Player in the browser you will use for SCC. See the *Sybase Control Center Installation Guide*.

Task

Sybase Control Center typically authenticates users through the operating system or an LDAP directory service. Consult your SCC administrator if you are not sure which login account to use for SCC.

Note: When logging in to a newly installed Sybase Control Center for which secure authentication has not been configured, use the sccadmin account (with no password, by default). For more information, see the *Sybase Control Center Installation Guide*.

1. Connect to the Sybase Control Center server. In your Web browser, enter: `https://scc-hostname:8283/scc`.

2. Enter your user name and password, and click **Login**.

   Tip: If you use a Windows account to log in to SCC, enter your user name in the format username@domain. Omit top-level domain extensions such as .com or .net—for example, enter fred@sybase, not fred@sybase.com.

See also
* *Logging out of Sybase Control Center* on page 86

Logging out of Sybase Control Center

When you finish working in Sybase Control Center, end your login session.

From the main menu bar, select **Application > Log Out.**

Alternatively, click **Log Out** in the upper-right corner of the window.

Note: If an administrator has configured the automatic logout feature, Sybase Control Center logs you out if your session is idle (no typing or mouse movement) for longer than the timeout period, which is set by the administrator.

If no automatic logout period is configured,

* A login session left open on a screen that refreshes (a monitor screen or a data collection job screen, for example) remains open indefinitely.
* A login session left open on a screen that does not change expires after 30 minutes. The next time you make a request of the server, SCC logs you out.

See also
* *Logging in to Sybase Control Center* on page 85

Setting Up Security

Configure login authentication and map roles.

Read about security and follow these procedures before you configure Sybase Control Center product modules.

Note: These security topics are intended for use in a production environment. If you are evaluating or testing SCC, see the *Installation Guide* for instructions on getting started quickly.

1. *Security*

Sybase Control Center can authenticate user logins through an LDAP server, through the operating system, or both.

2. *Configuring Authentication for Windows*

Authentication through the Windows operating system is enabled by default, but it requires some configuration. First, set Sybase Control Center to create an account when a Windows user logs in to Sybase Control Center.

3. *Configuring a Pluggable Authentication Module (PAM) for UNIX*

Set up Sybase Control Center to support username and password login using accounts on the UNIX operating system. Optionally, have Sybase Control Center create an account when a UNIX user first logs in to Sybase Control Center.

4. *Configuring an LDAP Authentication Module*

Configure an LDAP authentication module for Sybase Control Center by editing the security properties file to point to the correct LDAP server.

5. *Mapping Sybase Control Center Roles to LDAP or OS Groups*

To grant Sybase Control Center privileges to users who are authenticated through LDAP or the operating system, associate roles used in Sybase Control Center with groups in LDAP or the operating system.

6. *Encrypting a Password*

Use the passencrypt utility to encrypt passwords and other values that must be kept secure while stored in text files.

7. *Configuring Ports*

(Optional) Use the scc --port command to assign Sybase Control Center services to new ports.

Security

Sybase Control Center can authenticate user logins through an LDAP server, through the operating system, or both.

*   Sybase Control Center can be configured to authenticate through any LDAP server that supports the inetOrgPerson (RFC 2798) schema.
*   When Sybase Control Center authenticates through the operating system, it uses the operating system of the Sybase Control Center server machine (not the client).

Although you can create native user accounts in Sybase Control Center, Sybase does not recommend this approach to authentication. It is simpler and safer to configure Sybase Control Center to authenticate using existing LDAP, Windows, or UNIX login accounts.

Sybase strongly recommends that you use a common authentication provider for all Sybase products, including Sybase Control Center. A common authentication provider ensures that single sign-on works for users of Sybase Control Center and its managed servers.

Sybase Control Center requires each authenticated login account to have a predefined role. When a login is authenticated, roles for the login are retrieved by the security module and are mapped to Sybase Control Center predefined roles. Authorization is resolved through the mappings between the security module native roles and Sybase Control Center roles. You can enable mappings by creating a "sybase" group in your operating system or LDAP server and adding all Sybase Control Center users, or by modifying the Sybase Control Center `roles-map.xml` file to configure the mapping of native roles to Sybase Control Center roles. The security module authenticates the logins and authorizes access to managed resources.

Sybase Control Center provides a set of predefined login modules for authentication. All login modules are defined in the `<install_location>/SCC-3_2/conf/csi.properties` file. The syntax is defined by the Sybase Common Security Infrastructure (CSI) framework. You can configure the different login modules to customize security strength. The login modules are:

*   Simple Login – defines a user name, password, and a list of roles. The default user name is "sccadmin" with a blank password and a native role of "sccAdminRole". You can create additional accounts by adding simple login modules to `csi.properties`. However, Sybase does not recommend the use of simple login modules for authentication in production environments.

    Note: Add a password for the sccadmin account as soon as possible after you install Sybase Control Center. See the *Sybase Control Center Installation Guide* for instructions.

*   NT Proxy Login – delegates authentication to the underlying Windows operating system. When you log in to Sybase Control Center through an NT Proxy Login module, enter your user name in the format *username@nt-domain-name*. For example, user@sybase. Windows authentication is enabled by default, but it requires some configuration.
*   UNIX Proxy Login – delegates authentication to the underlying UNIX or Linux operating system using Pluggable Authentication Modules (PAM). When you log in to Sybase

Control Center through a UNIX PAM, enter your UNIX user name and password. UNIX authentication is enabled by default, but it requires some configuration.

* LDAP Login – delegates authentication to an LDAP server you specify. When you log in to Sybase Control Center through an LDAP server, enter your LDAP user name and password. LDAP authentication is not enabled by default; you must configure the login module.

See also
* *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 44
* *Configuring an LDAP Authentication Module* on page 91
* *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 100

Configuring Authentication for Windows

Authentication through the Windows operating system is enabled by default, but it requires some configuration. First, set Sybase Control Center to create an account when a Windows user logs in to Sybase Control Center.

This task is optional. However, if you choose not to create Sybase Control Center accounts automatically as described here, you must enter them manually. Sybase Control Center needs the accounts for purposes of setting authorization (user privileges).

1. Log in to Sybase Control Center using an account with administrative privileges. (The login account or its group must have sccAdminRole.)
2. Select **Application > Administration > Security.**
3. Check the box labeled **Automatically add SCC login records for authenticated logins**.
4. Check the box labeled **Automatically grant sccUserRole to newly created logins**.
5. Click **OK** to close the Security dialog.

Next

There are two next steps:

* If you opted not to automatically create Sybase Control Center login accounts, enter each account into Sybase Control Center manually.
* Whether you add accounts automatically or manually, you must grant privileges to any login accounts that require more than basic user access. You can grant privileges by assigning Sybase Control Center roles directly to the login accounts, or by assigning the login accounts to groups and mapping Sybase Control Center roles to the groups. The group approach is generally more efficient.

See also
* *Configuring an LDAP Authentication Module* on page 91

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 100
- *Adding a Login Account to the System* on page 110

Configuring a Pluggable Authentication Module (PAM) for UNIX

Set up Sybase Control Center to support username and password login using accounts on the UNIX operating system. Optionally, have Sybase Control Center create an account when a UNIX user first logs in to Sybase Control Center.

1. Using a login account with root privileges, configure the pluggable authentication module for your platform:

Platform	Action
Solaris	Append the contents of the `<SCC-install-dir>/utility/su-nos/pam.conf` file (provided with Sybase Control Center) to the `/etc/pam.conf` file on your Solaris platform.
Linux	Copy the `<SCC-install-dir>/utility/linux/sybase-ua` file (provided with Sybase Control Center) to the `/etc/pam.d` directory on your Linux platform.
	Note: The `sybase-ua` file provided with Sybase Control Center is not compatible with the most recent SUSE Linux versions. For SUSE 11 and later, see the example at the end of this topic.

Note: In the table above, the portion of the path that indicates the operating system might differ slightly from what is shown.

2. If the host UNIX system is not using a directory lookup for authentication (yp or NIS, for example) and authentication is carried out against the local `/etc/passwd` file, change the permissions on `/etc/shadow` to provide read access to the login account that executes SCC.
3. (Skip if you configured a PAM before starting Sybase Control Center) Restart Sybase Control Center.
4. (Optional) If you want Sybase Control Center to create an account when a UNIX user logs in to Sybase Control Center, execute these steps. If you choose not to create Sybase Control Center accounts automatically, you must enter them manually. Sybase Control Center needs the accounts for purposes of setting authorization (user privileges).
   a) Log in to Sybase Control Center using an account with administrative privileges (sccAdminRole).
   b) Select **Application > Administration > Security.**
   c) Check the box labeled **Automatically add SCC login records for authenticated logins**.
   d) Click **OK** to close the Security dialog.

Example: PAM for SUSE Linux 11 and later

For SUSE 11 and later, do not use the `sybase-ua` file provided with Sybase Control Center. Instead, in your `/etc/pam.d` directory, create a `sybase-ua` file that contains:

```
# sybase-ua PAM Configuration (SUSE style)
auth        include        common-auth
account     include        common-account
password    include        common-password
session     include        common-session
```

Next

There are two next steps:

- If you opted not to automatically create Sybase Control Center login accounts, enter each account into Sybase Control Center manually.
- Whether you add accounts automatically or manually, you must also grant privileges to the login accounts. You can grant privileges by assigning Sybase Control Center roles directly to the login accounts, or by assigning the login accounts to groups and mapping Sybase Control Center roles to the groups. The group approach is generally more efficient.

Configuring an LDAP Authentication Module

Configure an LDAP authentication module for Sybase Control Center by editing the security properties file to point to the correct LDAP server.

1. Open the `<SCC-install-dir>\conf\csi.properties` file.
2. Uncomment the LDAP module in the properties file by removing the # symbol at the beginning of each line (or, if necessary, add an LDAP module to the file). The sample module below specifies the LDAP server that will provide user authentication.

   The sample module shows the properties used for an OpenDS LDAP server. See the example at the end for values that work for ActiveDirectory. Configuration properties you can use in the LDAP module are described in a subtopic.

   Each line of the LDAP server module of the properties file must begin with "CSI.loginModule." followed by a module number. (The module number in this sample is 7.) The module number you assign must be unique in the properties file, and you must use the same module number in every line of the module.

```
CSI.loginModule.
7.options.AuthenticationSearchBase=ou=users,dc=example,dc=com
CSI.loginModule.7.options.BindDN=cn=Directory Manager
CSI.loginModule.7.options.BindPassword=secret
CSI.loginModule.7.options.DefaultSearchBase=dc=example,dc=com
CSI.loginModule.7.options.ProviderURL=ldap://localhost:10389
CSI.loginModule.
7.options.RoleSearchBase=ou=groups,dc=example,dc=com
CSI.loginModule.7.options.ServerType=openldap
CSI.loginModule.7.options.moduleName=LDAP Login Module
CSI.loginModule.7.controlFlag=sufficient
```

```
CSI.loginModule.
7.provider=com.sybase.ua.services.security.ldap.LDAPLoginModule
```

Note: Change the values of bolded lines only.

3. Save the file.

4. If your LDAP server's SSL certificate is signed by a nonstandard certificate authority (for example, if it is a self-signed certificate), use the **keytool** utility to configure your JVM or JDK to trust the certificate. Execute a command similar to this:

```
keytool -import -keystore <sybase-dir>/shared/JRE-6_0_6/bin/
keytool/lib/security/cacerts -file
<your cert file and path> -alias ldapcert -storepass changeit
```

LDAP configuration values for ActiveDirectory

For an ActiveDirectory server, use these values for configuration properties in your LDAP login module:

```
ServerType: msad2K
DefaultSearchBase: dc=<domainname>,dc=<tld> or o=<company
name>,c=<country code>
                E.g. dc=sybase,dc=com or o=Sybase,c=us
ProviderUrl: ldaps://<hostname>:<port>
                E.g.:  ldaps://myserver:636
AuthenticationFilter: (&(userPrincipalName={uid})
(objectclass=user))
BindDN: <User with read capability for all users>
BindPassword: <Password for BindDN user>
RoleFilter: (|(objectclass=groupofnames) (objectclass=group))
controlFlag: sufficient
```

Next
There are two additional steps:

- Set up roles and passwords for LDAP
- Map Sybase Control Center role to LDAP groups

See also
- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 100

Setting Up Roles and Passwords
Set the initial user roles and passwords required for Sybase Control Center to authenticate through an LDAP server.

Prerequisites
Configure an LDAP authentication module.

Task

1. Open the `<SCC-install-dir>\conf\roles-map.xml` file and add an LDAP login module.

   Insert an LDAP login module similar to this at the end of the security-modules portion of the file, just before `</security-modules>`:

   ```
       <module name="LDAP Login Module">
         <role-mapping modRole="sybase"
   uafRole="uaAnonymous,uaPluginAdmin,sccUserRole" />
         <role-mapping modRole="administrators"
   uafRole="uaAnonymous,sccAdminRole" />
       </module>
   ```

2. Ensure that the roles defined in the LDAP repository match the roles defined in `roles-map.xml`.

3. In the `<SCC-install-dir>\conf\csi.properties` file, set the BindPassword and ProviderURL properties with values used in your deployment.

   Sybase recommends that you encrypt sensitive values before saving them in `csi.properties`.

Next
Map Sybase Control Center roles to LDAP groups.

See also
- *LDAP Configuration Properties* on page 94
- *Encrypting a Password* on page 101

LDAP Configuration Properties

Use these properties in your `csi.properties` file to control your LDAP service.

Property	Default Value	Description
ServerType	None	Optional. The type of LDAP server you are connecting to: • `sunone5` -- SunOne 5.x OR iPlanet 5.x • `msad2k` -- Microsoft ActiveDirectory, Windows 2000 • `nsds4` -- Netscape Directory Server 4.x • `openldap` -- OpenLDAP Directory Server 2.x The value you choose establishes default values for these other authentication properties: • RoleFilter • UserRoleMembership • RoleMemberAttributes • AuthenticationFilter • DigestMD5Authentication • UseUserAccountControl
ProviderURL	`ldap://local-host:389`	The URL used to connect to the LDAP server. Use the default value if the server is: • Located on the same machine as your product that is enabled with the common security infrastructure. • Configured to use the default port (389). Otherwise, use this syntax for setting the value: `ldap://<hostname>:<port>`

Property	Default Value	Description
DefaultSearchBase	None	The LDAP search base that is used if no other search base is specified for authentication, roles, attribution and self registration: 1. `dc=<domainname>,dc=<tld>` For example, a machine in sybase.com domain would have a search base of dc=sybase,dc=com. 2. `o=<company name>,c=<country code>` For example, this might be o=Sybase,c=us for a machine within the Sybase organization.
SecurityProtocol	None	The protocol to be used when connecting to the LDAP server. To use an encrypted protocol, use "ssl" instead "ldaps" in the url. **Note:** ActiveDirectory requires the SSL protocol when setting the value for the password attribute. This occurs when creating a user or updating the password of an existing user.
AuthenticationMethod	simple	The authentication method to use for all authentication requests into LDAP. Legal values are generally the same as those of the java.naming.security.authentication JNDI property. Choose one of: • simple — For clear-text password authentication. • DIGEST-MD5 — For more secure hashed password authentication. This method requires that the server use plain text password storage and only works with JRE 1.4 or later.

Property	Default Value	Description
AuthenticationFilter	For most LDAP servers: `(&(uid={uid})(object-class=person))` or For Active Directory email lookups: `(&(userPrinci-palName={uid})(object-class=user))[ActiveDirec-tory]` For Active Directory Windows username lookups: `(&(sAMAc-count-Name={uid})(object-class=user))`	The filter to use when looking up the user. When performing a username based lookup, this filter is used to determine the LDAP entry that matches the supplied username. The string "{uid}" in the filter is replaced with the supplied username.
AuthenticationScope	onelevel	The authentication search scope. The supported values for this are: • `onelevel` • `subtree` If you do not specify a value or if you specify an invalid value, the default value is used.
AuthenticationSearchBase	none	The search base used to authenticate users. If this value is not specified, the LDAP DefaultSearch-Base is used.

Property	Default Value	Description
BindDN	none	The user DN to bind against when building the initial LDAP connection.
		In many cases, this user may need read permissions on all user records. If you do not set a value, anonymous binding is used. Anonymous binding works on most servers without additional configuration.
		However, the LDAP attributer may also use this DN to create the users in the LDAP server. When the self-registration feature is used, this user may also need the requisite permissions to create a user record. This behavior can occur if you do not set useUserCredentialsToBind to `true`. In this case, the LDAP attributer uses this DN to update the user attributes.
BindPassword	none	BindPassword is the password for BindDN, which is used to authenticate any user. BindDN and BindPassword are used to separate the LDAP connection into units.
		The AuthenticationMethod property determines the bind method used for this initial connection.

		If you use an encrypted the password using the CSI encryption utility, append .e to the property name. For example:
		
		`CSI.loginModule.7.options.`
		`BindPassword.e=1-AAAAEgQQOLL+LpX`
		`JO8fO9T4SrQYRC9lRT1w5ePfdczQTDs`
		`P8iACk9mDAbm3F3p5a1wXWKK8+NdJuk`
		`nc7w2nw5aGJlyG3xQ==`

RoleSearchBase	none	The search base used to retrieve lists of roles. If this value is not specified, the LDAP DefaultSearchBase is used.

Property	Default Value	Description
RoleFilter	For SunONE/iPlanet: `(&(object-class=ldapsu-bentry) (ob-jectclass=nsro-ledefinition))` For Netscape Directory Server: `(object-class=groupof-names) (object-class=groupofu-niquenames))` For ActiveDirectory: `(object-class=groupof-names) (object-class=group))`	The role search filter. This filter should, when combined with the role search base and role scope, return a complete list of roles within the LDAP server. There are several default values depending on the chosen server type. If the server type is not chosen or this property is not initialized, no roles are available.
RoleMemberAttributes	For Netscape Directory Server: member,unique-member	The role's member attributes defines a comma-delimited list of attributes that roles may have that define a list of DN's of people who are in the role. These values are cross referenced with the active user to determine the user's role list. One example of the use of this property is when using LDAP groups as placeholders for roles. This property only has a default value when the Netscape server type is chosen.
RoleNameAttribute	cn	The attribute for retrieved roles that is the common name of the role. If this value is "dn" it is interpreted specially as the entire dn of the role as the role name.
RoleScope	onelevel	The role search scope. The supported values for this are: • `onelevel` • `subtree` If you do not specify a value or if you specify an invalid value, the default value is used.

Property	Default Value	Description
UserRoleMembershipAttributes	For iPlanet/SunONE: nsRoleDN For ActiveDirectory: memberOf For all others: none	The user's role membership attributes property is used to define an attribute that a user has that contains the DN's of all of the roles as user is a member of. These comma-delimited values are then cross-referenced with the roles retrieved in the role search base and search filter to come up with a list of user's roles.
UserFreeformRoleMembershipAttributes	None	The "freeform" role membership attribute list. Users who have attributes in this comma-delimited list are automatically granted access to roles whose names are equal to the attribute value. For example, if the value of this property is "department" and user's LDAP record has the following values for the department attribute, { "sales", "consulting" }, then the user will be granted roles whose names are "sales" and "consulting".
Referral	ignore	The behavior when a referral is encountered. The valid values are those dictated by LdapContext, for example, "follow", "ignore", "throw".
DigestMD5AuthenticationFormat	DN For OpenLDAP: Username	The DIGEST-MD5 bind authentication identity format.
UseUserAccountControlAttribute	For most LDAP servers: false For ActiveDirectory: true	The UserAccountControl attribute to be used for detecting disabled user accounts, account expirations, password expirations and so on. ActiveDirectory also uses this attribute to store the above information.
controlFlag	optional	Indicates whether authentication with this login module is sufficient to allow the user to log in, or whether the user must also be authenticated with another login module. Rarely set to anything other than "sufficient" for any login module. **Note:** controlFlag is a generic login module option rather than an LDAP configuration property.

See also
*   *Setting Up Roles and Passwords* on page 92

Mapping Sybase Control Center Roles to LDAP or OS Groups

To grant Sybase Control Center privileges to users who are authenticated through LDAP or the operating system, associate roles used in Sybase Control Center with groups in LDAP or the operating system.

You can configure Sybase Control Center to enable users to authenticate through their local operating system or through an LDAP server. To make this type of authentication work, SCC roles must be mapped to groups that exist in the system providing authentication (LDAP or the operating system) or in the login module.

By default, SCC assumes there is a "sybase" group in the authenticating system and maps the LDAP or OS "sybase" group to SCC roles to provide basic privileges. The table lists additional default mappings of LDAP and OS groups to SCC roles.

Login Module	OS Group	Sybase Control Center Roles
UNIX Proxy	root	uaAnonymous, uaAgentAdmin, uaOSAdmin
	sybase	uaAnonymous, uaPluginAdmin, sccUserRole
	user	uaAnonymous, uaUser
	guest	uaAnonymous, uaGuest
NT Proxy	Administrators	uaAnonymous, uaAgentAdmin, uaOSAdmin
	sybase	uaAnonymous, uaPluginAdmin, sccUserRole
	Users	uaAnonymous, uaUser
	Guests	uaAnonymous, uaGuest
LDAP	sybase	uaAnonymous, uaPluginAdmin, sccUserRole

There are two ways to accomplish the mapping:

- (Recommended) Add a "sybase" group to the operating system or LDAP server Sybase Control Center is using to authenticate users, and add all users who need to access Sybase Control Center to the "sybase" group.
- Configure Sybase Control Center to use an existing group in LDAP or the operating system by editing the `roles-map.xml` file. This option is described here.

1. If Sybase Control Center is running, shut it down.
2. In a text editor, open:

   `<SCC-install-directory>/conf/roles-map.xml`
3. Locate the appropriate login module: UNIX or NT (for Windows).

4. Copy the line that maps the "sybase" group and paste it into the module just above the original sybase line.

5. Change "sybase" to the name of the group in your operating system to which Sybase Control Center users belong.

   For example, if the group is SCCusers, the new line should look like this:

   ```
   <role-mapping modRole="SCCusers"
   uafRole="uaAnonymous,uaPluginAdmin,sccUserRole" />
   ```

6. Save the file and exit.

7. Start Sybase Control Center.

See also
- *Configuring an LDAP Authentication Module* on page 91
- *Configuring Authentication for Windows* on page 89
- *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 44
- *Assigning a Role to a Login or a Group* on page 106
- *User Authorization* on page 106

Encrypting a Password

Use the **passencrypt** utility to encrypt passwords and other values that must be kept secure while stored in text files.

You can safely store an encrypted password in a properties file. Enter the password in clear text (unencrypted) when you execute **passencrypt** and when you use the password to log in.

passencrypt, which is located in the Sybase Control Center bin directory, uses the DES encryption algorithm.

1. Open a command window and change to the bin directory:

   Windows: `cd <SCC-install-directory>\bin`
   UNIX: `cd <SCC-install-directory>/bin`

2. To encrypt a password, enter **passencrypt**. Enter your new password at the resulting prompt.
   passencrypt encrypts the password you enter (which does not appear on the screen) and displays the password in encrypted form.

3. Copy the encrypted password.

4. Paste the encrypted password where needed.

See also
- *Setting Up Roles and Passwords* on page 92

Configuring Ports

(Optional) Use the **scc --port** command to assign Sybase Control Center services to new ports.

Prerequisites

Check for port conflicts between Sybase Control Center and other software running on the same host.

Task

Sybase Control Center cannot function properly if other services use its ports. If you discover a conflict with any port listed in the right column below, you can either reconfigure the other service's port or reconfigure Sybase Control Center as described here.

Port Name	Description	Service Names	Property Names	Default Port
db	Database port Present on SCC server	SccSADataserver Messaging Alert Scheduler	com.sybase.asa.server.port messaging.db.port alert.database.port org.quartz.data-Source.ASA.URL	3638
http	Web HTTP port Present on SCC server	EmbeddedWebContainer	http.port	8282
https	Web HTTPS (secure HTTP) port Present on SCC server	EmbeddedWebContainer	https.port	8283
jiniHttp	JINI HTTP server Present on SCC server and SCC agent	Jini	httpPort	9092
jiniRmid	JINI remote method invocation daemon Present on SCC server and SCC agent	Jini	rmidPort	9095
msg	Messaging port Present on SCC server	Messaging	messaging.port	2000

Port Name	Description	Service Names	Property Names	Default Port
rmi	RMI port Present on SCC server and SCC agent	RMI	port	9999
tds	Tabular Data Stream™ port (used to communicate with other Sybase products) Present on SCC server and SCC agent	Tds	tdsPort	9998

1. Shut down Sybase Control Center.
2. Execute **scc --info ports** to display a list of Sybase Control Center services, their properties, and their assigned ports.
3. To reassign a port, enter a command in one of these formats:

   ```
   scc --port port-name=port-number
   scc --port service-name:property-name=port-number
   ```

   Use the first, simpler format unless you want to configure the database services to use different ports. (By default, they all use the same port.)
4. Start Sybase Control Center.
5. Execute **scc --info ports** again to confirm that the port has been reassigned.

Examples

Set all four database services (data server, messaging, database alert, and scheduler) to the same port, 3639. (The database is SQL Anywhere, used by the Sybase Control Center internal repository.)

```
scc --port db=3639
```

Set only the database messaging service to port 3639.

```
scc --port Messaging:messaging.db.port=3639
```

Set the HTTP port to 9292.

```
scc --port http=9292
```

Set the Jini RMI daemon to port 9696.

```
scc --port jiniRmid=9696
```

Set the main Sybase Control Center messaging service to port 2001.

```
scc --port msg=2001
```

Set the RMI port to 9991.

```
scc --port rmi=9991
```

Set the Tabular Data Stream port to 9997.

```
scc --port tds=9997
```

Note: scc commands that include a port-setting option (**-p** or **--port**) do not start Sybase Control Center. To start SCC, execute a separate **scc** command.

See also
* *scc Command* on page 81

Configuring the E-mail Server

(Optional) Specify the e-mail server for Sybase Control Center to use to send e-mail alert notifications.

Prerequisites
Launch Sybase Control Center and log in using an account with administrative privileges. (The login account or its group must have sccAdminRole.)

Task

1. From the menu bar, select **Application > Administration.**
2. Select **General Settings**.
3. Click the **E-mail** tab.
4. Enter the name of the e-mail server through which Sybase Control Center will send alert notifications.
5. Change the default e-mail server port only in consultation with your e-mail administrator.
6. (Optional) Click **Customize e-mail settings** to display options for setting the domain name and e-mail sender for alert e-mail notifications.
7. (Optional) Enter your domain name (for example, mycompany.com).

   Most e-mail servers do not require SCC to provide an explicit domain name. Try providing a domain name here if your first attempt to configure e-mail alerts fails.
8. (Optional) Change the default e-mail sender name.

   This name appears in the "From" field of SCC e-mail alert messages. Do not use spaces; use hyphens or underscore characters instead.

   Tip: If you have multiple SCC servers, configure their sender names so you can tell which SCC an alert is coming from. For example, SybaseControlCenter_Boston or SCC_test11.

9. (Optional) If you entered anything in the **E-mail Domain name** or **E-mail sender name** fields, click **Apply** to make the test e-mail option reappear.

10. (Optional) To dispatch a test message, enter an e-mail address in the **Test e-mail address** field and click **Send**.

   If the test e-mail is received, you have properly configured the server for e-mail alert notifications.

11. Click **OK** (to apply the change and close the properties dialog) or **Apply** (to apply the change and leave the dialog open).

Next
(Optional) Configure automatic logout.

See also
- *Launching Sybase Control Center* on page 70
- *Logging in to Sybase Control Center* on page 85

Configuring the Automatic Logout Timer

(Optional) Set Sybase Control Center to end login sessions when users are inactive for too long.

Prerequisites
Launch Sybase Control Center and log in using an account with administrative privileges. (The login account or its group must have sccAdminRole.)

Task

1. From the menu bar, select **Application > Administration.**

2. Select **General Settings**.

3. Click the **Auto-Logout** tab.

4. Enter the number of minutes after which an idle user will be automatically logged out.

   Enter 0 or leave the box empty to disable automatic logout.

5. Click **OK** (to apply the change and close the properties dialog) or **Apply** (to apply the change and leave the dialog open).

See also
- *Launching Sybase Control Center* on page 70
- *Logging in to Sybase Control Center* on page 85

User Authorization

The authorization mechanism in Sybase Control Center employs login accounts and task-based roles.

Access to Sybase Control Center is controlled by login accounts. You grant permissions to a login account by assigning predefined roles that control tasks the user can perform in Sybase Control Center, such as administration and monitoring of particular types of Sybase servers. The roles can be assigned directly to login accounts or to groups; a login account inherits the roles of any group to which it belongs. Component product modules assign some roles automatically.

Sybase Control Center classifies roles as follows:

- System roles – define how a user can interact with Sybase Control Center.
- Product roles – define how a user can interact with a particular managed resource in Sybase Control Center, for example the Replication Server named RepBoston01.

Note: The tools described here are for managing SCC-enabled login accounts; you cannot use them to manage accounts and groups that are native to your managed resource.

See also
- *Authenticating a Login Account for a Managed Resource* on page 126

Assigning a Role to a Login or a Group

Use the security configuration options to add one or more roles to a Sybase Control Center login account or to a group. Roles enable users to perform tasks such as monitoring servers or administering Sybase Control Center.

Prerequisites
You must have administrative privileges (sccAdminRole) to perform this task. To assign a monitoring role for a server, first register the server.

Task
Assign the sccAdminRole to any login account that will perform administrative tasks in Sybase Control Center.

1. From the application menu bar, select **Application > Administration.**
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Click **Logins** or **Groups**.
4. In the table, select the login account or group to which you want to assign a role.
5. Click the **Roles** tab.

6. In the **Available roles for resource** list, select the role, then click **Add**. For example, to grant administrative privileges, add the SCC Service:sccAdminRole. To grant monitoring privileges, add the MonitorRole for the desired server and server type.

    Note: Sybase Control Center product modules assign certain roles automatically, so you might not need to add a MonitorRole.

    If a role appears in the **Has following roles** list, this account or group has already been configured with that role.

7. Click **OK**.

See also
*   *Removing a Role from a Login or a Group* on page 107

Removing a Role from a Login or a Group

Use the security configuration options to remove one or more roles from a Sybase Control Center login acount or from a group.

Prerequisites

You must have administrative privileges to perform this task.

Task

1. From the menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Click **Logins** or **Groups**.
4. Select the login account or group from which you want to remove a role.
5. Click the **Roles** tab.
6. Select the role, then click **Remove**.
7. Click **OK**.

See also
*   *Assigning a Role to a Login or a Group* on page 106

Adding a Group

Use the security configuration options to create a new group.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

Groups can make roles easier to manage. Rather than assigning roles to individual users, assign roles to groups and add users to the groups or remove them as needed.

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Select **Groups**.
4. Click **Create Group**.
5. Enter a group name and a description.
6. Click **Finish**.

See also
- *Removing a Group* on page 108
- *Adding a Login Account to a Group* on page 109
- *Removing a Login Account from a Group* on page 109

Removing a Group

Use the security configuration options to remove a group.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Select **Groups**.
4. Select the group to remove.
5. Click **Delete**.
6. Click **OK** to confirm the deletion.

See also
- *Adding a Group* on page 107
- *Adding a Login Account to a Group* on page 109
- *Removing a Login Account from a Group* on page 109

Adding a Login Account to a Group

Use the security configuration options to add one or more login accounts to a group.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Click **Groups**.
4. Select the group to which you want to assign an account.
5. Click the **Membership** tab.
6. Select the account, then click **Add**.
7. Click **OK**.

See also
- *Adding a Group* on page 107
- *Removing a Group* on page 108
- *Removing a Login Account from a Group* on page 109

Removing a Login Account from a Group

Use the security configuration options to remove one or more login accounts from a group.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties, expand the **Security** folder.
3. Select **Groups**.
4. Select the group from which to remove members.
5. Click the **Membership** tab.
6. Select the login, then click **Remove**.
7. Click **OK**.

See also

- *Adding a Group* on page 107
- *Removing a Group* on page 108
- *Adding a Login Account to a Group* on page 109

Adding a Login Account to the System

Use the security configuration options to create a native login account in Sybase Control Center.

Prerequisites

- You must have administrative privileges (sccAdminRole) to perform this task.
- If you intend to use LDAP or the operating system to authenticate users, configure the appropriate authentication module.

Task

Note: Sybase does not recommend that you create a native login account for every Sybase Control Center user. It is more efficient to configure Sybase Control Center to authenticate users through their user accounts in LDAP or the operating system.

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Select **Logins**.
4. Click **Create Login**.
5. Enter a login name and expiration for the new account. Expiration is optional.
6. Click **Next**.
7. Select **Specify new user information**.
8. Enter details about the user:
   - Title
   - First name*
   - M.I. (middle initial)
   - Last name*
   - Suffix
   - E-mail address*
   - Phone
   - Ext.
   - Fax
   - Mobile
   - Supports text messaging (checkbox)

*You must fill in the **First Name**, **Last Name**, and **E-mail Address** fields.

9. Click **Finish**.

Note: If you are using the predefined Simple Login module for authentication, the default login accounts, "sccadmin" and "sccuser," come with blank passwords. To change or modify the passwords, configure the `csi.properties` file as described in the *Installation Guide*.

Next
Grant privileges to the new login account. You can grant privileges by assigning Sybase Control Center roles directly to the login accounts, or by assigning the login accounts to groups and mapping Sybase Control Center roles to the groups. The group approach is generally more efficient.

See also

Removing a Login Account from the System
Use the security configuration options to delete a Sybase Control Center login account.

Prerequisites
You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Select **Logins**.
4. Select the login to delete.
5. Click **Delete**.
6. Click **OK** to confirm the deletion.

See also

Modifying a User Profile

Use the security configuration options to suspend a login account, impose an expiration date, or modify the account's user information.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration.**
2. In the Sybase Control Center Properties dialog, expand the **Security** folder.
3. Select **Logins**.
4. Select the login account to modify.
5. Click the **General** tab.
6. To suspend this account, click **Login disabled**.
7. To set the date on which this account will stop working, click the calendar icon next to the **Expiration** field and select a data.
8. Click **Apply**.
9. Click the **User Info** tab.
10. Edit the user information.

    When this user configures e-mail alert subscriptions, Sybase Control Center automatically populates the subscription dialog with the e-mail address you enter here.

11. Click **Apply**.

See also

*   *Adding a Login Account to the System* on page 110
*   *Removing a Login Account from the System* on page 111

Logins, Roles, and Groups

Sybase Control Center includes predefined login accounts and roles.

In Sybase Control Center, a login account identifies a user who can connect to the application. An account may have roles that specify the tasks the user is allowed to perform.

Sybase Control Center is designed to delegate user authentication to the operating system or to an LDAP directory service. Delegation requires some configuration, however, so Sybase Control Center comes with two predefined login accounts. Sybase recommends using the predefined accounts only for installing and setting up Sybase Control Center. These accounts are not intended for use in a production environment.

Table 12. Predefined accounts

Login name	Description
sccadmin	Can use all the administration features in Sybase Control Center
sccuser	Test account with no special privileges

A role is a predefined profile that can be assigned to a login account or a group. Roles control the access rights for login accounts. Sybase Contol Center comes with predefined roles that are intended for use in production environments.

Table 13. Predefined roles

Role	Description
sccUserRole	Provides nonadministrative access to Sybase Control Center. Required for every user.
sccAdminRole	Provides administrative privileges for managing Sybase Control Center.

Monitoring privileges for SCC product modules are assigned automatically.

A group is made up of one or more login accounts; all the accounts in a group have the roles granted to the group. In Sybase Control Center you can create groups to suit your business requirements.

Get Started

Configure

Configure Sybase Control Center for Sybase IQ.

1. *Configuring Sybase IQ for Monitoring*

   To enable one or more Sybase IQ users to authenticate a Sybase IQ server with Sybase Control Center, create and populate the SCC_MONITOR group.

2. *Configuring Sybase IQ for Administration*

   To perform administration tasks, you must have the correct authority or group membership, and you may need to register the server's Sybase Control Center agent.

3. *Registering a Sybase IQ Server*

   Make Sybase Control Center aware of a Sybase IQ resource (for example, a server that can be monitored) and its connection information by registering the resource.

4. *Importing Resources for Batch Registration*

   (Optional) Import and register multiple servers from an interfaces or sql.ini file.

5. *Registering and Authenticating a Sybase Control Center Agent*

   Register and authenticate the Sybase Control Center agent for a managed server.

6. *Creating a Perspective*

   Create a perspective in which you can add and manage resources.

7. *Adding a Resource to a Perspective*

   Add one or more resources to the current perspective.

8. *Authenticating a Login Account for a Managed Resource*

   Specify the login account Sybase Control Center will use when it connects to your server or agent to collect monitoring data or manage the resource.

9. *Changing Update Frequency for Statistics and Charts*

   You can control the rate at which data on monitor screens and charts is refreshed, the amount of time covered by charts, and the multiplex nodes included in charts.

10. *Setting Up Statistics Collection*

    Use the Properties view of your managed resource to create a data collection job and add a schedule to the job.

11. *Creating an Alert*

    Use the Add Alert wizard to create an alert instance for your resource.

12. *Optional Configuration Steps*

    Perform additional configuration, including user authorization, alerts, data collection scheduling, backups, and setting purging options for the repository.

Configuring Sybase IQ for Monitoring

To enable one or more Sybase IQ users to authenticate a Sybase IQ server with Sybase Control Center, create and populate the SCC_MONITOR group.

Prerequisites

You are logged in to Sybase IQ as a user with DBA authority.

Task

To monitor a resource with SCC, you must authenticate the resource as a Sybase IQ user with DBA authority or membership in the SCC_MONITOR group.

Note: If you are performing a quick start, you need only authenticate your resource with Sybase Control Center using a Sybase IQ account with DBA authority (such as DBA). You can skip the steps below until you do a complete production set-up of SCC.

1. Using Interactive SQL or another SQL command tool, execute
   scc_iq_monitor_privileges_setup.sql, located in the directory `SCC-3_2/plugins/IQMAP`.
   The script creates the SCC_MONITOR group and grants a set of permissions.
2. Assign one or more Sybase IQ users or groups to the SCC_MONITOR group. You can do this by either of these methods:

   - Using Interactive SQL or another SQL command tool, execute **grant membership in group SCC_MONITOR to** <user/group>
   - In a user interface tool such as Sybase Control Center or Sybase Central, add the user to the SCC_MONITOR group

Configuring Sybase IQ for Administration

To perform administration tasks, you must have the correct authority or group membership, and you may need to register the server's Sybase Control Center agent.

Prerequisites

- You have monitoring permissions
- The Sybase IQ server you want to manage has a Sybase Control Center agent running on the same machine

Task

Before you can perform management tasks on a Sybase IQ server, your login account must have the required permissions. To perform certain tasks, including starting and stopping the Sybase IQ server, syncing the server, and adding secondary nodes, you must also register and authenticate the server's SCC agent.

1. Assign the authorities required for administrative tasks to users who will perform those tasks. You can do this by granting the authorities directly to user accounts or by granting the authorities to a group and adding users to the group.
2. Register and authenticate the SCC agent associated with the Sybase IQ server you want to manage.

   Note: You must register the managed server before you can register its SCC agent.

See also

- *Registering a Sybase IQ Server* on page 120
- *Registering and Authenticating a Sybase Control Center Agent* on page 123
- *Granting an Authority to a User or Group* on page 286

Required Authorities/Groups for Administration

Menu items are disabled unless you have the correct authority, or group membership, for a particular administration task.

Administration Category	Administration Task	Required Authorities/ Groups
Monitoring	Monitor	DBA authority or SCC_MONI- TOR group

Administration Category	Administration Task	Required Authorities/ Groups
Multiplex administration	Add secondary node	DBA authority (or Multiplex Admin and Backup and Space Admin authority)
	Drop secondary node	DBA authority
	Include a secondary node	DBA authority (or Multiplex Admin and Backup authority)
	Exclude a secondary node	DBA or Multiplex Admin authority
	Synchronize a secondary node	DBA authority
	Configure server	DBA or Multiplex Admin authority
	Change reader to writer, or writer to reader	DBA or Multiplex Admin authority
	Change designated failover node	DBA or Multiplex Admin authority
	Manual failover	DBA or Multiplex Admin authority
	Change multiplex name	DBA or Multiplex Admin authority
	Modify multiplex property sheet	DBA or Multiplex Admin authority
Simplex and multiplex server administration	Start server	n/a
	Stop server	DBA authority
	Multiplex server property sheet	IQ server is authenticated
Logical server administration	Create logical server	DBA or Multiplex Admin authority
	Delete logical server	DBA or Multiplex Admin authority
	Configure logical server	DBA or Multiplex Admin authority

Administration Category	Administration Task	Required Authorities/ Groups
	Edit logical server policy	DBA or Multiplex Admin authority
	Modify logical server property sheet	DBA or Multiplex Admin authority
Database administration	Create database	Utility DB user and password must be passed from the user to the Sybase Control Center agent.
	Edit database option	DBA authority
	Edit license management	DBA authority
	Modify database property sheet	DBA authority
Dbspace administration	Create dbspace	DBA or Space Admin authority
	Delete dbspace	DBA or Space Admin authority
	Pre-allocate catalog dbspace	DBA or Space Admin authority
	Change dbspace name, mode, status, striping, stripe size	DBA or Space Admin authority
	Create db file	DBA or Space Admin authority
	Delete db file	DBA or Space Admin authority
	Empty file	DBA or Space Admin authority
	Change file name, path, mode, size	DBA or Space Admin authority
	Modify dbspace property sheet	DBA or Space Admin authority
Security administration	Create user/group	DBA (or User Admin & Perms Admin) authority
	Delete user/group	DBA or User Admin authority
	Change user to group	DBA or Perms Admin authority
	Change group to user	DBA or Perms Admin authority
	Edit user/group options	DBA authority
	Edit user/group properties	DBA (or User Admin & Perms Admin) authority

Administration Category	Administration Task	Required Authorities/ Groups
	Edit user/group authorities	DBA or Perms Admin authority
	Manage parent groups	DBA or Perms Admin authority
	Manage group members	DBA or Perms Admin authority
	Create login policy	DBA or User Admin authority
	Delete login policy	DBA or User Admin authority
	Edit login policy options and comment	DBA or User Admin authority
	Modify user/group property sheet	DBA (or User Admin & Perms Admin) authority
	Modify login policy property sheet	DBA or User Admin authority

Registering a Sybase IQ Server

Make Sybase Control Center aware of a Sybase IQ resource (for example, a server that can be monitored) and its connection information by registering the resource.

Prerequisites
Ensure that the Sybase IQ server does not have multiple databases. Sybase Control Center for Sybase IQ supports a maximum of one database per server.

Task

1. In the Resource Explorer, select **Resources > Register.**
2. Specify:

Table 14. New resource type details

Field	Description
Resource Name	(Required) Name of the resource to register. Enter the actual name of the server, including uppercase and lowercase letters. If the name registered in Sybase Control Center does not exactly match the server name, some monitoring functions may not work.

Field	Description
Resource Type	Select a resource type: • **IQ Multiplex** – register a Sybase IQ multiplex server. • **IQ Server** – register a Sybase IQ simplex server. Recommended. Can handle both simplex and multiplex servers.
Description	A brief description to help you identify the resource.

3. Click **Next**.
4. Specify the connection information for your resource:

Table 15. New resource connection details

Field	Description
Host Name	(Required) Host name of the Sybase IQ server.
IQ Port Number	(Required) Port number on server host.
Database	Name of the database.
Character Set	Character set to use for the connection.
Language	Language to use for the connection.

5. (Optional) Enter a user name and password that SCC can use to authenticate with this resource to retrieve its software version. The credentials are used only for this purpose, then discarded.

   If you prefer not to authenticate now, click **I do not want to supply authentication information**.

   This step enables SCC to display the correct version information for the server before the server is formally authenticated (later in the configuration process).

6. Click **Next**.

7. (Optional) Click **Add this resource to the current perspective**. You must add a resource to a perspective (not necessarily the current perspective) before you can manage or monitor it.

8. (Optional) Click **Open the resource explorer to view this new resource**. (This option is not present when the Resource Explorer is open.)
   The resource is added to the Resource Explorer even if you choose not to view it.

9. Click **Finish**.

See also
- *Configuring Sybase IQ for Administration* on page 117
- *Authenticating a Login Account for a Managed Resource* on page 126
- *Registering and Authenticating a Sybase Control Center Agent* on page 123

- *Manage a Multiplex Server* on page 258
- *Common Display Options* on page 7
- *Resources* on page 175
- *Unregistering a Resource* on page 175

Importing Resources for Batch Registration

(Optional) Import and register multiple servers from an `interfaces` or `sql.ini` file.

Prerequisites

Copy the `interfaces` or `sql.ini` file to a location on or accessible from the machine that hosts your Web browser.

Task

An `interfaces` (UNIX) or `sql.ini` file (Windows) is a list of Sybase servers and their ports; it may contain other connection information as well. The file is created during the installation of a Sybase server:

- Windows: `%SYBASE%\ini\sql.ini`
- Unix: `$SYBASE/interfaces`

For more information on `interfaces` files, see the appendix on configuration files in *Configuration Guide Open Client and Open Server 15.0 for UNIX.*

For more information on `sql.ini` files, see the chapter on network communications using sql.ini in the Adaptive Server Enterprise 15.0 *Configuration Guide for Windows.*

Note: The Import Resources wizard imports servers in batches of a single type (Adaptive Server, Sybase IQ, or Replication Server, for example). If your `interfaces` or `sql.ini` file includes resources of more than one type, you must perform this procedure for each resource type.

1. In the application menu, select **View > Open > Resource Explorer.**
2. In the Resource Explorer, select **Resources > Import.**
    The Import Resources wizard opens; **Interfaces file** is already selected.
3. Click **Next.**
    The Directory Service Connection page appears.
4. Click **Browse** and navigate to the interfaces file you want to import from.
    You cannot type in the **File name** field.
5. Click **Next.**
6. On the Import Resource Type page, select the type of server you want to import.

7. On the Resource Selection page, click to select the servers you want to import.

   Select only servers of the type you chose on the Import Resource Type page. If you import servers with incorrect types, Sybase Control Center will not be able to monitor or manage them properly.

8. Resources of your chosen type may require connection parameters in addition to those present in the file—RSSD host name and port for Replication Server, for example, or character set and language for Adaptive Server. Enter any required connection parameters.

9. Click **Next**.

10. (Optional) Click **Add these resources to the current perspective**. You must add a resource to a perspective (not necessarily the current perspective) before you can manage or monitor it.

11. Click **Next**.
    The Confirmation page displays a list of the resources you have selected.

12. Click **Finish** if you are ready to import, or click **Back** to return to the previous screens and change your selections.
    When you click **Finish**, Sybase Control Center imports and registers the resources and displays a summary page.

13. Click **Close** to finish the wizard.

The newly imported resources appear in the Resource Explorer. If you elected to add them to the current perspective, the resources also appear in the Perspective Resources view.

See also
- *Resources* on page 175
- *Unregistering a Resource* on page 175

Registering and Authenticating a Sybase Control Center Agent

Register and authenticate the Sybase Control Center agent for a managed server.

The Sybase Control Center agent runs on a managed server and enables Sybase Control Center to manage it. The SCC agent is installed automatically as part of the Sybase server.

To perform certain administrative tasks, including starting and stopping a Sybase IQ server, syncing the server, and adding secondary nodes, you must register and authenticate the server's SCC agent.

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select either:

- **IQ Servers**
- **IQ Servers > Multiplex Management > Multiplex Servers**

4. Select **Resource > Register Agent**.

5. Enter the IQ server host name and agent port (the default port is 9999) and click **OK**.

6. In the Administration Console, select the same IQ server or multiplex server and select **Resource > Authenticate Agent**.

7. Enter the Sybase Control Center agent user (the default is `uafadmin`) and password (by default, there is no password—leave the field blank).

Next

For instructions on changing the password for the SCC agent's default uafadmin account, see the topic on setting passwords in the *Sybase Control Center Installation Guide.*

See also

Viewing Sybase Control Center Agent Connection Information

View Sybase Control Center agent connection information in the server properties.

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select either:

- **IQ Servers**
- **IQ Servers > Multiplex Management > Multiplex Servers**

4. Select the server or multiplex node from the right pane and either:

- Click the arrow to the right of the name and select **Properties**, or
- From the Administration Console menu bar, select **Resource > Properties**.

5. Select **Agent** from the left pane.

Area	Description
Agent Page	**Agent registered** – Indicates if the Sybase Control Center agent is registered: true or false.
	Agent authenticated – Indicates if the Sybase Control Center agent is authenticated: true or false.
	Agent status – Status of the Sybase Control Center agent: Running, Stopped, or Unknown.
	Agent host – Name of the host machine where the Sybase Control Center agent is running.
	Agent port – Port number on the host machine where the Sybase Control Center agent is running.
	Agent user – User name for authentication of the agent. Default is *uafadmin*.
	Agent process owner – The user name that owns the agent process.
	Agent home – The home directory of the Sybase Control Center agent.
	Agent version – The version of the Sybase Control Center agent.
	SCC agent plugin version – The agent plugin version of the Sybase Control Center agent.
	IQ directory – Installation directory of the IQ server with which the Sybase Control Center agent is associated.
	IQ version – Version of the IQ server with which the Sybase Control Center agent is associated.

Creating a Perspective

Create a perspective in which you can add and manage resources.

1. From the application menu bar, select **Perspective > Create.**
2. Enter a name for your perspective. The name can contain up to 255 characters.
3. Click **OK**.

See also

Adding a Resource to a Perspective

Add one or more resources to the current perspective.

Add servers or other resources to a perspective so you can monitor and manage them along with other resources in the same perspective.

1. From the Sybase Control Center toolbar, click the **Launch Resource Explorer** icon.
2. Select the resources to add to your perspective. Use **Shift-click** or **Control-click** to select multiple resources.
3. Perform one of these actions:

   - Select **Resources > Add Resources to Perspective.**
   - Drag and drop resources from the Resource Explorer onto the Perspective Resources view. You can select and drag multiple resources.

See also
- *Removing a Resource from a Perspective* on page 176
- *Resources* on page 175

Authenticating a Login Account for a Managed Resource

Specify the login account Sybase Control Center will use when it connects to your server or agent to collect monitoring data or manage the resource.

Perform this task for each resource registered with Sybase Control Center.

Note: You can also authenticate a server during administrative tasks like creating an alert or a collection job.

1. Connect a browser to Sybase Control Center and log in.
2. If the Perspective Resources view is not open, click the **Show/Hide Perspective Resources View** icon in the toolbar.
3. In the Perspective Resources view, select your resource and select **Resource > Authenticate** from the view menu.
4. Select **Use my current SCC login** or **Specify different credentials**.
5. If you chose **Specify different credentials**, enter the login and password for Sybase Control Center to use to connect to your resource.
6. If the selected server is a Replication Server, also enter the RSSD user name and password.
7. Click **OK** to save and exit the dialog.

See also
- *Registering and Authenticating a Sybase Control Center Agent* on page 123
- *Registering a Sybase IQ Server* on page 120
- *Manage a Multiplex Server* on page 258
- *User Authorization* on page 106

Changing Update Frequency for Statistics and Charts

You can control the rate at which data on monitor screens and charts is refreshed, the amount of time covered by charts, and the multiplex nodes included in charts.

1. In the Perspective Resources view, select a resource, click the arrow, and select **Monitor Node**.

   For a multiplex resource, select **Monitor Multiplex**.

2. In the left pane of the IQ Node Level Monitor or IQ Multiplex Level Monitor view, select **Settings**.

3. For **Screen Refresh Interval**, enter the number of seconds between refreshes. You can also use the up and down arrows to select the appropriate number of seconds.

   The default is 30 seconds.

   Note: Refreshing a screen redraws it with the most recent data in the Sybase Control Center repository. This setting has no effect on the frequency of data collection.

4. For **Chart Trend Period**, enter the number of minutes of data to appear in charts.

   The minimum number of minutes is 5, and the maximum number is 999999999. The default is 30 minutes.

   Because data is added to a chart only when it is open, a chart contains data starting from when you opened it. Each refresh interval adds new data to the end of the graph. A chart trend period of 30 minutes shows the statistics trend over the last 30 minutes, even if the view has been open longer than 30 minutes.

5. For multiplex-level monitoring:

   a) For **Maximum Number of Nodes to Show in Chart**, enter the maximum number of nodes to include in any monitoring charts.

      The default is 10 nodes.

   b) Click **Select Nodes**, then choose the nodes to include in the monitoring charts. You cannot select more than the maximum number of nodes.

   c) Click **OK**.

6. Click **Apply**.

See also
- *Modifying the Data Collection Interval for a Job* on page 165

Setting Up Statistics Collection

Use the Properties view of your managed resource to create a data collection job and add a schedule to the job.

Statistics gathering consumes system resources intensively; the more collection jobs you run, the greater the burden on your server. For best performance, Sybase recommends these guidelines for scheduling data collection jobs:

* Schedule only one collection job for each collection.
* Set the collection interval to 5 minutes or more. (The default is 5 minutes.)

Data collections for a managed resource do not run until the resource is authenticated.

1. In the Perspective Resources view, select a resource, click its drop-down arrow, and select **Resource > Properties.**
2. Select **Collection Jobs**.
3. Click **Create Job**.
4. If this resource has not yet been authenticated, you see the Authentication page. Enter a user name and password that Sybase Control Center can use to log in to the resource. Click **Authenticate** to verify your credentials. Data collections can run only on an authenticated resource.
5. On the Collection Information page, select the data collection that this job will run.
6. (Optional) If you do not want SCC to save data collected for this job in the repository, unselect **Save data collected from this job**.

   If you choose not to save collection data, SCC updates any open views (the heat chart or a node monitor, for example) when the job runs. If the job runs when no views are open, the data is not captured.

   This option cannot be modified once the job is created. If you need to change it, drop the data collections and add it again.
7. Click **Next**.
8. (Optional) If you do not want to create a schedule yet, unselect **Create a schedule for this job**.
9. Specify details for the new schedule:

Field	Description
Name	A name for this schedule
Description	A description of this schedule

10. Choose to start the job **Now** or **Later**.

11. Specify the duration of this schedule. The job can run:

- **Once**
- **Repetitively** at an interval you specify

Field	Description
Repeat interval	Time period (in seconds, minutes, hours, or days) between job executions

- **Until** a stop date that you specify, at an interval you specify

Field	Description
Repeat interval	Time period (in seconds, minutes, hours, or days) between job executions
Stop date	Date and time the job should stop running

Note: Enter dates and times using your local time. Sybase Control Center converts your times for remote time zones if necessary.

12. Click **Finish**.

See also

About Statistics

Understand availability and performance statistics in Sybase Control Center.

The statistics you work with in Sybase Control Center can be divided into two types:

- Availability statistics are concerned with present conditions; they help you determine whether a resource you are monitoring (a server or an agent, for example) is running and functioning properly.
- Performance statistics are concerned with behavior of the same resources over time. They describe the flow of data through your environment. You can use performance statistics to spot trends, identify problems like resource bottlenecks, and make plans.

Sybase Control Center includes predefined key performance indicators (KPIs) for each product module; these KPIs are grouped into collections. KPIs such as server status, which serves as an availability statistic when it is fresh, have long-term value as historical performance statistics.

Availability statistics appear on the heat chart and on resource monitoring screens in each product module.

Performance statistics appear on the statistics chart and on resource monitoring screens in each product module.

Some KPIs are included in the default collection for each product module. To make other KPIs available to the heat chart, statistics chart, and resource monitoring views, you must set up collection jobs in the scheduler. See the data collections help topic for information on data collections and the KPIs contained in them.

Several configuration options affect the collection and display of data in Sybase Control Center:

*   Collection repeat interval—The frequency of data collection. Set this on the collection job in the scheduler.
*   Screen refresh interval—The period between screen refreshes. Refreshing the screen redraws it with the newest available data. Set the screen refresh interval in the product module. (May not be settable in all product modules.)
*   Chart trend period—The period over which data is displayed in historical charts. Set the trend period in the product module. (May not be settable in all product modules.)

See also
*   *Sybase IQ Data Collections* on page 130
*   *Key Performance Indicators for Sybase IQ* on page 139

Sybase IQ Data Collections

Predefined data collections you can schedule for Sybase IQ. Collected statistics appear on Sybase Control Center monitoring screens and trigger user-configured alerts.

Sybase provides alternative versions of several data collections to avoid duplication. The alternatives (which are all described in the tables below) are:

*   All Statistics Collection and All Stats w/o Availability
*   Engine Statistics Collection and Engine Stats w/o Availability
*   Connection Statistics Collection and Connection Stats w/o Availability
*   Cache Statistics Collection and Cache Stats w/o Availability
*   All Multiplex Statistics Collection and All MPX Stats w/o Availability

The "w/o Availability" collections omit statistics that are collected by the Availability Statistics Collection or the Multiplex Availability Collection, which are scheduled by default. They are otherwise identical to their counterparts.

Minimizing duplication:

*   Reduces the occurrence of duplicate alerts
*   Improves the performance of both Sybase Control Center and managed servers

A statistic is considered a duplicate only when its key performance indicator is collected more than once in the same time interval. For example, if you schedule both the All Statistics

Collection and the Availability Statistics Collection at 15-minute intervals, the data collected for the Availability Statistics Collection is duplicated—all of its KPIs are also in the All Statistics Collection. If you schedule the All Statistics Collection at 30-minute intervals and the Availability Statistics Collection at 5-minute intervals, however, the data is not duplicated.

Table 16. Data Collections for Server-Type (Simplex) Resources

Collection	Description	KPIs
All Statistics Collection	Contains all statistics, including server availability and performance statistics.	All KPIs listed in this table.
Availability Statistics Collection	Contains server availability statistics used in the heat chart. This is the default collection; it is automatically scheduled when you authenticate a Sybase IQ simplex server. **Important:** Sybase strongly recommends that you leave this collection running for each monitored simplex server. Suspend it only if you schedule another collection that includes all of this collection's KPIs, such as All Statistics Collection.	• Server Status • SCC Agent Status • Total CPU Usage • System CPU Usage • User CPU Usage • IQ Memory Allocated • Catalog Cache in Use • Main Cache in Use • Temporary Cache in Use • Available Connections
All Stats w/o Availability	Contains all statistics except those collected by the Availability Statistics Collection. Schedule this collection to collect all statistics without duplication when you also use the Availability Statistics Collection, which is scheduled by default.	All KPIs listed in this table with the exception of those included in the Availability Statistics Collection, above.

Collection	Description	KPIs
Engine Statistics Collection	Contains performance statistics for the Sybase IQ engine.	• Total CPU Usage • System CPU Usage • User CPU Usage • IQ Memory Allocated • IQ Maximum Memory Allocated • IQ Threads in Use • IQ Threads Available
Engine Stats w/o Availability	Contains all statistics in the Engine Statistics Collection except those that are also collected by the Availability Statistics Collection. To avoid duplication with the Availability Statistics Collection, schedule this collection rather than the Engine Statistics Collection.	• IQ Maximum Memory Allocated • IQ Threads in Use • IQ Threads Available
Connection Statistics Collection	Contains connection performance statistics.	• Active Connections • Available Connections • Active User Connections • Active INC Incoming Connections • Active INC Outgoing Connections • User Connections Per Minute • User Disconnections Per Minute
Connection Stats w/o Availability	Contains all statistics in the Connection Statistics Collection except those that are also collected by the Availability Statistics Collection. To avoid duplication with the Availability Statistics Collection, schedule this collection rather than the Connection Statistics Collection.	• Active Connections • Active User Connections • Active INC Incoming Connections • Active INC Outgoing Connections • User Connections Per Minute • User Disconnections Per Minute

Collection	Description	KPIs
Transaction Statistics Collection	Contains transaction performance statistics.	• Active Transactions • Active User Transactions • Active INC Transactions • Active Load Table Statements
DBSpace Statistics Collection	Contains statistics for each dbspace.	• Percentage of Available DBSpace • DBSpace Size in Use
DBSpace File Statistics Collection	Contains statistics for each dbspace file.	• Percentage of Available DBSpace File • DBSpace File Size in Use
Store I/O Statistics Collection	Contains performance statistics for store reads and writes.	• Catalog Store Disk Reads Per Second • Catalog Store Disk Writes Per Second • Temporary Store Disk Reads Per Second • Temporary Store Disk Writes Per Second • Main Store Disk Reads Per Second • Main Store Disk Writes Per Second

Collection	Description	KPIs
Cache Statistics Collection	Contains performance statistics for caches.	• Catalog Cache Hits Per Second • Catalog Cache Reads Per Second • Catalog Cache Size • Catalog Cache in Use • Catalog Cache Pinned • Catalog Cache Pinned Percent • Catalog Cache Dirty Pages Percent • Temporary Cache Hits Per Second • Temporary Cache Reads Per Second • Temporary Cache Size • Temporary Cache in Use • Temporary Cache Pinned • Temporary Cache Pinned Percent • Temporary Cache Dirty Pages Percent • Main Cache Hits Per Second • Main Cache Reads Per Second • Main Cache Size • Main Cache in Use • Main Cache Pinned • Main Cache Pinned Percent • Main Cache Dirty Pages Percent

Collection	Description	KPIs
Cache Stats w/o Availability	Contains all statistics in the Cache Statistics Collection except those that are also collected by the Availability Statistics Collection. To avoid duplication with the Availability Statistics Collection, schedule this collection rather than the Cache Statistics Collection.	• Catalog Cache Hits Per Second • Catalog Cache Reads Per Second • Catalog Cache Size • Catalog Cache Pinned • Catalog Cache Pinned Percent • Catalog Cache Dirty Pages Percent • Temporary Cache Hits Per Second • Temporary Cache Reads Per Second • Temporary Cache Size • Temporary Cache Pinned • Temporary Cache Pinned Percent • Temporary Cache Dirty Pages Percent • Main Cache Hits Per Second • Main Cache Reads Per Second • Main Cache Size • Main Cache Pinned • Main Cache Pinned Percent • Main Cache Dirty Pages Percent
Op/Req Statistics Collection	Contains performance statistics related to operations and requests.	• IQ Waiting Operations • IQ Active Operations • Requests Per Second • Unscheduled Requests
Network Statistics Collection	Contains network-related performance statistics.	• Bytes Received Per Second • Bytes Received Uncompressed Per Second • Bytes Sent Per Second • Bytes Sent Uncompressed Per Second • Available Communication Buffers • Total Communication Buffers

Table 17. Data Collections for Multiplex-Type Resources

Collection	Description	KPIs
All Multiplex Statistics Collection	Contains all multiplex-related statistics and all statistics for each node. When scheduled, all node statistics are collected.	• Multiplex Status The following KPIs for each node: • Server Status • SCC Agent Status • Multiplex Role • Multiplex Status • Multiplex Failover Node • INC Status • Total CPU Usage • System CPU Usage • User CPU Usage • IQ Memory Allocated • IQ Maximum Memory Allocated • IQ Threads in Use • IQ Threads Available • Active Connections • Available Connections • Active User Connections • Active INC Incoming Connections • Active INC Outgoing Connections • User Connections Per Minute • User Disconnections Per Minute • Active Transactions • Active User Transactions • Active INC Transactions • Active Load Table Statements • IQ Waiting Operations • IQ Active Operations • Requests Per Second • Unscheduled Requests • Bytes Received Per Second • Bytes Received Uncompressed Per Second • Bytes Sent Per Second

Collection	Description	KPIs
		• Bytes Sent Uncompressed Per Second • Available Communication Buffers • Total Communication Buffers The following KPIs for each dbspace: • Percentage of Available DBSpace • DBSpace Size in Use The following KPIs for each dbspace file: • Percentage of Available DBSpace File • DBSpace File Size in Use The following KPIs for stores: • Catalog Store Disk Reads Per Second • Catalog Store Disk Writes Per Second • Temporary Store Disk Reads Per Second • Temporary Store Disk Writes Per Second • Main Store Disk Reads Per Second • Main Store Disk Writes Per Second The following KPIs for caches: • Catalog Cache Hits Per Second • Catalog Cache Reads Per Second • Catalog Cache Size • Catalog Cache in Use • Catalog Cache Pinned • Catalog Cache Pinned Percent • Catalog Cache Dirty Pages Percent • Temporary Cache Hits Per Second • Temporary Cache Reads Per Second • Temporary Cache Size • Temporary Cache in Use

Collection	Description	KPIs
		• Temporary Cache Pinned • Temporary Cache Pinned Percent • Temporary Cache Dirty Pages Percent • Main Cache Hits Per Second • Main Cache Reads Per Second • Main Cache Size • Main Cache in Use • Main Cache Pinned • Main Cache Pinned Percent • Main Cache Dirty Pages Percent
Multiplex Availability Collection	Contains multiplex status, availability statistics, and multiplex role/status/failover/INC status of each node. When scheduled, all node availability statistics are collected. This is the default collection; it is automatically created when you authenticate a multiplex-type resource. **Important:** Sybase strongly recommends that you leave this collection running for each monitored multiplex server. Suspend it only if you schedule another collection that includes all of this collection's KPIs, such as All Multiplex Statistics Collection.	• Multiplex status The following KPIs of each node: • Server Status • SCC Agent Status • Total CPU Usage • System CPU Usage • User CPU Usage • IQ Memory Allocated • Catalog Cache in Use • Main Cache in Use • Temporary Cache in Use • Available Connections • Multiplex Role • Multiplex Status • Multiplex Failover Node • INC Status

Collection	Description	KPIs
All MPX Stats w/o Availability	Contains all statistics except those collected by the Multiplex Availability Collection. Schedule this collection to collect all statistics without duplication when you also use the Multiplex Availability Collection, which is scheduled by default.	All KPIs listed in this table with the exception of those included in the Multiplex Availability Collection, above.

See also
- *About Statistics* on page 129
- *Key Performance Indicators for Sybase IQ* on page 139
- *Creating an Alert* on page 146

Key Performance Indicators for Sybase IQ

Key performance indicators (KPIs) provide the statistics that appear on Sybase IQ screens and charts in Sybase Control Center.

Each Sybase IQ data collection includes a subset of the KPIs listed here.

Table 18. Sybase IQ Availability Statistics

KPI	Description
Server Status (Resource State)	Status of the Sybase IQ server. Valid states are: • UNKNOWN (0) • STOPPED (1) • PENDING (2) • RUNNING (3) • WARNING (4) • ERROR (5)
Total CPU Usage	Sybase IQ server total CPU usage percentage, including both system and user usage.
CPU System Usage	Sybase IQ server system CPU usage percentile.
CPU User Usage	Sybase IQ server user CPU usage percentile.
IQ Memory Allocated	Memory allocated by the Sybase IQ server, in megabytes.

KPI	Description
Catalog Cache in Use	Percentage of catalog cache in use.
Temporary Cache in Use	Percentage of temporary cache in use.
Main Cache in Use	Percentage of main cache in use.
Available Connections	Number of concurrent connections available.

Table 19. Sybase Control Center Agent Availability Statistics

KPI	Description
SCC Agent Status (Resource State)	Status of the registered and authenticated SCC agent that enables SCC to manage the Sybase IQ server. Valid states are: • UNKNOWN (0) • STOPPED (1) • PENDING (2) • RUNNING(3) • WARNING (4) • ERROR (5)

Table 20. Multiplex Availability Statistics

KPI	Description
Multiplex Status (Resource State)	Status of the multiplex. Valid states are: • UNKNOWN (0) • STOPPED (1) (A multiplex is considered stopped if there is no running node.) • PENDING (2) • RUNNING (3) (A multiplex is considered to be running if at least one node is running.) • WARNING (4) • ERROR (5)

Table 21. Multiplex Node Properties

KPI	Description
Multiplex Role (Resource Secondary State)	Indicates the multiplex role of the node. Valid states are: • COORDINATOR (0) • WRITER (1) • READER (2) • SINGLE_SERVER (3)
Multiplex Status	Indicates the multiplex status of the node. Valid values are: • INCLUDED (0) • EXCLUDED (1)
Multiplex Failover Node	Indicates whether the node is a designated failover node. Valid values are: • FALSE (0) • TRUE (1)

Table 22. Multiplex Link Availability Statistics

KPI	Description
INC Status (Resource State)	Status of inter-node communication between a secondary node and the coordinator. Valid states are: • UNKNOWN (0) • STOPPED (1) • PENDING (2) • ACTIVE (3) • WARNING (4) • ERROR (5)

Table 23. Overview Statistics

KPI	Description
Server Status	Server state or status. Valid values are: • UNKNOWN (0) • STOPPED (1) • PENDING (2) • RUNNING (3) • WARNING (4) • ERROR (5)
Total CPU Usage	Sybase IQ server total CPU usage percentage, including both system and user usage.
IQ Memory Allocated	Memory allocated by the Sybase IQ server, in megabytes.
Number of Active Connections	Total number of active connections including user and inter-node-communication connections.

Table 24. Engine Statistics

KPI	Description
Total CPU Usage	Sybase IQ server total CPU usage percentage, including both system and user usage.
CPU System Usage	Sybase IQ server system CPU usage percentile.
CPU User Usage	Sybase IQ server user CPU usage percentile.
IQ Memory Allocated	Memory allocated by the Sybase IQ server, in megabytes.
IQ Maximum Memory Allocated	Maximum memory allocated by the Sybase IQ server, in megabytes.
IQ Threads in Use	Number of threads used by the Sybase IQ server.
IQ Threads Available	Number of threads available in the Sybase IQ server.

Table 25. Connection Statistics

KPI	Description
Active Connections	Total number of active connections including user and inter-node-communication connections.
Available Connections	Number of concurrent connections available.
Active User Connections	Number of active user connections.
Active INC Incoming Connections	Number of inter-node-communication incoming connections.
Active INC Outgoing Connections	Number of inter-node-communication outgoing connections.
User Connections Per Minute	Number of user connections per minute.
User Disconnections Per Minute	Number of user disconnections per minute.

Table 26. Transaction Statistics

KPI	Description
Active Transactions	Total number of active transactions, including user and INC transactions.
Active User Transactions	Number of active user transactions.
Active INC Transactions	Number of active inter-node communication transactions.
Active Load Table Statements	Number of active **LOAD TABLE** statements.

Table 27. DBSpace and DBSpace File Statistics

KPI	Description
Percentage of Available DBSpace	Percentage of the dbspace that is available.
DBSpace Size in Use	Dbspace size in use (MB).
Percentage of Available DBSpace File	Percentage of space available in the dbspace file.
DBSpace File Size in Use	Dbspace file size in use (MB).

Table 28. Store Input and Output Statistics

KPI	Description
Catalog Store Disk Reads Per Second	Number of kilobytes per second that have been read from the catalog store.
Catalog Store Disk Writes Per Second	Number of kilobytes per second that have been written to the catalog store.
Temporary Store Disk Reads Per Second	Number of kilobytes per second that have been read from the temporary store.
Temporary Store Disk Writes Per Second	Number of kilobytes per second that have been written to the temporary store.
Main Store Disk Reads Per Second	Number of kilobytes per second that have been read from the main store.
Main Store Disk Writes Per Second	Number of kilobytes per second that have been written to the main store.

Table 29. Cache Statistics

KPI	Description
Catalog Cache Hits Per Second	Number of catalog cache hits per second.
Catalog Cache Reads Per Second	Number of catalog cache page lookups per second.
Catalog Cache Size	Current catalog cache size, in megabytes.
Catalog Cache in Use	Percentage of catalog cache in use.
Catalog Cache Pinned	Number of pinned catalog cache pages.
Catalog Cache Pinned Percent	Percentage of catalog cache pinned.
Catalog Cache Dirty Pages Percent	Percentage of catalog cache dirty pages.
Temporary Cache Hits Per Second	Number of temporary cache hits per second.
Temporary Cache Reads Per Second	Number of temporary cache page lookups per second.
Temporary Cache Size	Current temporary cache size, in megabytes.
Temporary Cache in Use	Percentage of temporary cache in use.
Temporary Cache Pinned	Number of pinned temporary cache pages.

KPI	Description
Temporary Cache Pinned Percent	Percentage of temporary cache pinned.
Temporary Cache Dirty Pages Percent	Percentage of temporary cache dirty pages.
Main Cache Hits Per Second	Number of main cache hits per second.
Main Cache Reads Per Second	Number of main cache page lookups per second.
Main Cache Size	Current main cache size, in megabytes.
Main Cache in Use	Percentage of main cache in use.
Main Cache Pinned	Number of pinned main cache pages.
Main Cache Pinned Percent	Percentage of main cache pinned.
Main Cache Dirty Pages Percent	Percentage of main cache dirty pages.

Table 30. Operations and Requests Statistics

KPI	Description
IQ Waiting Operations	Number of Sybase IQ operations waiting for the resource governor.
IQ Active Operations	Number of active concurrent operations admitted by the Sybase IQ resource governor.
Requests Per Second	Number of times per second the server has been entered to allow it to handle a new request or continue processing an existing request.
Unscheduled Requests	Number of requests that are currently queued up waiting for an available server thread.

Table 31. Network Statistics

KPI	Description
Bytes Received Per Second	Number of bytes per second received during client/server communications.
Bytes Received Uncompressed Per Second	Number of bytes per second that would have been received during client/server communications if compression was disabled.
Bytes Sent Per Second	Number of bytes per second sent during client/server communications.

KPI	Description
Bytes Sent Uncompressed Per Second	Number of bytes per second that would have been sent during client/server communications if compression was disabled.
Available Communication Buffers	Number of available network communication buffers.
Total Communication Buffers	Total number of network communication buffers.

See also
- *About Statistics* on page 129
- *Sybase IQ Data Collections* on page 130
- *Graphing Performance Counters* on page 160
- *Creating an Alert* on page 146

Creating an Alert

Use the Add Alert wizard to create an alert instance for your resource.

Prerequisites

- You must have administrative privileges (sccAdminRole) to perform this task.
- Specify an e-mail server for Sybase Control Center to use for alerts. You cannot create e-mail subscriptions to alerts without an e-mail server.
- Schedule data collections. Alerts for each product module are based on one or more data collections. If the correct collection or collections are not scheduled to run, the alert system cannot function and no alerts are generated. See the data collections topic for your product module for information on which collections you need to schedule to enable alerts.
- (Optional) If you want this alert to trigger the execution of a shell script, copy the script to a location on or accessible from the machine that hosts your Sybase Control Center server. Set permissions to make the script executable.

Warning! Use caution in writing scripts. A poorly designed script can cause a blocking situation, creating a deadlock in your Sybase Control Center server.

Task

1. In the Perspective Resources view, click the server or other resource and select **Resource > Properties** in the view's menu bar.
2. Select **Alerts** in the left pane and click **Add**.

The Add Alert Wizard opens. If the selected resource supports child alerts, the wizard opens to the Resource page. If the resource does not support child alerts, the wizard opens to the Type page—in that case, skip to step *5*.

3. On the Resource page of the wizard, select the object on which to set the alert. Expand the folder representing the server or agent to select lower-level child objects.

4. Click **Next**.

5. On the Type page, select the alert type and click **Next**.

   For this step and the next one, see the topic on key performance indicators for information on what this alert monitors and how it is triggered. (Each alert is based on a KPI.)

6. Based on the type of alert you selected, do one of the following:

   • For a state-based alert – select a severity level for each alert state.

   > **Note:** You can associate only one severity level with each state.

   • For a threshold-based alert – review and if necessary adjust the range of values that defines each severity.

7. Click **Next**.

8. (Optional) Enter the storm suppression period. Storm suppression blocks redundant alert notifications and script executions resulting from the same condition for the specified period of time. Enter this value in seconds, minutes, or hours and click **Next**.

9. (Optional) To configure this alert to trigger the execution of a script:

   a) **Alert Severity** specifies the severity level that triggers the script. Select **Critical**, **Warning**, or both.

      Critical is typically more serious than Warning.

   b) Browse to the location of the script.

      > **Note:** In UNIX, make sure the script is executable. You cannot select a script unless it has execute permission.

   c) If the script requires parameter values, click **Select Parameters** to enter them in the **Execution Parameters** box.

      You can include a number of predefined substitution parameters, which are replaced by values from the alert. The parameter values are passed on the command line to the script. See the example (below) and the substitution parameters topic (linked below) for more information.

      > **Note:** When you test a script, Sybase Control Center supplies test values for the **%Severity%** and **%Source_Application%** parameters ("Testing" and "TestScriptExecution," respectively). Any test values you supply for these parameters are discarded. This prevents the test results from being confused with real script results after testing and in the SCC repository.

   d) (Optional) Click **Test** to perform a test execution of your script.

If your script takes parameters, the test may fail if parameter values are missing or incorrect.

e) Click **Next**.

If the selected resource has sibling resources (databases or devices of the same type, for example) that support this alert type, you see the Duplicates page. If the selected resource has no identical siblings, you see the Subscription page.

10. (Optional) On the Duplicates page, select any resources that should use this alert definition as a template for their own alerts. Click the box at the top of the list to select all the resources listed. Then click **Next**.

This step saves time when you need to configure similar alerts for several resources of the same type.

11. (Optional) On the Subscription page, specify e-mail addresses if you want this alert to issue e-mail notifications when it fires.

The e-mail addresses default to the address in your user profile, but you can override the defaults.

For both critical and warning alerts:

Table 32. Alert subscription details

Option	Description
E-mail	To send an e-mail notification when this alert fires, click the **E-mail Message** box and enter the e-mail address of one user or list.
Escalation E-mail	To escalate this alert (by sending another e-mail notification if this alert has not been responded to after a specified period of time), click the **Escalation E-mail** box and enter the e-mail address of one user or list. You cannot enter an escalation address unless you enter an address for primary notification first.
Time Period	Specify how long to wait, following the initial alert notification, before Sybase Control Center sends an e-mail notification to the escalation address. (The same notification is sent again to the original notification address.) Select a time unit (hours, minutes, or seconds) and enter a number.

12. Click **Finish**.

If you are creating duplicate or child alerts, the **Cancel** button is activated; click it to interrupt the creation of further alerts. (The primary alert, at a minimum, is always created before the operation can be cancelled.) If you do not want to keep the duplicate or child alerts (if any) created before you cancelled the operation, drop them manually.

Note: Click **Cancel** to stop the creation of duplicate alerts.

Examples: Alert-triggered scripts

This sample script is a Windows `.bat` file. It outputs the parameter values you pass to it to a text file. Windows batch files support only nine arguments. (Arg0, the name of the script, is not counted.)

```
@echo off
@echo. >> stest.txt
@echo %date% %time% >> stest.txt
@echo arg0: %0 >> stest.txt
@echo arg1: %1 >> stest.txt
@echo arg2: %2 >> stest.txt
@echo arg3: %3 >> stest.txt
@echo arg4: %4 >> stest.txt
@echo arg5: %5 >> stest.txt
@echo arg6: %6 >> stest.txt
@echo arg7: %7 >> stest.txt
@echo arg8: %8 >> stest.txt
@echo arg9: %9 >> stest.txt
@echo. >> stest.txt
```

This is a sample execution parameter string for the script above:

```
Time:%Time%
Severity:%Severity%
Resource:%Resource%
Server:%Top_resource%
KPI:%KPI%
State:%Current_state%
URL:%SCC_URL%
```

The script's ouput might look like this:

```
Tue 12/15/2009 14:54:45.58
arg0: C:\project\sccmain\script-test.bat
arg1: Time:"Mon Dec 21 21:30:04 2009"
arg2: Severity:CRITICAL
arg3: Resource:"SCC Tester 1"
arg4: Server:"SCC Tester 1"
arg5: KPI:kpi_scc_mostate_primary
arg6: State:ERROR
arg7: HYPERLINK "http://ik-scc.sybase.com:8282/scc"URL:http://ik-
scc.sybase.com:8282/scc
arg8:
arg9:
```

This is a UNIX script. It also outputs the parameter values you pass to it to a text file.

```
#!bin/sh
outfile=c:/testing/latest/scriptTest.out
echo> $outfile
echo `date` >> $outfile
count=1
while [ "$1" ]
do
  echo arg$count: $1 >> $outfile
  shift
  count=`expr $count + 1`
```

```
done
echo --- DONE --- >> $outfile
```

See also

- *Setting Up Statistics Collection* on page 128
- *Optional Configuration Steps* on page 157
- *Sybase IQ Data Collections* on page 130
- *Key Performance Indicators for Sybase IQ* on page 139
- *Assigning a Role to a Login or a Group* on page 106
- *Configuring the E-mail Server* on page 104
- *Alerts* on page 167
- *Testing an Alert-Triggered Script* on page 170

Sybase IQ Alerts, Collections, and KPIs

Lists and describes alerts you can use for Sybase IQ.

The alerts are based on the same key performance indicators (KPIs) that are collected for the Sybase IQ node level monitor displays, and for the Statistics Chart.

Note: To configure alerts for a multiplex, register each node for which you want to set alerts and add the node to a perspective. Then set up alerts for each node using the Resource Properties view (select a server in the Perspective Resources view and select **Resource > Properties > Alerts**). Do not create alerts in the Resource Properties view for the multiplex.

Alert	Description	Alert Type
Active INC incoming connections	Number of internode communication (INC) incoming connections.	Threshold
Active INC outgoing connections	Number of INC outgoing connections.	Threshold
Active INC transactions	Number of active INC transactions.	Threshold
Active load table statements	Number of active load table statements.	Threshold
Active transactions	Total number of active transactions, including user and INC transactions.	Threshold
Active user connections	Number of active user connections.	Threshold
Active user transactions	Number of active user transactions.	Threshold

Alert	Description	Alert Type
Available communication buffers	Number of available network communication buffers.	Threshold
Bytes received	Number of bytes per second received during client/server communications.	Threshold
Bytes received uncompressed	Number of bytes per second that would have been received during client/server communications if compression was disabled.	Threshold
Bytes sent	Number of bytes per second sent during client/server communications.	Threshold
Bytes sent uncompressed	Number of bytes per second that would have been sent during client/server communications if compression was disabled.	Threshold
Catalog cache dirty pages percent	Percentage of pages in the catalog cache where data has been modified and stored in the buffer cache and has not yet been written to disk.	Threshold
Catalog cache hits	Number of catalog cache hits per second.	Threshold
Catalog cache in use percent	Percentage of the catalog cache in use.	Threshold
Catalog cache pinned	Number of pinned catalog cache pages.	Threshold
Catalog cache pinned percent	Percentage of the catalog cache pinned.	Threshold
Catalog cache reads	Number of catalog cache page lookups per second.	Threshold
Catalog store disk reads	Number of kilobytes per second that have been read from the catalog store.	Threshold
Catalog store disk writes	Number of kilobytes per second that have been written to the catalog store.	Threshold
CPU total usage	Percentage of the Sybase IQ server CPU in use, including both system and user usage.	Threshold
Dbspace size in use	Dbspace size in use in MB.	Threshold
Dbspace file size in use	Dbspace file size in use in MB.	Threshold

Alert	Description	Alert Type
Dbspace percent available	Percentage of Dbspace size available.	Threshold
Dbspace file percent available	Percentage of Dbspace file size available.	Threshold
IQ active operations	Number of active concurrent operations admitted by the Sybase IQ resource governor.	Threshold
SCC Agent Availability	Indicates Sybase Control Center agent availability state. Possible values: UNKNOWN(0), STOPPED(1), and RUNNING(3).	State
IQ threads in use	Number of threads used by the Sybase IQ server.	Threshold
Main cache hits	Number of main cache hits per second.	Threshold
Main cache in use percent	Percentage of the main cache in use.	Threshold
Main cache pinned	Number of pinned main cache pages.	Threshold
Main cache pinned percent	Percentage of the main cache pinned.	Threshold
Main cache reads	Number of main cache page lookups per second.	Threshold
Main cache size	Current main cache size in megabytes.	Threshold
Main store disk reads	Number of kilobytes per second that have been read from the main store.	Threshold
Main store disk writes	Number of kilobytes per second that have been written to the main store.	Threshold
Number of active connections	Total number of active connections, including user and internode communication connections.	Threshold
Number of connections available	Number of concurrent connections available.	Threshold
Number of IQ threads available	Number of threads available in the Sybase IQ server.	Threshold
Requests per second	Number of times per second the server has been accessed to handle a new request or continue processing an existing request.	Threshold

Alert	Description	Alert Type
Server availability	Status of the Sybase IQ server.	State
Temp cache dirty pages percent	Percentage of pages in the temporary caches where data has been modified and stored in the buffer cache and has not yet been written to disk.	Threshold
Temp cache hits	Number of temporary cache hits per second.	Threshold
Temp cache in use percent	Percentage of the temporary cache in use.	Threshold
Temp cache pinned	Number of pinned temporary cache pages.	Threshold
Temp cache pinned percent	Percentage of the temporary cache pinned.	Threshold
Temp cache reads	Number of temporary cache page lookups per second.	Threshold
Temp cache size	Current temporary cache size in megabytes.	Threshold
Temp store disk reads	Number of kilobytes per second that have been read from the temporary store.	Threshold
Temp store disk writes	Number of kilobytes per second that have been written to the temporary store.	Threshold
Total communication buffers	Total number of network communication buffers.	Threshold
Unscheduled requests	Number of requests that are currently queued up waiting for an available server thread.	Threshold
User connections per minute	Number of user connections per minute.	Threshold
User disconnections per minute	Number of user disconnections per minute.	Threshold

See also
- *Alert Types and Severities for Sybase IQ* on page 153
- *Alert-Triggered Scripts* on page 154
- *Substitution Parameters for Scripts* on page 155

Alert Types and Severities for Sybase IQ

Learn about the properties that define and control alerts.

An alert's type determines what causes it to fire.

Table 33. Alert types

Type	Description
State	A state alert fires when the metric on which it is based changes to a particular state. The possible states are running, pending, stopped, warning, error, and unknown.
Threshold	A threshold alert fires when the metric on which it is based passes a preset level.

Alert severities control when an alert is issued. You can configure the states or threshold values for each alert.

Table 34. Alert severities

Severity	Description
Normal	No alert is issued.
Warning	A problem has given cause for concern. An alert is issued; you can subscribe to alerts that fire at the Warning level.
Critical	A serious problem exists. An alert is issued; you can subscribe to alerts that fire at the Critical level.

See also
- *Sybase IQ Alerts, Collections, and KPIs* on page 150
- *Alert-Triggered Scripts* on page 154
- *Substitution Parameters for Scripts* on page 155

Alert-Triggered Scripts

You can write a shell script and configure an alert to execute the script.

Use scripts to help manage and respond to alerts. A script might trigger a visual alarm in a control center or send an e-mail message about the alert to a list of addresses (a way of supplementing the alert subscription feature, which accepts a single address).

When you configure an alert to execute a script, you:

- Specify the states or thresholds that set off the alert
- Specify the severity level that triggers execution of the script
- Supply an execution parameter string to be passed to the script

Scripts are executed under the login account used to start Sybase Control Center. Make sure that account has permissions that allow it to perform the actions contained in all scripts.

When a script executes, Sybase Control Center logs the start time, end time, and status and exit codes to the alert services log. Log location:

- In a standard installation:
  ```
  SCC-3_2\log\alert-server.log
  ```
- In a shared disk installation:
  ```
  SCC-3_2\instances\<instance-name>\log\alert-server.log
  ```

Warning! Use caution in writing scripts. A poorly designed script can cause a blocking situation, creating a deadlock in your Sybase Control Center server.

See also
- *Sybase IQ Alerts, Collections, and KPIs* on page 150
- *Alert Types and Severities for Sybase IQ* on page 153
- *Substitution Parameters for Scripts* on page 155
- *Testing an Alert-Triggered Script* on page 170
- *Alerts* on page 167

Substitution Parameters for Scripts

In the execution parameter string you supply to be passed to your shell script, you can include substitution parameters that are replaced at execution time with values from the alert that triggers the script.

Substitution parameters are available for both state-based and threshold-based alerts.

Table 35. Substitution Parameters for State-Based Alerts

Parameter	Description
%Alert%	A three-part name supplied by the alert system. The parts are the name of this alert, the name of the resource, and the name of the key performance indicator (KPI) on which this alert is based.
%Current_state%	The current state of the resource on which this alert is configured.
%KPI%	The name of the KPI on which this alert is based.
%Resource%	The name of the resource with which this alert is associated.
%SCC_URL%	A link to Sybase Control Center, where more information about the alert may be available.
%Severity%	The severity of this alert: critical or warning.
%Source_application%	The SCC product module that generated this alert.
%Time%	The date and time at which the alert fired, in this format: Tue Sep 15 10:10:51 2009

Parameter	Description
%Server%	The name of the alerted resource's top-level parent resource—usually the server. This is valuable when the alerted resource is a component of a larger system (a database in a server, for example). If the alerted resource has no parent, %Server% and %Resource% have the same value.

Table 36. Substitution Parameters for Threshold-Based Alerts

Parameter	Description
%Alert%	A three-part name supplied by the alert system. The parts are the name of this alert, the name of the resource, and the name of the key performance indicator (KPI) on which this alert is based.
%Datapoint%	The current value, on the alerted resource, of the KPI on which this alert is based.
%KPI%	The name of the KPI on which this alert is based.
%Resource%	The name of the resource with which this alert is associated.
%SCC_URL%	A link to Sybase Control Center, where more information about the alert may be available.
%Severity%	The severity of this alert: critical or warning. (Critical is more serious.)
%Source_application%	The SCC product module that generated this alert.
%Threshold%	The threshold value at which this alert fires.
%Time%	The date and time at which the alert fired, in this format: Tue Sep 15 10:10:51 2009
%Server%	The name of the alerted resource's top-level parent resource. This is valuable when the alerted resource is a component of a larger system (a database in a server, for example). If the alerted resource has no parent, %Server% and %Resource% have the same value.

See also

Optional Configuration Steps

Perform additional configuration, including user authorization, alerts, data collection scheduling, backups, and setting purging options for the repository.

Table 37. Configuration areas

Configuration area	Description	Topic
User authorization	Set up groups of users or assign roles. Make sure there are users with administrative privileges (sccAdminRole).	*User Authorization* on page 106
Authentication	Add authentication modules to allow Windows, UNIX, and LDAP users to log in to Sybase Control Center.	*Setting up Security* on page 87
Alerts	Modify alert thresholds and subscriptions and delete alerts.	*Alerts* on page 167
Data collection	Modify collection intervals and schedules, suspend and resume the schedule, and delete collection jobs.	*Job Scheduling* on page 162
Resources	Unregister resources, add them to perspectives, or remove them.	*Resources* on page 175
Perspectives	Create, remove, and rename perspectives.	*Perspectives* on page 177
Instances	Enable or disable shared-disk mode and deploy, remove, refresh, or convert SCC agent or server instances running from a shared disk.	*Instances* on page 180
Repository	Set purging options and schedule backups of the repository database.	*Repository* on page 189

See also

* *Creating an Alert* on page 146

Configure

Manage and Monitor

Sybase Control Center allows you to manage and monitor resource availability status, view performance statistics, and provides various log information for system administrators to debug application errors.

Heat Chart

The heat chart displays status and availability statistics for managed resources in the current perspective.

The heat chart displays the state of resources in your perspective—whether the resources are running, suspended, or down. In addition, the heat chart lists the type of each resource and provides statistical data, including the start time of the last data collection.

In the Perspective Heat Chart view, you can filter the resources that you want to see and search and sort the results by column. You can also select a resource and pull down its context menu to see monitoring and administrative options that vary based on the resource type.

Heat chart data is collected directly from managed servers, tagged with the date and time when it was collected, and stored in the Sybase Control Center repository.

Displaying Resource Availability

Use the heat chart to view availability information on the servers in the current perspective.

1. From the application menu bar, select **View > Open > Heat Chart.**
2. (Optional) To display tools for filtering (narrowing the list of resources in the heat chart) or changing the columns, select **View > Filter** from the Perspective Heat Chart menu bar. The Filter and Column tools appear in the left pane.
3. (Optional) To use filtering, select **View > Filter** from the view's menu bar and enter a search term in the **Filter string** field.

   The search term can be any string that appears in the tabular portion of the heat chart, such as the name, or part of the name, of a server or a resource type (ASE Server, for example).
4. (Optional) Select a filtering setting:

   - **Match case** – search for resources whose displayed data includes the search term, including uppercase and lowercase letters; or
   - **Exact match** – search for resources whose displayed data includes an item identical to the search term.
5. (Optional) Select a column from the **Filter on** list to restrict your search to that column.

---

6. (Optional) Click **Columns** to customize your heat chart.

7. (Optional) Unselect any column that should not appear in your heat chart.

8. (Optional) Click the sorting arrow in the column headers to sort the column values in either ascending or descending order.

9. (Optional) Click the resource's row and pull down the menu to the right of the resource name to view options for the selected resource.

10. (Optional) To resize the Filter and Columns tools pane, move your mouse over the border between the tools pane and the resource table. When the mouse cursor changes to a resize icon, click and drag the border to the left or the right.

11. (Optional) To hide the Filter and Columns tools, unselect **View > Filter.**

Historical Performance Monitoring

Monitor performance data to determine whether your environment is working efficiently.

Obtain detailed information about the status of the resources in your environment. You can create performance graphs that illustrate resource performance over a specified period of time.

Graphing Performance Counters

To show performance trends, generate a graph for any set of performance counters.

Prerequisites

Verify that statistical data to be graphed has been collected. To verify data collection, go to the Collection Jobs page of the Resource Properties view and check the History tab for a collection job. You can also look at the resource monitor: if data appears there, data is being collected.

Task

Tip: Data collections start running when a resource is authenticated. A recently authenticated resource might not have accumulated enough data to make a useful graph.

1. In the Perspective Resources view, click a resource and select **Resource > Launch Statistics Chart** in the view menu bar.

2. Expand the folders in the Statistics tab and select the key performance indicator (KPI) you want to graph.

3. Click **Graph Statistic** or drag the KPI onto the Chart tab.
   The Chart tab displays the graphed data, while the KPI with its corresponding value and the date and time it was collected appear in the Data tab.

4. (Optional) Repeat to add KPIs to the graph.

5. (Optional) Use the slider at the bottom of the Chart tab to control the amount of time covered by the graph, ranging from a minute to a year.

6. (Optional) Use <<, <, >, and >> to move the displayed graph to an earlier or later time, depending on how the slider is set.

> **Tip:** The statistics chart displays data covering a fixed period of time, and that period does not change automatically. If you are viewing the most recent statistics and want to keep the graph current, adjust the displayed time period as new statistics are collected.

7. (Optional) You can click the date/time labels that appear above the slider. Use these to change the start and end time and the chart time span.

8. (Optional) Click **Clear Graph** to remove all the graphed statistics and start anew.

> **Note:** You can graph a maximum of five statistics with no more than two distinct units of measure. By default, only 24 hours of statistics are available; change the repository purge options to save statistics for a longer period.

See also
- *Setting Up Statistics Collection* on page 128
- *Key Performance Indicators for Sybase IQ* on page 139
- *Configuring Repository Purging* on page 194

Manage Sybase Control Center

Manage Sybase Control Center for Sybase IQ.

Administration Console

Use the Administration Console to browse and manage the selected resources in a perspective.

Browsing and Managing Resources
Create new resources or browse and manage existing resources.

Prerequisites
If you want to view or manage existing resources, register at least one resource and add it to a perspective.

Task

The Administration Console enables you to view and manage both servers and resources below the server level, such as processes, databases, and devices.

1. Launch the Administration Console.

- To populate the Administration Console with information on one or more resources: in the Perspective Resources view, select the resources and select **Resource > Administration Console**. This method is the most efficient.
- To populate the Administration Console with information on all the resources in the current perspective: from the main menu bar, select **View > Open > Administration Console**. If you are monitoring a large number of resources, the Administration Console may take a few minutes to load.

2. Expand the objects in the left pane to explore the hierarchy of resource types.
3. Select a high-level resource type (a logical server, for example) in the hierarchy.

   The Administration Console displays a list of resources of that type. You can use the Folder menu to create another server of the same type, or to refresh the view.
4. Select various objects in the hierarchy.
   Information about each selected object appears in the table in the right pane.
5. In either the right or the left pane, select an object.
   A dropdown arrow appears to the right of the name. If the selected object is in the right pane, the **Resource** menu becomes active.
6. Click the dropdown arrow to display a menu of actions you can perform on that object. If the selected object is in the right pane, use the **Resource** menu to display the same actions.

   Note: Some managed objects have no actions.

Job Scheduling

A schedule defines a data collection job and specifies how often the job executes in your system.

In Sybase Control Center, collection jobs provide the data that appears on monitoring screens and charts. A collection is a set of key performance indicators (KPIs). When the scheduler runs a collection job, it gathers the value of each KPI in the collection and tags the data with the date and time it was gathered. The data is stored in the repository and displayed. Each product module has predefined collections that you can schedule.

You can define schedules as one-time or repeating. You can modify the schedule for a job based on a number of attributes such as:

- Repeat interval
- Date
- Time

The job history displays the status of jobs executed each day.

See also

Executing and Stopping a Data Collection Job

Use the Properties view to execute or stop a data collection job.

Most of the time, data collection jobs should run on a schedule; you should rarely need to start or stop a job manually.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties**.
2. Select **Collection Jobs**.
3. Select the job and:

   - To execute a job immediately, click **Execute**.

   - To stop a job, click **Stop**, then click **Yes** to confirm.

See also
- *Deleting a Data Collection Job* on page 163
- *Resuming and Suspending a Data Collection Job* on page 164
- *Adding a New Schedule to a Job* on page 164
- *Modifying the Data Collection Interval for a Job* on page 165
- *Resuming and Suspending the Scheduler* on page 166
- *Viewing the Job Execution History* on page 166

Deleting a Data Collection Job

Use the Properties view for a resource to delete one or more data collection jobs.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties.**
2. Select **Collection Jobs**.
3. Select the job and click **Delete**.
4. Click **OK** to confirm the deletion.

See also
- *Executing and Stopping a Data Collection Job* on page 163
- *Resuming and Suspending a Data Collection Job* on page 164
- *Adding a New Schedule to a Job* on page 164
- *Modifying the Data Collection Interval for a Job* on page 165
- *Resuming and Suspending the Scheduler* on page 166
- *Viewing the Job Execution History* on page 166

Resuming and Suspending a Data Collection Job

Use the Properties view for a resource to resume or suspend a data collection job.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties.**
2. Select **Collection Jobs**.
3. Select the job (a top-level item in the Collection Jobs table). On the **General** tab:

   • To resume a job, click **Resume**.
   • To suspend a job, click **Suspend**, then click **Yes** to confirm the suspension.

   Tip: If the **General** tab is grayed out, you have selected a schedule (child) rather than a job (parent) in the Collection Jobs table. Select the parent job to display the **General** tab.

See also

Adding a New Schedule to a Job

Use the Properties view for a resource to add more than one schedule to a job.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties.**
2. Select **Collection Jobs**.
3. Select the job.
4. Click **Add Schedule**.
5. Specify details for the new schedule:

Field	Description
Name	A name for this schedule
Description	A description of this schedule

6. Choose to start the job **Now** or **Later**.
7. Specify the duration of this schedule. The job can run:

   • **Once**
   • **Repetitively** at an interval you specify

Field	Description
Repeat interval	Time period (in seconds, minutes, hours, or days) between job executions

- **Until** a stop date that you specify, at an interval you specify

Field	Description
Repeat interval	Time period (in seconds, minutes, hours, or days) between job executions
Stop date	Date and time the job should stop running

Note: Enter dates and times using your local time. Sybase Control Center converts your times for remote time zones if necessary.

8. Click **Finish**.

See also

- *Executing and Stopping a Data Collection Job* on page 163
- *Deleting a Data Collection Job* on page 163
- *Resuming and Suspending a Data Collection Job* on page 164
- *Modifying the Data Collection Interval for a Job* on page 165
- *Resuming and Suspending the Scheduler* on page 166
- *Viewing the Job Execution History* on page 166

Modifying the Data Collection Interval for a Job

Use the Properties view for a managed resource to modify the data collection schedule.

1. In the Perspective Resources view, select a server (or other resource).
2. In the view's menu bar, select **Resource > Properties.**
3. Select **Collection Jobs**.
4. Expand a job folder and select a schedule.
5. On the **Schedule** tab, modify the Repeat interval field.
6. Click **Apply**.

See also

- *Executing and Stopping a Data Collection Job* on page 163
- *Deleting a Data Collection Job* on page 163
- *Resuming and Suspending a Data Collection Job* on page 164
- *Adding a New Schedule to a Job* on page 164
- *Resuming and Suspending the Scheduler* on page 166
- *Viewing the Job Execution History* on page 166

Resuming and Suspending the Scheduler

Use the scheduler settings to resume or suspend all scheduled jobs.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu bar, select **Application > Administration.**
2. In the Sybase Control Center Properties dialog, select **Scheduler**.
3. Do one of the following:

   - To resume the scheduler, click **Resume**.
   - To suspend the scheduler, click **Suspend**.
4. Click **OK**.

See also

- *Executing and Stopping a Data Collection Job* on page 163
- *Deleting a Data Collection Job* on page 163
- *Resuming and Suspending a Data Collection Job* on page 164
- *Adding a New Schedule to a Job* on page 164
- *Modifying the Data Collection Interval for a Job* on page 165
- *Viewing the Job Execution History* on page 166

Viewing the Job Execution History

Use the Properties view to display a data collection job's execution history.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties.**
2. Select **Collection Jobs**.
3. Select a job.
4. Click the **History** tab.

See also

- *Executing and Stopping a Data Collection Job* on page 163
- *Deleting a Data Collection Job* on page 163
- *Resuming and Suspending a Data Collection Job* on page 164
- *Adding a New Schedule to a Job* on page 164
- *Modifying the Data Collection Interval for a Job* on page 165
- *Resuming and Suspending the Scheduler* on page 166

Alerts

You can configure Sybase Control Center to notify you when a resource requires attention.

You do this by setting up a predefined alert that is triggered when a performance counter enters a particular state or passes a threshold value that you set. When the alert goes off, it generates an alert notification.

An alert notification takes the form of a visual indicator in the Alert Monitor and, optionally, an e-mail message. The Alert Monitor displays information about the alert, including the resource name, alert severity, value, and date. You can resolve the alert or allow it to escalate.

Configure, monitor, and control alerts for managed resources by:

- Enabling and disabling alert subscriptions for resources
- Configuring shell scripts to run when alerts fire
- Setting alert state or threshold triggers
- Responding to an alert by resolving it, adding notes if desired
- Modifying or deleting alerts
- Viewing alert history

See also
- *Alert-Triggered Scripts* on page 154
- *Creating an Alert* on page 146
- *Assigning a Role to a Login or a Group* on page 106
- *Configuring the E-mail Server* on page 104

Types, Severities, and States

Learn about the properties that define and control alerts.

An alert's type determines what causes it to fire.

Table 38. Alert types

Type	Description
State	A state alert fires when the metric on which it is based changes to a particular state. The possible states are running, pending, stopped, warning, error, and unknown.
Threshold	A threshold alert fires when the metric on which it is based passes a specified level.

Alert severities control when an alert is issued. You can configure the states or threshold values for each alert.

Table 39. Alert severities

Severity	Description
Normal	No alert is issued.
Warning	A problem has given cause for concern. An alert is issued; you can choose whether to subscribe to alerts that fire at the Warning level.
Critical	A serious problem exists. An alert is issued; you can choose whether to subscribe to alerts that fire at the Critical level.

State-based alerts use these states:

- Running
- Pending
- Unknown
- Warning
- Stopped
- Error

The definitions of these states vary by component and sometimes by alert. See the component-specific topics for details.

See also

Viewing Alerts
Display alert notifications and alerts that have been configured for a given resource.

- To display generated alerts (notifications):
    a) Select **View > Open > Alert Monitor** from the application menu bar.
    For a given alert, the Alert Monitor displays only the most recent unresolved notifications at each severity level. That is, if an alert fires five times at the warning level, only the notification of the fifth firing is listed—even if the previous four alerts remain unresolved.
    b) To display information about a generated alert, select the alert in the Alert Monitor and click **Properties**.

- To display configured alerts:
  a) In the Perspective Resources view, select a resource and select **Resource > Properties.**
  b) Click **Alerts** to view configured alerts for the selected resource.

  (This is a different route to the information displayed in the second step, above.)

See also
- *Types, Severities, and States* on page 167
- *Modifying an Alert* on page 169
- *Testing an Alert-Triggered Script* on page 170
- *Deleting an Alert* on page 171
- *Alert Subscriptions* on page 171
- *Alert Notifications* on page 173
- *Creating an Alert* on page 146

Modifying an Alert
Use the Properties view of your managed resource to modify an alert.

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select the alert to modify.
4. On the Thresholds tab, modify the threshold values. Click **OK** to save your changes.
5. On the Script tab, click **Modify** to change the alert severity at which script execution is triggered, the path to the script, the execution parameters, or the test values. Click **Finish** to save your changes.
6. On the Subscriptions tab, select a subscription and click **Modify** to change its e-mail address or escalation address. Click **Finish** to save your changes.
7. On the Storm Supression tab, pull down the menu to change the units and enter a value for the storm suppression period.
8. Click **OK** (to apply the changes and close the properties dialog) or **Apply** (to apply the changes and leave the dialog open).

See also
- *Types, Severities, and States* on page 167
- *Viewing Alerts* on page 168
- *Testing an Alert-Triggered Script* on page 170
- *Deleting an Alert* on page 171
- *Alert Subscriptions* on page 171
- *Alert Notifications* on page 173
- *Creating an Alert* on page 146

Testing an Alert-Triggered Script
Execute a script to make sure it works properly.

Prerequisites
Configure an alert with a script.

Task

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select the alert to test.
4. On the Script tab, click **Modify**.
5. If the script requires parameter values, click **Select Parameters** to enter them in the **Execution Parameters** box.

   You can include a number of predefined substitution parameters, which are replaced by values from the alert. The parameter values are passed on the command line to the script. For the test execution, use values that test all the parameters used by the script. See the substitution parameters topic (linked below) for more information.

   Note: When you test a script, Sybase Control Center supplies test values for the **%Severity %** and **%Source_Application%** parameters ("Testing" and "TestScriptExecution," respectively). Any test values you supply for these parameters are discarded. This prevents the test results from being confused with real script results after testing and in the SCC repository.

6. Click **Test** to perform a test execution of your script.

   If your script takes parameters, the test may fail if parameter values are missing or incorrect.

See also
- *Types, Severities, and States* on page 167
- *Viewing Alerts* on page 168
- *Modifying an Alert* on page 169
- *Deleting an Alert* on page 171
- *Alert Subscriptions* on page 171
- *Alert Notifications* on page 173
- *Alert-Triggered Scripts* on page 154
- *Substitution Parameters for Scripts* on page 155
- *Creating an Alert* on page 146

Deleting an Alert

Use the Properties view of your resource to delete an alert.

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select an alert and click **Drop**.
4. Click **Yes** to confirm the deletion.

See also
*   *Types, Severities, and States* on page 167
*   *Viewing Alerts* on page 168
*   *Modifying an Alert* on page 169
*   *Testing an Alert-Triggered Script* on page 170
*   *Alert Subscriptions* on page 171
*   *Alert Notifications* on page 173
*   *Creating an Alert* on page 146

Alert Subscriptions

When an alert subscription is configured, the alert notifies the specified user or group of users by e-mail message when the alert fires.

You can configure an alert subscription to send e-mail notifications when the alert reaches a severity of warning, a severity of critical, or both.

You can also configure an alert subscription to escalate after a period of time that you specify. If the alert is not resolved within the escalation period, Sybase Control Center e-mails an escalation message to the user or group whose address you provide for escalations, as well as to the primary subscriber. The escalation message is identical to the primary notification message. Sybase recommends that if you configure alert subscriptions to escalate, you do so only for the most urgent alerts, those with a severity of critical.

See also
*   *Types, Severities, and States* on page 167
*   *Viewing Alerts* on page 168
*   *Modifying an Alert* on page 169
*   *Testing an Alert-Triggered Script* on page 170
*   *Deleting an Alert* on page 171
*   *Alert Notifications* on page 173
*   *Creating an Alert* on page 146

Adding or Modifying an Alert Subscription

Use the Properties view to subscribe to an alert or edit an alert subscription.

Prerequisites

Specify the e-mail server to which Sybase Control Center will send e-mail alert notifications.

Task

Each alert can support one subscription. To change addresses, modify the alert's existing subscription.

Note: E-mail notifications are sent from an address of the form SybaseControlCenter@yourdomain—for example, SybaseControlCenter@Bigcompany.com. Make sure your mail system does not block or filter that address.

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select an alert instance.
4. On the **Subscriptions** tab:
   - Click **Add** to create a subscription, or
   - Select a subscription and click **Modify** to edit an existing subscription
5. Follow the instructions in the Add Alert Subscription wizard.

   For both critical and warning alerts:

Table 40. Alert subscription details

Option	Description
E-mail message	To send an e-mail notification when this alert fires, click the **E-mail message** box and enter the e-mail address of one user or list.
Escalation e-mail	To escalate this alert (by sending an e-mail notification to another address when this alert has not been responded to after a specified period of time), click the **Escalation e-mail** box and enter the e-mail address of one user or list.
Time period	Enter the amount of time to wait, following the initial alert notification, before Sybase Control Center sends an e-mail notification to the escalation address.

6. Click **Finish**.

See also
- *Unsubscribing from an Alert* on page 173

- *Enabling and Disabling Alert Subscription* on page 173

Unsubscribing from an Alert
Use the Properties view to unsubscribe from an alert.

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select an alert instance.
4. In the Subscriptions tab, select the alert subscription and click **Drop**.
   When you drop a regular subscription, any escalation subscription is also dropped.
   However, dropping an escalation does not affect the regular subscription.
5. Click **Yes** to confirm the deletion.

See also
- *Adding or Modifying an Alert Subscription* on page 172
- *Enabling and Disabling Alert Subscription* on page 173

Enabling and Disabling Alert Subscription
Use the Properties view to enable and disable alert subscription.

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select an alert instance.
4. In the **Subscriptions** tab, select an alert subscription and:

   - To enable subscription, click **Enable**.
   - To disable subscription, click **Disable**, then click **Yes** to confirm.

See also
- *Adding or Modifying an Alert Subscription* on page 172
- *Unsubscribing from an Alert* on page 173

Alert Notifications
An alert notification indicates that an alert has been generated.

Alert notifications are produced when alerts fire. An alert fires if the performance indicator on which it is based passes the threshold or state specified for the severity level of warning. If the performance indicator passes the threshold or state specified for the severity level of critical, the alert fires again and another notification is generated.

Detailed alert notifications appear in the Alert Monitor view. In addition, alerts appear as yellow ! symbols in the heat chart. You can set an alert to also send an e-mail message when it fires.

See also

- *Types, Severities, and States* on page 167
- *Viewing Alerts* on page 168
- *Modifying an Alert* on page 169
- *Testing an Alert-Triggered Script* on page 170
- *Deleting an Alert* on page 171
- *Alert Subscriptions* on page 171
- *Creating an Alert* on page 146

Displaying Alert History and Resolutions

Use the Properties view to see historical information about resolved and unresolved alerts.

The History tab on the Alerts page of the Resource Properties view displays information about every time this alert has fired. Each row of the table represents a single notification generated by the selected alert.

The Resolutions tab displays information about alerts that have been resolved (closed) by a Sybase Control Center administrator.

The History and Resolutions tabs display the 100 most recent alerts or alerts for the last 24 hours, whichever is reached first.

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. Select **Alerts**.
3. Select the alert instance.
4. Click the **History** tab.
5. (Optional) Click the **Resolutions** tab.

See also

- *Resolving Alerts* on page 174

Resolving Alerts

After you address the cause of an alert, resolve it to remove it from the list of active alerts in the Alert Monitor.

Prerequisites

You must be logged in as a user with Sybase Control Center administrative privileges (sccAdminRole) to resolve alerts.

Task

1. In the Perspective Resources view, select a resource and select **Resource > Properties.**
2. In the left pane, select **Alerts**.

3. Select an alert instance in the top table.

4. Click **Resolve**.

5. Enter an explanation of how you resolved the alert.

6. Click **Submit**.
The state of the alert (shown in the State column) changes to Normal. Notifications on this alert disappear from the Alert Monitor.

Note: See the Resolutions tab for details on resolved alerts.

See also
- *Displaying Alert History and Resolutions* on page 174

Resources

In Sybase Control Center, a resource is a unique Sybase product component or subcomponent. A server is the most common managed resource.

Sybase products comprise many components, including servers, agents, databases, devices, and processes. A managed resource is a product component or subcomponent that Sybase Control Center lets you monitor and administer. Two important tools for resource management are the Resource Explorer and the Perspective Resources view.

- The Resource Explorer lists resources that are registered with Sybase Control Center. The list may include resources that you have not yet added to a perspective. Registration enables Sybase Control Center to connect to the resource, log in, retrieve monitoring data, and issue commands. Resources are registered at the server or agent level, and registering a server or agent also makes Sybase Control Center aware of any subcomponents. You can register resources individually or register several at once by importing them in a batch.
- The Perspective Resources view lists registered resources that you have added to the current perspective. You must add a resource to a perspective to manage and monitor its availability and performance.

See also
- *Registering a Sybase IQ Server* on page 120
- *Importing Resources for Batch Registration* on page 122

Unregistering a Resource
Remove one or more servers or other resources from Sybase Control Center.

1. From the Sybase Control Center toolbar, click the **Launch Resource Explorer** icon.

2. In the Resource Explorer, select the resources you want to unregister. Use **Shift+click** or **Control+click** to select multiple resources.

3. Select **Resources > Unregister.**

4. Click **Yes** to confirm the removal.

See also
*   *Adding a Resource to a Perspective* on page 176
*   *Removing a Resource from a Perspective* on page 176
*   *Searching for Resources in the Resource Explorer* on page 177
*   *Registering a Sybase IQ Server* on page 120
*   *Importing Resources for Batch Registration* on page 122

Adding a Resource to a Perspective
Add one or more resources to the current perspective.

Add servers or other resources to a perspective so you can monitor and manage them along with other resources in the same perspective.

1. From the Sybase Control Center toolbar, click the **Launch Resource Explorer** icon.
2. Select the resources to add to your perspective. Use **Shift-click** or **Control-click** to select multiple resources.
3. Perform one of these actions:

*   Select **Resources > Add Resources to Perspective.**
*   Drag and drop resources from the Resource Explorer onto the Perspective Resources view. You can select and drag multiple resources.

See also
*   *Unregistering a Resource* on page 175
*   *Removing a Resource from a Perspective* on page 176
*   *Searching for Resources in the Resource Explorer* on page 177

Removing a Resource from a Perspective
Remove one or more resources from the current perspective.

1. To open the Perspective Resources view, click the **Show/Hide the Resource Browser** icon in the perspective toolbar.
2. In the Perspective Resources view, select the resources to remove. Use **Shift-click** or **Control-click** to select multiple resources.
3. Select **Resource > Remove.**
4. Click **Yes** to confirm the removal.

See also
*   *Unregistering a Resource* on page 175
*   *Adding a Resource to a Perspective* on page 176

Searching for Resources in the Resource Explorer

Search for all your managed resources or narrow your search for a particular resource.

1. Click the **Launch Resource Explorer** icon.
2. If the Filter pane is not visible in the Resource Explorer window, select **View > Filter** from the view's menu bar.
3. Enter your search term in the **Filter string** field.

   The search term can be any string that appears in the tabular portion of the Resource Explorer, such as the name, or part of the name, of a server or a resource type (ASE Server, for example).
4. (Optional) Select a filtering setting:

   - **Match case** – search for resources whose displayed data includes the search term, including uppercase and lowercase letters; or
   - **Exact match** – search for resources whose displayed data includes an item identical to the search term.
5. (Optional) Select a column from the **Filter on** list to restrict your search to that column.

See also

Perspectives

A perspective is a named container for a set of one or more managed resources. You can customize perspectives to provide the information you need about your environment.

As the main workspaces in the Sybase Control Center window, perspectives let you organize managed resources. You might assign resources to perspectives based on where the resources are located (continents, states, or time zones, for example), what they are used for, which group owns them, or which administrator manages them. Perspectives appear as tabs in the main window.

Every perspective includes a Perspective Resources view, which lists the resources in that perspective and provides high-level status and descriptive information. Use the View menu to switch from detail view to icon view and back.

You can open additional views—the heat chart, statistics chart, or alert monitor, for example —as needed to manage the perspective's resources. The views in a perspective display information only about resources in that perspective.

One resource can appear in many perspectives.

See also
* *Creating a Perspective* on page 125

Creating a Perspective
Create a perspective in which you can add and manage resources.

1. From the application menu bar, select **Perspective > Create.**
2. Enter a name for your perspective. The name can contain up to 255 characters.
3. Click **OK**.

See also
* *Removing a Perspective* on page 178
* *Renaming a Perspective* on page 178

Removing a Perspective
Delete a perspective window.

1. Select the perspective tab you want to delete.
2. In the main menu bar, select **Perspective > Delete.**
   The selected perspective disappears. If there are other perspectives, Sybase Control Center displays one.

See also
* *Creating a Perspective* on page 178
* *Renaming a Perspective* on page 178

Renaming a Perspective
Change the name of your perspective.

1. Select the perspective tab you want to rename.
2. From the main menu bar, select **Perspective > Rename.**.
3. Enter the new name for your perspective.
4. Click **OK**.

See also
* *Creating a Perspective* on page 178
* *Removing a Perspective* on page 178

Views

Use views to manage one or more resources within a perspective.

In Sybase Control Center, views are the windows you use to monitor and manage a perspective's resources. You can re-arrange, tile, cascade, minimize, maximize, and generally control the display of the views in your persective.

Each perspective includes these views:

- Perspective Resources
- Administration Console
- Heat chart
- Alert monitor
- Component log viewer
- Views that exist for each managed resource. These vary by resource type, but typically include the statistics chart, the properties view, and a monitoring view.

Managing a View

Open, close, minimize, maximize, or restore a view in the current perspective.

You can:

Task	Action
Open a view	Do one of the following: • In the Perspective Resources view, click a resource, pull down its menu using the handle to the right of the resource name, and select the view to open. • In the application menu bar, select **View > Open** and choose a view.
Close a view	Select the view to close. In the application menu bar, select **View > Close.** You can also click the **X** in the view's upper right corner.
Maximize a view	Click the box in the view's upper right corner. The view enlarges to fill the entire perspective window. Click the box again to return the view to its former size.
Minimize a view	Click the _ in the view's upper right corner. The view shrinks to a small tab at the bottom of the perspective window.
Minimize all views	In the application menu bar, select **View > Minimize All Views.**
Restore a view	Click the box on the minimized tab to maximize the view. Click the box again to return the view to its former (smaller) size so you can see other views at the same time.

Task	Action
Bring a view to the front	In the application menu bar, select **View > Select** and choose the view you want from the submenu.

See also
* *Arranging View Layout in a Perspective* on page 180

Arranging View Layout in a Perspective
Use the view layout options to manage your perspective space.

Click one of these icons from the Sybase Control Center toolbar:

* **Cascade all open views**
* **Tile all open views vertically**
* **Tile all open views horizontally**

In a cascade, views overlap; in tiling arrangements, they do not.

Alternatively, you can arrange view layouts from the Sybase Control Center menu bar. From the menu bar, select **Perspective > Arrange** and select your view layout.

See also
* *Managing a View* on page 179

Instances
Deploy, remove, refresh, or convert Sybase Control Center server or agent instances running from an installation on a shared disk.

Enabling and Disabling Shared-Disk Mode
Turn on or turn off shared-disk mode, which allows you to run multiple Sybase Control Center agents and servers from a single installation on a shared disk.

Prerequisites
Install Sybase Control Center on a shared disk. See the *Sybase Control Center Installation Guide*.

Task

Shared-disk mode affects the entire installation; do not enable or disable individual instances.

Disabling shared-disk mode leaves the instances' file systems intact under <SCC-install-directory>/instances, but the instances cannot run. If you reenable, the instances are able to run again.

1. Change to `SCC-3_2/bin`.

2. Enable or disable shared disk mode.

   To enable shared disk mode:
   ```
   sccinstance -enable
   ```

   To disable shared disk mode:
   ```
   sccinstance -disable
   ```

See also
- *Deploying an Instance from a Shared Disk Installation* on page 181
- *Refreshing or Converting an Instance* on page 182
- *Removing an Instance* on page 184
- *Shared-Disk Mode* on page 185
- *sccinstance Command* on page 185

Deploying an Instance from a Shared Disk Installation

(Optional) Create a Sybase Control Center server or agent from an installation on a shared disk.

Prerequisites

- Install Sybase Control Center on a shared disk.
- Enable shared-disk mode.

Task

1. Log in to the host on which you plan to run the SCC server or agent.

   Note: You can create an instance on one host and run it on another host, but doing so interferes with the predeployment checks run by **sccinstance**. Such a deployment might generate errors (port conflicts, for example). If you are confident that the errors are caused by problems that will not be present on the host where you plan to run the instance, use the **-force** option to create the instance.

2. Change to `SCC-3_2/bin`.

3. Create the instance as an SCC agent if you plan to run a managed server on this host. Create the instance as an SCC server if you plan to manage other Sybase servers from this host.

   To create an SCC agent called Boston-agent and configure it to run as a Windows service:

   ```
   sccinstance -create -agent -instance Boston-agent -service
   ```

   To create an SCC server called Boston and configure it to run as a Windows service:

   ```
   sccinstance -create -server -instance Boston -service
   ```

4. If other SCC instances will run on this host, change the port assignments for the new instance. Change the instance names and port values in the sample commands to suit your environment, but take care to specify ports that are not in use by another SCC instance or any other application or server.

    This command changes the port assignments for an SCC agent called myagent:
    ```
    sccinstance -refresh -instance myagent -portconfig
    rmi=8888,jiniHttp=9093,jiniRmi=9096,tds=9997
    ```

    This command changes the port assignments for an SCC server called myserver:
    ```
    sccinstance -refresh -server -instance myserver -portconfig
    rmi=8889,db=3640,
    http=7072,https=7073,jiniHttp=9094,jiniRmi=9097,msg=2002,tds=9996
    ```

5. (Optional) List the instances deployed from this installation:
    ```
    sccinstance -list
    ```

6. (Optional) If you are setting up an instance in UNIX, configure it to run as a service. (See *Starting and Stopping Sybase Control Center in UNIX*).

Next

When you manage and maintain instances, keep in mind that the directory structure for instances is different from that of singleton installations. In file paths in SCC help, replace `SCC-3_2` or `<scc-install-directory>` with `SCC-3_2/instances/<instance-name>`.

For example, the path to the log directory, `SCC-3_2/log`, becomes this for an instance called kalamazoo:

```
SCC-3_2/instances/kalamazoo/log
```

See also
*   *Enabling and Disabling Shared-Disk Mode* on page 180
*   *Refreshing or Converting an Instance* on page 182
*   *Removing an Instance* on page 184
*   *Shared-Disk Mode* on page 185
*   *sccinstance Command* on page 185

Refreshing or Converting an Instance
Refresh a Sybase Control Center server or agent deployed from an installation on a shared disk, or convert between server and agent.

Prerequisites
Shut down the instance.

Task

When you refresh an instance of an SCC server or agent, SCC recopies files from the main installation on the shared disk (SCC-3_2/) into the instance's subdirectories (SCC-3_2/instances/<instance-name>). In Windows, SCC recopies all the files that make up this instance; in UNIX, it recopies all this instance's services and plug-ins.

Refreshing an instance preserves configuration and logs but overwrites the repository, so historical performance data is lost.

As part of a refresh, you can:

- Convert a server to an agent
- Convert an agent to a server
- Reassign ports on the instance

Converting from an agent to a server adds server-related files to the instance; converting from a server to an agent removes files.

1. Change to SCC-3_2/bin.
2. Refresh the instance. Change the instance names and port values in the sample commands to suit your environment, but take care to specify ports that are not in use by another SCC instance or any other application or server.

   This command refreshes an SCC server called boston. If boston is an agent, it becomes a server after the refresh.

   ```
   sccinstance -refresh -server -instance boston
   ```

   This command refreshes an SCC agent called kalamazoo. If kalamazoo is a server, it becomes an agent after the refresh.

   ```
   sccinstance -refresh -agent -instance kalamazoo
   ```

   This command refreshes an SCC agent called kalamazoo and reassigns kalamazoo's RMI and TDS ports. If kalamazoo is a server, it becomes an agent after the refresh.

   ```
   sccinstance -refresh -agent -instance kalamazoo -portconfig
   rmi=7070,tds=7071
   ```

3. (Optional) Display the status of the refreshed instance. Replace the name in the sample command with your instance's name, or omit the **-instance** option to display the status of the instance on this host.

   ```
   sccinstance -instance kalamazoo
   ```

See also
- *Enabling and Disabling Shared-Disk Mode* on page 180

- *Deploying an Instance from a Shared Disk Installation* on page 181
- *Removing an Instance* on page 184
- *Shared-Disk Mode* on page 185
- *sccinstance Command* on page 185

Removing an Instance

Delete a Sybase Control Center server or agent deployed from an installation on a shared disk.

Prerequisites

Shut down the instance.

Task

Removing an SCC instance deletes the instance's files and directories (SCC-3_2/ instances/<instance-name> and its contents) from the installation.

You cannot restore a removed instance.

1. Change to SCC-3_2/bin.
2. Remove the instance. Change the instance names in the sample commands to suit your environment.

   This command removes an SCC server called porcupine if it is not running; if it is running, you see an error.

   ```
   sccinstance -remove -instance porcupine
   ```

   This command removes the SCC agent on the current host if it is not running. If the agent is running, the command returns an error.

   ```
   sccinstance -remove
   ```

See also
- *Enabling and Disabling Shared-Disk Mode* on page 180
- *Deploying an Instance from a Shared Disk Installation* on page 181
- *Refreshing or Converting an Instance* on page 182
- *Shared-Disk Mode* on page 185
- *sccinstance Command* on page 185

Shared-Disk Mode

Shared-disk mode lets you run multiple Sybase Control Center instances—SCC servers, SCC agents, or a mixture of the two—from a single installation of the product.

The shared-disk capability enables SCC servers or agents on the installation host or on remote hosts to access and execute from the same installation. This feature is especially useful if you plan to use SCC to manage Adaptive Server clusters or Sybase IQ multiplexes.

After installing SCC on a shared disk, use the **sccinstance** command to enable shared-disk mode and deploy instances. **sccinstance** copies the files needed for the instance into a new directory structure. The path takes the form `<SCC-install-directory>/instances/<instance-name>` (for example, `SCC-3_2/instances/SCCserver-1`).

You can specify a name for each instance. If you do not supply a name, the instance name defaults to the host name.

An instance runs on the host on which you start it. When shared-disk mode is enabled, SCC servers and agents run out of the `SCC-3_2/instances` subdirectories, not from the base file system.

In shared-disk mode, changes made to configuration files in the base file system (everything under `SCC-3_2` except the `SCC-3_2/instances` branch) are copied to any instance deployed thereafter. Previously deployed instances are not affected.

Use **sccinstance** to deploy, remove, refresh, or convert an instance; to configure an instance's ports; and to configure a Windows instance to run as a service. Perform other tasks, including configuring a UNIX instance to run as a service, and all other configuration, using the tools and procedures documented for all installations. Use tools provided by the UI wherever possible. When you must edit a file to change the configuration of an instance (for role mapping, for example), edit the copy of the file stored under `<SCC-install-directory>/instances/<instance-name>`.

See also

sccinstance Command

Use **sccinstance.bat** (Windows) or **sccinstance** (UNIX) to deploy an instance of Sybase Control Center from a shared-disk installation or to manage existing instances.

You can run multiple instances of Sybase Control Center, including SCC servers, SCC agents, or a mixture of the two, from a single installation on a shared disk.

Syntax

```
sccinstance[.bat]
[-agent]
[-c | -create]
[-d | -debug]
[-disable]
[-enable]
[-f | -force]
[-h | -help]
[-i | -instance [instance-name]]
[-l | -list]
[-plugins {plugin-ID,plugin-ID,...}]
[-portconfig {port-name=port-number,port-name=port-number, ...}]
[-refresh]
[-r | -remove]
[-s | -server]
[-service]
[-silent]
```

Parameters

- **-agent** – use with **-create** or **-refresh** to create or refresh an SCC agent. In a **-create** or **-refresh** command, **-agent** is the default, so you can omit it.
- **-create** – deploy a new instance. Use alone or with **-agent** to create an agent instance, or with **-server** to create a server instance.
- **-d | debug** – display debugging messages with the output of this command.
- **-disable** – turn off shared-disk mode for this installation. Generates an error if any instance is running.
- **-enable** – turn on shared-disk mode for this installation. Shared-disk mode is required if you intend to run more than one server or agent from a single installation of SCC.
- **-f | -force** – execute **sccinstance** even if there are potential conflicts (such as port clashes or a running SCC process).
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-instance** – specify an instance. Use with **-create**, **-remove**, or **-refresh**, or use alone to display the instance's status. You can omit **-instance** when you are addressing the only SCC instance or the only instance of the specified type (server or agent) on the current host.
- **-l | -list** – display a list of all instances deployed from this SCC installation.
- **-plugins {*plugin-ID,plugin-ID,...*}** – specify one or more product module plug-ins for this instance. An alternative to **-agent** and **-server**, **-plugins** is primarily for use by the SCC installation program. Use with **-create** or **-refresh**. Use commas to separate plug-in names.
- **-portconfig {*port-name=port-number*, *port-name=port-number*, ...}** – assign ports to services for this instance. Use only with **-create** or **-refresh**. For the *port-name* value, use a port name from the table below. If you plan to run more than one SCC instance on a host machine, you must reassign all the ports for every instance after the first.

  Port information:

Port Name	Description	Service Names	Property Names	Default Port
db	Database port Present on SCC server	SccSADataserver Messaging Alert Scheduler	com.sybase.asa.server.port messaging.db.port alert.database.port org.quartz.data-Source.ASA.URL	3638
http	Web HTTP port Present on SCC server	EmbeddedWebCon-tainer	http.port	8282
https	Web HTTPS (secure HTTP) port Present on SCC server	EmbeddedWebCon-tainer	https.port	8283
jiniHttp	JINI HTTP server Present on SCC server and SCC agent	Jini	httpPort	9092
jiniR-mid	JINI remote method invocation daemon Present on SCC server and SCC agent	Jini	rmidPort	9095
msg	Messaging port Present on SCC server	Messaging	messaging.port	2000
rmi	RMI port Present on SCC server and SCC agent	RMI	port	9999
tds	Tabular Data Stream™ port (used to communicate with other Sybase products) Present on SCC server and SCC agent	Tds	tdsPort	9998

- **-refresh** – recopy all the files that make up this instance (Windows) or all this instance's services and plug-ins (UNIX). Refreshing preserves any service or plug-in configuration in the deployed instance.

  You can also use **-refresh** to convert a server to an agent or an agent to a server (see the examples). Files are removed or added to change the function of the instance. Use alone or

with **-agent** to refresh an agent instance, or with **-server** to refresh a server instance. Generates an error if the instance is running.

- **-r | -remove** – delete an instance. Use alone or with **-instance**. Generates an error if the instance is running. You cannot restore a removed instance.
- **-s | -server** – use with **-create** or **-refresh** to create or refresh an SCC server, including any product modules available.
- **-service** – use with **-create** or **-remove** to create or remove a Windows service for this instance. You must be logged in to Windows as an administrator to use this option.
- **-silent** – suppress the output of **sccinstance**.

Examples

- **Deploy an SCC server instance** – enables shared-disk mode, deploys a server called Boston with a Windows service, and starts the Windows service:

```
sccinstance -enable
sccinstance -create -server -instance Boston -service
net start "Sybase Control Center 3.2.3 (Boston)"
```

Note: To create the service, you must log in to Windows as an administrator.

- **Deploy an SCC agent instance** – deploys an SCC agent on this host and configures a Windows service for it. The **-agent** option, because it is the default, is not required—the command does exactly the same thing without it.

```
sccinstance -create -agent -service
```

or

```
sccinstance -create -service
```

- **Deploy a server instance and reassign ports** – deploys the server on this host and configures nondefault RMI, HTTP, and HTTPS ports.

```
sccinstance -create -server -portconfig
rmi=8888,http=7070,https=7071
```

- **Refresh a server instance or convert an agent to a server** – refreshes the server on this host. If the instance on this host is an SCC agent, refreshing it as an SCC server converts it into a server.

```
sccinstance -refresh -server
```

- **Refresh an agent instance or convert a server to an agent** – refreshes the instance named kalamazoo. If kalamazoo is a server, refreshing it as an SCC agent converts it into an agent.

```
sccinstance -refresh -agent -instance kalamazoo
```

- **Remove a server instance** – removes the instance named porcupine if it is not running:

```
sccinstance -remove -instance porcupine
```

- **Display status** – displays the status of the instance on this host:

```
sccinstance
```

- **List all instances** – displays a list of all SCC server and agent instances deployed from this SCC installation:

```
sccinstance -list
```

- **Scenario: Remove an instance by force** – suppose you have inadvertently deployed two SCC agent instances on the same host:

```
$ sccinstance -list
2 SCC instances deployed:
SCC instance node1 deployed in agent mode for host node1 RMI port
9999
SCC instance node2 deployed in agent mode for host node2 RMI port
9999
```

Both instances use the same RMI port. You must either reassign ports for one instance or remove it. But you get an error if you try remove an instance when another instance is running on the same host:

```
$ sccinstance -instance node2 -remove
[ERROR] Command execution failed.
[ERROR] SCC instance node2 could not be removed because it is
running. Shut
down the SCC before removing the instance.
```

Use the **-force** option to override the error and force the removal of the second agent instance:

```
$ sccinstance -instance node2 -remove -force
Removing SCC instance node2 ...
SCC instance node2 was successfully removed.
```

Permissions

sccinstance permission defaults to all users, except as noted for certain parameters.

See also
- *Enabling and Disabling Shared-Disk Mode* on page 180
- *Deploying an Instance from a Shared Disk Installation* on page 181
- *Refreshing or Converting an Instance* on page 182
- *Removing an Instance* on page 184
- *Shared-Disk Mode* on page 185

Repository

The Sybase Control Center embedded repository stores information related to managed resources, as well as user preference data, operational data, and statistics.

You can back up the repository database on demand, schedule automatic backups, restore the repository from backups, and configure repository purging options. Full and incremental backups are available. A full backup copies the entire repository. An incremental backup copies the transaction log, capturing any changes since the last full or incremental backup.

By default, Sybase Control Center saves backups as follows:

- Each full backup is stored in its own subdirectory in `<SCC-install-directory>/backup`.
- Each incremental backup is stored in a file in `<SCC-install-directory>/backup/incremental`.

Sybase recommends that you periodically move backup files to a secondary storage location to prevent the installation directory from becoming too large.

Scheduling Backups of the Repository

Configure full and incremental backups of the repository to occur automatically.

Prerequisites

Determine your backup strategy, including when to perform full backups and incremental backups. For example, you might schedule incremental backups every day and a full backup every Saturday.

You must have administrative privileges (sccAdminRole) to perform this task.

Task

A full backup copies the entire repository. An incremental backup copies the transaction log, capturing any changes since the last full or incremental backup.

1. From the main menu, select **Application > Administration.**
2. In the left pane, select **Repository**.
3. Click the **Full Backup** tab.
4. (Optional) To change the directory in which backups will be stored, click **Browse** and navigate to the desired directory.
5. Select **Schedule a Regular Backup**.
6. Specify the day you want scheduled backups to begin. Enter a **Start date** or click the calendar and select a date.
7. (Optional) Use the **Time** and **AM/PM** controls to specify the time at which backups occur.
8. Specify how often backups occur by setting the **Repeat interval** and selecting hours, days, or weeks.
9. (Optional) To purge the repository after each backup, select **Run a repository purge after the backup completes**.
10. If you include purging in the backup schedule, go to the **Size Management** tab and unselect **Automatically purge the repository periodically** to disable automatic purging.
11. Click **Apply** to save the schedule.

12. Click the **Incremental Backup** tab and repeat the steps above to schedule incremental backups to occur between full backups.

Next
Set purging options on the Size Management tab.

See also
- *Modifying the Backup Schedule* on page 191
- *Forcing an Immediate Backup* on page 192
- *Restoring the Repository from Backups* on page 192
- *Configuring Repository Purging* on page 194

Modifying the Backup Schedule
Suspend or resume repository backups or change the backup schedule.

Prerequisites
You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu, select **Application > Administration.**
2. In the left pane, select **Repository**.
3. Choose the type of backup to modify:

   - Click the **Full Backup** tab, or
   - Click the **Incremental Backup** tab.

4. (Optional) To suspend or resume the backup schedule, select or unselect **Schedule a Regular Backup**.
   When you unselect (uncheck) this option, the scheduling area is grayed out and scheduled backups no longer occur. However, the schedule is preserved and you can reinstate it at any time.
5. To change the backup schedule, edit the **Start date**, **Time**, **Repeat interval**, or units. You can also select or unselect **Run a repository purge after the backup completes**.
6. Click **Apply** to save the schedule.

See also
- *Scheduling Backups of the Repository* on page 190
- *Forcing an Immediate Backup* on page 192
- *Restoring the Repository from Backups* on page 192
- *Configuring Repository Purging* on page 194

Forcing an Immediate Backup

Perform an unscheduled full or incremental backup of the repository.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

1. From the main menu, select **Application > Administration.**
2. In the left pane, select **Repository**.
3. Choose the type of backup to run:

    - Click the **Full Backup** tab, or
    - Click the **Incremental Backup** tab.

4. Click **Back up Now**.

Sybase Control Center saves the backup to the directory shown in the Location field.

See also

- *Scheduling Backups of the Repository* on page 190
- *Modifying the Backup Schedule* on page 191
- *Restoring the Repository from Backups* on page 192
- *Configuring Repository Purging* on page 194

Restoring the Repository from Backups

Load backup files into the repository database to revert undesirable changes or to recover from a catastrophic failure.

If you configured Sybase Control Center to store backups somewhere other than the default location, change the source directory in the copy commands in this procedure.

1. Shut down Sybase Control Center.
2. Copy the most recent full backup from `<SCC-install-directory>/backup/<generated_directory_name>` to `<SCC-install-directory>/services/Repository`. For example:

    Windows:

    ```
    copy C:\sybase\SCC-3_2\backup\repository.
    270110161105\scc_repository.db
    C:\sybase\SCC-3_2\services\Repository
    ```

    UNIX:

```
cp <SCC-install-directory>/backup/repository.270110161105/
scc_repository.db
<SCC-install-directory>/services/Repository
```

3. If you have no incremental backups to load,

   a) Also copy the log file from `<SCC-install-directory>/backup/`
      `<generated_directory_name>` to `<SCC-install-directory>/`
      `services/Repository`. For example:

      Windows:

      ```
      copy C:\sybase\SCC-3_2\backup\repository.
      270110161105\scc_repository.log
      C:\sybase\SCC-3_2\services\Repository
      ```

      UNIX:

      ```
      cp <SCC-install-directory>/backup/repository.270110161105/
      scc_repository.log
      <SCC-install-directory>/services/Repository
      ```

   b) Skip to step *5* on page 193.

4. Start the repository database using the **-ad** option, which directs it to load transaction logs
   (incremental backups) from the `incremental` directory. (The database loads full
   backups automatically.) For example:

   Windows:

   ```
   cd <SCC-install-directory>\services\Repository

   ..\..\bin\sa\bin_<platform>\dbsrv11.exe scc_repository -ad
   <SCC-install-directory>\backup\incremental
   ```

   UNIX:

   ```
   cd <SCC-install-directory>/services/Repository

   ../../bin/sa/bin_<platform>/dbsrv11 scc_repository -ad
   <SCC-install-directory>/backup/incremental
   ```

   The repository database loads the full backup and any subsequent incremental backups
   present in the `incremental` directory. Incremental backups are loaded in date order.
   After loading and saving, the database shuts down.

5. Start Sybase Control Center.
   If you loaded incremental backups, SCC starts normally (that is, no further recovery
   occurs). If you copied a full backup to the `Repository` directory, the database recovers
   the repository from the full backup.

Example: Loading incremental backups into the repository database

These commands start SQL Anywhere® on a 32-bit Windows machine:

```
% cd C:\sybase\SCC-3_2\services\Repository

% ..\..\bin\sa\bin_windows32\dbsrv11.exe scc_repository -ad
C:\sybase\SCC-3_2\backup\incremental
```

These commands start SQL Anywhere on a 64-bit machine running Solaris:

```
$ cd /opt/sybase/SCC-3_2/services/Repository
$ ../../bin/sa/bin_sunsparc64/dbsrv11 scc_repository -ad
/opt/sybase/SCC-3_2/backup/incremental
```

See also
- *Scheduling Backups of the Repository* on page 190
- *Modifying the Backup Schedule* on page 191
- *Forcing an Immediate Backup* on page 192
- *Configuring Repository Purging* on page 194

Configuring Repository Purging

Change repository purging options.

Prerequisites

You must have administrative privileges (sccAdminRole) to perform this task.

Task

As you decide how to purge your repository, consider that:

- Purging keeps the repository from absorbing too much disk space.
- By default, purging is enabled. It occurs once a day and purges data older than one day.
- Statistics and alert history can help you detect trends in server performance and user behavior. The Sybase Control Center statistics chart can graph performance data over a period of a year or more if the data is available. If you have enough disk space, consider saving data for a longer period of time or disabling the purging of statistics or alert history.
- Changing the purge frequency and other options might affect Sybase Control Center performance.

Note: If you configure purging as part of a scheduled backup of the repository, disable automatic purging on the Size Management tab.

1. From the main menu bar, select **Application > Administration.**
2. Select **Repository**.
3. Click the **Size Management** tab.
4. To turn automatic purging on or off, click **Automatically purge the repository periodically**.

   Turn this option off if purging is configured as part of your scheduled full or incremental backups.
5. Click purge options to turn them on or off:
   - **Purge statistics**

- • **Purge alert history**

6. In **Purge data older than**, enter the number of days after which to purge repository data.
7. Click **Apply**, then **OK**.

See also
- *Scheduling Backups of the Repository* on page 190
- *Modifying the Backup Schedule* on page 191
- *Forcing an Immediate Backup* on page 192
- *Restoring the Repository from Backups* on page 192

Logging

Logging helps Sybase Control Center administrators identify and track errors and other system events by recording messages about the events in log files.

Sybase Control Center maintains these logs:

- • The client log – captures messages about activities in the browser-based client components. These messages are generated by the component product modules to display information that is pertinent to the user but not critical enough to warrant a pop-up. Sybase also uses the client log to trace client browser operations.
- • Server logs – capture messages about activities during the initialization sequence, such as starting services; auditing messages recording logins and logouts; errors such as missed scheduled events; and other events on the server. Server logs include:
  - • Component logs, which record only events concerning individual product modules
  - • The SCC agent log, which is a composite log. In an SCC server, the agent log records events in all product modules and in the Sybase Control Center framework. In an SCC agent, the agent log records events in the agent.
- • The repository log – captures information about inserts and updates that have occurred in the Sybase Control Center repository, a SQL Anywhere database. This log is in `SCC-3_2\log\repository.log`.
- • The alert services log – captures information about alert service status and events, including execution of alert-triggered scripts (start time, end time, and status and exit codes). This log is in `SCC-3_2\log\alert-server.log`.

Viewing the Sybase Control Center for Sybase IQ Log
View event logs for Sybase Control Center for Sybase IQ.

Sybase Control Center for Sybase IQ uses Log4J for message logging. The Sybase Control Center for Sybase IQ log files are located at:

- • Windows – `%SYBASE%\SCC-3_2\plugins\IQMAP\log\iqmap.log`
- • UNIX – `$SYBASE/SCC-3_2/plugins/IQMAP/log/iqmap.log`

---

1. Display the log file using a log viewer or another method of your choice.

2. Look for entries of interest such as login attempts or the failure of a service to start.

See also
- *Modifying the Sybase IQ Log Configuration* on page 196
- *Viewing Sybase Control Center Server Logs* on page 196
- *Viewing the Sybase Control Center Client Log* on page 197
- *Changing the Logging Level* on page 197
- *Logging or Message Levels* on page 198
- *Changing Logging Configuration* on page 199

Modifying the Sybase IQ Log Configuration

Change the log level or logging configuration settings for Sybase Control Center for Sybase IQ.

1. Navigate to `%SYBASE%\SCC-3_2\plugins\IQMAP`.

2. Open the `IQMapLog4j.properties` file, and modify the settings as needed.

3. Save and close the `IQMapLog4j.properties` file.

4. Restart the SCC server.

Messages related to SCC for Sybase IQ are recorded on the console and the `iqmap.log` file. The `iqmap.log` file is located in `%SYBASE%\SCC-3_2\plugins\IQMAP\log`.

See also
- *Viewing the Sybase Control Center for Sybase IQ Log* on page 195
- *Viewing Sybase Control Center Server Logs* on page 196
- *Viewing the Sybase Control Center Client Log* on page 197
- *Changing the Logging Level* on page 197
- *Logging or Message Levels* on page 198
- *Changing Logging Configuration* on page 199

Viewing Sybase Control Center Server Logs

View event logs for the Sybase Control Center server.

Sybase Control Center logs events to several places:

- The console from which Sybase Control Center is launched.
- The Sybase Control Center agent log: `<SCC-install-directory>/log/agent.log`
- The repository log: `<SCC-install-directory>/log/repository.log`

- The component log for each installed Sybase Control Center product module. The path to the component log takes this form: `<SCC-install-directory>/plugins/<component>/log/<component>.log`

1. Display one of the log files using a log viewer or a method of your choice.
2. Look for entries of interest such as login attempts or the failure of a service to start.

   On the console and in the Sybase Control Center agent log file, some components prepend the component name to log entries.

See also
- *Viewing the Sybase Control Center for Sybase IQ Log* on page 195
- *Modifying the Sybase IQ Log Configuration* on page 196
- *Viewing the Sybase Control Center Client Log* on page 197
- *Changing the Logging Level* on page 197
- *Logging or Message Levels* on page 198
- *Changing Logging Configuration* on page 199

Viewing the Sybase Control Center Client Log

Display the event log for the current session of your Sybase Control Center browser client.

In the perspective tab window (the main window), do either of the following to display the client log:

- Enter **Ctrl+Alt+L**.
- Select **View > Open > Log Window.**

Note: The client log reader displays the 100 most recent log messages for the current login session.

See also
- *Viewing the Sybase Control Center for Sybase IQ Log* on page 195
- *Modifying the Sybase IQ Log Configuration* on page 196
- *Viewing Sybase Control Center Server Logs* on page 196
- *Changing the Logging Level* on page 197
- *Logging or Message Levels* on page 198
- *Changing Logging Configuration* on page 199

Changing the Logging Level

Adjust the logging level that determines which events Sybase Control Center records in the server logs. This task requires you to restart Sybase Control Center.

If you are having a problem with Sybase Control Center, you might be able to discover the cause of the problem by changing the server logging level so that more events are recorded.

1. Shut down Sybase Control Center.
2. Restart Sybase Control Center using the -m option to change the logging level. In `<SCC-installation-dir>/bin`, enter:

```
scc -m <logging-level>
```

   The logging levels are OFF (logs nothing), FATAL (logs only the most severe events), ERROR, WARN, INFO, DEBUG, and ALL (logs everything).
3. Examine the server log for clues about what might be causing the problem.
4. When you have resolved the problem, set the logging level back to WARN, the default. Your log may become unmanageably large if you leave it at the DEBUG or ALL level.

Example

These commands, which must be executed in the installation directory, start Sybase Control Center with the logging level set to debug:

```
Windows: bin\scc -m DEBUG
UNIX: bin/scc -m DEBUG
```

See also
- *Viewing the Sybase Control Center for Sybase IQ Log* on page 195
- *Modifying the Sybase IQ Log Configuration* on page 196
- *Viewing Sybase Control Center Server Logs* on page 196
- *Viewing the Sybase Control Center Client Log* on page 197
- *Logging or Message Levels* on page 198
- *Changing Logging Configuration* on page 199
- *Starting and Stopping Sybase Control Center in Windows* on page 72
- *Starting and Stopping Sybase Control Center in UNIX* on page 74

Logging or Message Levels
Describes values you can use to control the types of events that are logged by Sybase Control Center.

These are the logging levels, from highest to lowest. The higher the level, the more serious an event must be to be logged. Each level includes all the levels above it—for example, if you set the logging level to WARN, you log events for the WARN, ERROR, and FATAL levels.

OFF	Nothing is logged. This is the highest level.
FATAL	Logs only very severe error events that lead the server to abort. This is the highest level at which events are logged.
ERROR	Logs error events that might allow the server to continue running.
WARN	Logs potentially harmful situations. WARN is the default logging level during normal operation (that is, after system initialization).

INFO	Logs informational messages that track the progress of the server in a coarse-grained fashion. INFO is the default logging level during the system initialization process.
DEBUG	Logs a larger set of events that provides a finer-grained picture of how the server is operating. This level is recommended for troubleshooting.
ALL	Logs all loggable events. This is the lowest level.

See also
- *Viewing the Sybase Control Center for Sybase IQ Log* on page 195
- *Modifying the Sybase IQ Log Configuration* on page 196
- *Viewing Sybase Control Center Server Logs* on page 196
- *Viewing the Sybase Control Center Client Log* on page 197
- *Changing the Logging Level* on page 197
- *Changing Logging Configuration* on page 199
- *scc Command* on page 81

Changing Logging Configuration

Edit the logging configuration file, log4j.properties, to modify Sybase Control Center logging.

You can change the names, locations, or maximum size of the log files as well as the number of log files backed up.

Options for the **scc** command let you change the overall Sybase Control Center log message level when you start SCC, but if you choose the DEBUG level, the large volume of log messages generated may be inconvenient. Editing the log properties file gives you finer control; you can set logging levels for each Sybase Control Center component separately. Sybase recommends making such changes only if you are familiar with log4j and you are working with Sybase technical support; DEBUG-level log messages are not likely to be meaningful to you. (If you have not used log4j before, a good place to start is *http://logging.apache.org/log4j/1.2/manual.html*.)

1. Shut down Sybase Control Center.
2. Make a backup copy of the log4j.properties file located in <SCC-installation-directory>/conf.
3. Open the log4j.properties file for editing.
4. Change values in the file to suit your needs. For example:

To	Modify
Change the name or location of a log file	• Agent log – log4j.appender.agent.File • Repository log – log4j.appender.repository.File • Collection statistics log – log4j.appender.collection-stats.File • Alert server log – log4j.appender.alert.File • Gateway log – log4j.appender.gateway.File
Change the maximum size that a log file can reach before Sybase Control Center creates a new file	• Agent log – log4j.appender.agent.MaxFileSize • Repository log – log4j.appender.repository.MaxFileSize • Collection statistics log – log4j.appender.collection-stats.MaxFileSize • Alert server log – log4j.appender.alert.MaxFileSize • Gateway log – log4j.appender.gateway.MaxFileSize
Change the number of log files Sybase Control Center backs up before deleting the oldest file	• Agent log – log4j.appender.agent.MaxBackupIndex • Repository log – log4j.appender.repository.MaxBackupIndex • Collection statistics log – log4j.appender.collection-stats.MaxBackupIndex • Alert server log – log4j.appender.alert.MaxBackupIndex • Gateway log – log4j.appender.gateway.MaxBackupIndex

5. Save and exit the file.
6. Start Sybase Control Center to make the logging changes take effect.

See also
- *Viewing the Sybase Control Center for Sybase IQ Log* on page 195
- *Modifying the Sybase IQ Log Configuration* on page 196
- *Viewing Sybase Control Center Server Logs* on page 196
- *Viewing the Sybase Control Center Client Log* on page 197
- *Changing the Logging Level* on page 197
- *Logging or Message Levels* on page 198
- *Starting and Stopping Sybase Control Center in Windows* on page 72
- *Starting and Stopping Sybase Control Center in UNIX* on page 74

Sybase Control Center Console

The console is a command-line interface for displaying details about the status of the Sybase Control Center server and its subsystems.

When you use the **scc** command to start Sybase Control Center, it displays start-up messages and then displays the console prompt.

Note: The console prompt does not appear if you start Sybase Control Center as a service, if you direct the output of **scc** to a file, or if you start Sybase Control Center in the background.

See also
- *Launching Sybase Control Center* on page 70

Console Commands
Use the Sybase Control Center console to get status information on Sybase Control Center and its ports, plug-ins, and services.

help Command
Display syntax information for one or more Sybase Control Center console commands.

Syntax
```
help [command_name]
```

Parameters

- **command_name** – optional. status, info, or shutdown. If you omit *command_name*, **help** returns information on all the console commands.

Examples

- **Example 1** – returns information on the **status** command:

```
help status
```

Permissions

help permission defaults to all users. No permission is required to use it.

See also
- *info Command* on page 201
- *shutdown command* on page 202
- *status Command* on page 203

info Command
Display information about specified parts of the Sybase Control Center server.

If you enter **info** with no parameters, it returns information for every parameter.

Syntax
```
info [-a | --sys]
[-D | --sysprop [system-property]]
[-e | --env [environment-variable]]
```

```
[-h | --help]
[-m | --mem]
[-p | --ports]
[-s | --services]
```

Parameters

- **-a | --sys** – optional. List all the services known to Sybase Control Center, indicate whether each service is enabled, and list other services on which each service depends.
- **-D | --sysprop [***system-property***]** – optional. Display information about the specified Java system property. Omit the system-property argument to return a list of all Java system properties and their values.
- **-e | --env [***environment-variable***]** – optional. List all the environment variables in the Sybase Control Center Java VM process environment. Omit the environment-variable argument to return a list of environment variables and their values.
- **-h | --help** – optional. Display information about the **info** command.
- **-m | --mem** – optional. Display information about the server's memory resources.
- **-p | --ports** – optional. List all the ports on which the Sybase Control Center agent and its services listen, indicate whether each port is in use, and show the service running on each port.
- **-s | --services** – optional. List all Sybase Control Center services, indicate whether each service is enabled, and list other services on which each service depends.

Examples

- **Example 1** – displays information about ports on this Sybase Control Center server:

```
info -p
```

Permissions

info permission defaults to all users. No permission is required to use it.

See also

- *help Command* on page 201
- *shutdown command* on page 202
- *status Command* on page 203

shutdown command
Stop the Sybase Control Center server if it is running.

Syntax

```
shutdown
```

Examples

- **Example 1 –** shuts down Sybase Control Center:

  ```
  shutdown
  ```

Permissions

shutdown permission defaults to all users. No permission is required to use it.

See also
- *help Command* on page 201
- *info Command* on page 201
- *status Command* on page 203

status Command
Display the status of the SCC agent, plug-in, or service components of Sybase Control Center.

Syntax

```
status [-a | --agent]
[-h | --help]
[-p | --plugin [plugin-name]]
[-s | --service [service-name]]
```

Parameters

- **-a | --agent –** display the status of the Sybase Control Center agent component.
- **-h | --help –** display information about the **info** command.
- **-p | --plugin [*plugin-name*] –** display the status of the specified Sybase Control Center plug-in (for example, ASEMap, the Adaptive Server® management module). Omit the plugin-name argument to return a list of plug-ins.
- **-s | --service [*service-name*] –** display the status of the specified Sybase Control Center service (for example, the Alert service or the Messaging service). Omit the service-name argument to return a list of services.

Examples

- **Example 1 –** displays status information on the Repository service:

  ```
  status --service Repository
  ```

Permissions

status permission defaults to all users. No permission is required to use it.

See also
*   *help Command* on page 201
*   *info Command* on page 201
*   *shutdown command* on page 202

Manage and Monitor the Sybase IQ Environment

Manage and monitor all single-node and multiplex servers in the Sybase IQ environment.

Monitor a Simplex Server

Statistics allow you to monitor the availability and performance of a simplex server.

Viewing Overview Statistics

Display high-level statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **Overview**.
3. Select the tab for the required information.

Note: Click a column header to sort the data by that column.

To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Tab	Description
Server	**State** – current status of the server. Valid states include: • Unknown • Stopped • Running **Host** – host name where the server is running. **Port** – port number where the server is running. **Server name** – name of the server. **Database** – name of the Sybase IQ database. **Server type** – type of server the database is on. Server types include Single Server, Coordinator, Reader, and Writer. **Server version** – version of the Sybase IQ server. **Platform** – operating system running on the server host.
Activities	**CPU total usage (%)** – total CPU usage percentage, including both system and user usage. **Active connections** – total number of active connections, including user and internode communication connections. **Connections available** – number of connections available for users and internode communication connections. **Active requests** – number of active requests on the server. **IQ threads in use** – number of threads being used by the Sybase IQ server. **Active transactions** – number of active transactions. **Number of committed transactions** – number of committed transactions. **Oldest transaction (minutes)** – elapsed age, in minutes, of the oldest transaction.
Caches	**Catalog cache reads (per second)** – number of catalog cache page lookups per second. **Main cache size (MB)** – size of the main cache, in megabytes. **Temp cache size (MB)** – size of the temporary cache, in megabytes. **Remaining heap size (MB)** – size of the remaining heap allocation, in megabytes.

Tab	Description
Version usage	**Number of committed versions** – the number of table versions in the server.
	Total version space used (MB) – total space consumed by all the table versions.
	Oldest version ID – the oldest version identifier on the server.
	Number of active versions – total number of active write table versions on the server.
	Total active version space created (MB) – amount of data created by active write transactions.
	Total active version space to be destroyed (MB) – amount of data destroyed by active write transactions. If these transactions commit, the destroyed data becomes an old version and is eventually dropped. If the transactions roll back, the created data is released.
Details	**Server full version** – version of the IQ server software, including the date and time.
	Platform version – version of the operating system installed on the server host.
Alerts	Any alerts for the selected server. While the montior is open, alerts are displayed as they are created.
CPU history chart	Percentage of total CPU usage over a period of time.
IQ memory chart	Allocation of the IQ memory between the main cache, temporary cache, and remaining heap.
Disk usage chart	Available and used space for the main and temporary stores.

See also

Viewing All Statistics

Display all the statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **All statistics**.
   For each group of statistics, expand the heading to list the individual key performance indicators (KPIs) in that group. The list shows:
   - Name of the KPI
   - Current value of the KPI
   - Unit of the value
   - Description of the KPI

See also

Viewing Engine Statistics

Display the engine statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **Engine**.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Engine statistics table	Displays the statistics for the engine. The calculations are continuously updated based on live server information. The engine statistics include: **CPU total usage** – Percentage of CPU total usage. **CPU system usage** – Percentage of CPU system usage. **CPU user usage** – Percentage of CPU user usage. **Total memory allocated (MB)** – Total amount of memory (in megabytes) allocated for the main cache, temporary cache, and remaining heap. **Max memory allocated (MB)** – Maximum amount of memory (in megabytes) allocated for the main cache, temporary cache, and remaining heap. **Main cache (MB)** – Total size of the main cache, in megabytes. **Temp cache (MB)** – Total size of the temporary cache, in megabytes. **IQ threads in use** – Number of IQ threads in use. **IQ threads available** – Number of available IQ threads.
IQ memory chart	Shows the allocation of the IQ memory between the main cache, temporary cache, and remaining heap.
CPU history chart	Displays the percentage of total, system, and user CPU usage over a period of time.

See also

Viewing Connection Statistics

Display the connection statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.

2. In the left pane of the IQ Node Level Monitor view, select **Connections**.

3. Select the tab for the required information.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Connections tab > Active connections	Displays all users currently connected to the server, including: **User ID** – the ID of the connected user. **Name** – the name of the connected user. **Creation time** – the date and time the connection was established. **Connection ID** – the ID of the connection. **Client IP address** – the IP address of the client that made the connection. **Connection or cursor** –indicates a connection or an active cursor
Connections tab > Connection details	Select a User ID to display that user's specific connection details, including: **Statement being executed** – the statement executed by the selected user. **Last request time** – the last time a request was executed on the server by the selected user. **Request type** – type of request executed by the selected user. **Last command time** – the last time a command was executed on the server by the selected user. **Command type** – type of command executed by the selected user. **IQ temporary store usage (KB)** – number of temporary store kilobytes used during the connection. **IQ temporary work space usage (KB)** – number of temporary workspace kilobytes used during the connection. **Cursor count** – number of open Sybase IQ cursors on the connection. **Thread in use** – number of threads in use by the selected user.
Connections tab > Associated transactions	Displays the transaction ID, creation time, and state of transactions executed by the selected user.

Area	Description
Connection Statistics tab > Active connections	Displays a chart of the number of user, internode incoming, and other connections to the server.
Connection Statistics tab > User connections/disconnections per minute	Displays the number of user connections and disconnections per minute.

See also

<u>**Viewing Transaction Statistics**</u>
Display transaction statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **Transactions**.
3. Select the tab for the required information.

Note: Click a column header to sort the data by that column.

To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Transactions tab > Transactions & versions	Displays all transactions that are currently on the server, and their version details.

Area	Description
Transactions tab > Transaction details	Transaction details include:
	Transaction ID – the unique identification number for the selected transaction.
	Connection ID – the connection identification number for the selected transaction.
	Statement being executed – the statement executed by the selected transaction.
	IQ main store space created by transaction (KB) – the amount of main store space created by the selected transaction.
	IQ main store space dropped by transaction (KB) – the amount of main store space dropped by the selected transaction.
	IQ temporary store space created by transaction (KB) – the amount of temporary store space created by the selected transaction.
	IQ temporary store space dropped by transaction (KB) – the amount of temporary store space created by the selected transaction.
	Cursor count – number of open Sybase IQ cursors on the transaction.
	IQ threads – number of threads being used by the transaction.
	IQ govern priority – numeric priority of the transaction in the queue.
	Connection or cursor – identifies whether the transaction is a connection or a cursor.
	Connection or cursor create time – date and time the connection or cursor was created.
Transaction statistics tab > Active transactions	Displays a chart of the number of user, internode communication (identified on the chart as INC), and other transactions on the server.

Area	Description
Transaction statistics tab > Other statistics	Displays additional details on the transaction, including: **Number of committed transactions** – total number of committed transactions on the server. **Number of active load statements** – total number of active load statements on the server. **Oldest active transaction time (minutes)** – elapsed time, in minutes, since the oldest active transaction's creation.

See also

Viewing Store I/O Statistics

Display the store I/O statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **Store I/O**.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Disk Reads	Number of disk reads per second on the catalog store, main store, and temporary store.
Disk Writes	Number of disk writes per second on the catalog store, main store, and temporary store.

See also
*   *Viewing Overview Statistics* on page 204
*   *Viewing All Statistics* on page 207
*   *Viewing Engine Statistics* on page 207
*   *Viewing Connection Statistics* on page 208
*   *Viewing Transaction Statistics* on page 210
*   *Viewing Cache Statistics* on page 213
*   *Viewing Table Version Statistics* on page 214
*   *Viewing Operations and Requests Statistics* on page 216
*   *Viewing Network Statistics* on page 217

Viewing Cache Statistics

Display the cache statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **Caches**.
3. Select the tab for the required information.

Note: Hover the mouse pointer over any line or bar graph to display information for that graph.

Tab	Description
Cache size	**Catalog cache** – the megabyte allocation for the catalog cache, and the number of megabytes in use.
	Main cache – the megabyte allocation for the main cache, and the number of megabytes in use.
	Temporary cache – the megabyte allocation for the temporary cache, and the number of megabytes in use.
Cache reads	Number of cache reads per second for a period of time. The cache reads for the catalog, main, and temporary caches appear.

Tab	Description
Cache statistics	Displays the cache statistics for the catalog, main, and temporary caches. Each cache type includes statistics for:
	Size (mb) – total size of the cache, in megabytes.
	In use (%) – percentage of the cache being used.
	Reads (per second) – number of cache reads per second.
	Hits (per second) – number of hits, per second, to the the catalog, main, and temporary caches.
	Dirty pages (%) – percentage of pages in the catalog, main, and temporary caches where data has been modified and stored in the buffer cache and has not yet been written to disk.
	Pinned – number of pinned catalog, main, and temporary cache pages.
	Pinned (%) – percentage of the catalog, main, and temporary chaches pinned.

See also

Viewing Table Version Statistics
Display the table version statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.
2. In the left pane of the IQ Node Level Monitor view, select **Table Versions**.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Area	Description
Statistics	**Number of committed versions** – the number of table versions in the server.
	Total version space used (MB) – total space consumed by all the table versions.
	Oldest version ID – the oldest version identifier on the server.
	Number of active versions – total number of active write table versions on the server.
	Total active version space created (MB) – amount of data created by active write transactions.
	Total active version space to be destroyed (MB) – amount of data destroyed by active write transactions. If these transactions commit, the destroyed data becomes an old version and is eventually dropped. If the transactions roll back, the created data is released.
Table versions	**Version ID** – the table version identifier.
	Server name – the name of the server.
	Connection ID – the connection ID using this table version.
	MinKBRelease – the minimum amount of space returned once this version is no longer in use.
	MaxKBRelease – the maximum amount of space returned once this version is no longer in use.
	WasReported – indicates whether the server has received usage information for this version.

See also

- *Viewing Operations and Requests Statistics* on page 216
- *Viewing Network Statistics* on page 217

Viewing Operations and Requests Statistics

Display the operation and request statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.

2. In the left pane of the IQ Node Level Monitor view, select **Operations & Requests**.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Opera-tions	**Total operations** – the total number of IQ operations of any type. **Active operations** – the number of active IQ operations. **Waiting operations** – the number of IQ operations waiting for the resource governor.
Requests	**Requests** – the number of times per second the server has been accessed to handle a new request or continue processing an existing request. **Active requests** – the number of active requests. **Unscheduled requests** – the number of requests that are currently in the queue, waiting for an available server thread.

See also
- *Viewing Overview Statistics* on page 204
- *Viewing All Statistics* on page 207
- *Viewing Engine Statistics* on page 207
- *Viewing Connection Statistics* on page 208
- *Viewing Transaction Statistics* on page 210
- *Viewing Store I/O Statistics* on page 212
- *Viewing Cache Statistics* on page 213
- *Viewing Table Version Statistics* on page 214
- *Viewing Network Statistics* on page 217

Viewing Network Statistics

Display the network statistics for the selected Sybase IQ server.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Node**.

2. In the left pane of the IQ Node Level Monitor view, select **Network**.

3. Select the tab for the required information.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Network us-age	**Bytes received** – amount of data, in bytes, received by the server.
	Bytes received uncompressed – amount of uncompressed data, in bytes, received by the server.
	Bytes sent – amount of data, in bytes, sent by the server.
	Bytes sent uncompressed – amount of uncompressed data, in bytes, sent by the server.
	Free communication buffers – number of free communication buffers.
	Total communication buffers – total number of communication buffers.
Data transfer history	Displays the amount of data, in kilobytes, sent and received by the server over time.
Buffer usage	Shows the total number of communication buffers, and a graph displaying the number of used and free communication buffers.

See also
- *Viewing Overview Statistics* on page 204
- *Viewing All Statistics* on page 207
- *Viewing Engine Statistics* on page 207
- *Viewing Connection Statistics* on page 208
- *Viewing Transaction Statistics* on page 210
- *Viewing Store I/O Statistics* on page 212
- *Viewing Cache Statistics* on page 213
- *Viewing Table Version Statistics* on page 214
- *Viewing Operations and Requests Statistics* on page 216

Manage a Server

Create, start, or stop a server, view properties, change server configuration, or generate administration scripts.

See also
*   *Converting a Simplex Server to Multiplex* on page 260

Editing Server Properties

Change general, configuration, and agent information, and options and values of properties for the selected server.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA authority.

Task

1. In the Perspective Resources view, select a resource, then select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select either:

    *   **IQ Servers**, or
    *   **IQ Servers > Multiplex Management > Multiplex Servers**
3. Select the resource in the right pane and either:

    *   Click the drop-down arrow to the right of the name and select **Properties**, or,
    *   From the Administration Console menu bar, select **Resource > Properties**.
4. Edit server properties.

    Note: Configuration changes may take several minutes to complete. When you make a configuration change, a status dialog box displays the progress of the change.

Area	Description
General page	**Server name** – Name of the server. **Host** – Host name where the server is running. **Port** – Port where the server is running. **Database** – Database name. **Server version** – Server version number. **Platform** – Platform the server is running on. **Platform version** – Version of the operating system where the server is running. **Type** (simplex server) – Server type: Single Server. **Type** (multiplex node) – Server type: Reader, Writer, or Coordinator. **Status** (multiplex node) – Server status: Include or Excluded. **Designated Failover Node** (multiplex node) – True or False. If true, then this server is the designated failover node.
Configuration page (multiplex node only)	**Server name** – Name of the server. **Public Host/Ports** (multiplex node) – Public host names and port numbers. **Private Host/Ports** (multiplex node) – Information about private hosts and ports to be used. **Database file path** – Location of the .db file.
Options page (multiplex node only)	**Current time** – The current time. **Refresh** – Click Refresh to update the current time. **Quitting time** – Enter a time for the database server to shut down. Use the same format as the current time: YYYY-MM-DD HH:NN:SS.SS **Disable new connections** – Prevent other users from connecting to the database. This may be useful for some maintenance operations. **Remember last statement** – Instruct the database server to capture the most recently prepared SQL statement for each connection to a database on the server.

Area	Description
Agent page	**Agent registered** – Indicates if the Sybase Control Center agent is registered: true or false.
	Agent authenticated – Indicates if the Sybase Control Center agent is authenticated: true or false.
	Agent status – Status of the Sybase Control Center agent: Running, Stopped, or Unknown.
	Agent host – Name of the host machine where the Sybase Control Center agent is running.
	Agent port – Port number on the host machine where the Sybase Control Center agent is running.
	Agent user – User name for authentication of the agent. Default is *uafadmin*.
	Agent process owner – The user name that owns the agent process.
	Agent home – The home directory of the Sybase Control Center agent.
	Agent version – The version of the Sybase Control Center agent.
	SCC agent plugin version – The agent plugin version of the Sybase Control Center agent.
	IQ directory – Installation directory of the IQ server with which the Sybase Control Center agent is associated.
	IQ version – Version of the IQ server with which the Sybase Control Center agent is associated.
Server Properties page	**Properties** – The name, value, and description of all server properties.

See also
- *Starting a Server* on page 221
- *Stopping a Server* on page 221
- *Editing the Server Configuration File* on page 222
- *Generating Administration Scripts* on page 228

Starting a Server

Start a simplex server or multiplex node.

Prerequisites

- The Sybase IQ server is stopped.
- The Sybase Control Center agent is registered, authenticated, and running on the server.

Task

1. In the Perspective Resources view, select a resource, then select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select either:

    - **IQ Servers**, or
    - **IQ Servers > Multiplex Management > Multiplex Servers**
3. Select the single server or multiplex node from the right pane and either:

    - Click the drop-down arrow to the right of the name and select **Start Server**, or,
    - From the Administration Console menu bar select **Resource > Start Server**.
4. Click **Finish**.

See also
- *Editing Server Properties* on page 218
- *Stopping a Server* on page 221
- *Editing the Server Configuration File* on page 222
- *Generating Administration Scripts* on page 228

Stopping a Server

Stop a simplex server or multiplex node.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- The Sybase IQ server is authenticated and running.

Task

1. In the Perspective Resources view, select a resource, then select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select either:

- **IQ Servers**, or,
- **IQ Servers > Multiplex Management > Multiplex Servers**

3. Select the single server or multiplex node from the right pane and either:

- Click the drop-down arrow to the right of the name and select **Stop Server**, or,
- From the Administration Console menu bar select **Resource > Stop Server**.

4. Click **Finish**.

See also
- *Editing Server Properties* on page 218
- *Starting a Server* on page 221
- *Editing the Server Configuration File* on page 222
- *Generating Administration Scripts* on page 228

Editing the Server Configuration File

Change server configuration settings including administrative, memory, connection, and debugging settings.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- The Sybase Control Center agent for the server is running and authenticated.

Task

Editing the configuration file edits the `params.cfg` file located in the database directory for the server. You cannot edit a custom configuration file. If `params.cfg` does not exist, the Sybase Control Center agent automatically generates it when you adjust any configuration values and click **OK** or **Apply**.

Important: The Start Server wizard starts the server using the `params.cfg` file. You cannot start a server using a custom configuration file with a different name.

1. In the Perspective Resources view, select a resource, then select **Resource > Administration Console**.

2. In the left pane of the Administration Console, select either:

- **IQ Servers**, or,
- **IQ Servers > Multiplex Management > Multiplex Servers**

3. Select a server from the right pane and either:

- Click the drop-down arrow to the right of the name and select **Edit Configuration File**, or,

- From the Administration Console menu bar select **Resource > Edit Configuration File**.

The Config file editor window appears.

4. Adjust the configuration values and click **Apply**.

Area	Description
Admin	**Disable triggers** – Disable firing of triggers.
	Checkpoint timeout period – Set the maximum length of time, in minutes, that the database server runs without doing a checkpoint.
	Maximum recovery time – Set the maximum length of time, in minutes, that the database server takes to recover from system failure.
	Start database permission – Specify permission required to start the database: "dba", "all" or "none".
	Stop database permission Specify permission required to stop the database: "dba", "all" or "none".
	Load/unload permission – Set LOAD/UNLOAD permission to "dba", "all" or "none".
	Utility permission – Set utility commands (e.g., DROP DATABASE) permission to "utility_db", "dba", "all" or "none".

Area	Description
Memory	**Display cache sizing statistics** – Displays cache size changes.
	Disable automatic cache resizing – Enforces a static cache size. Disables automatic cache resizing.
	Initial cache size – Sets the initial memory reserved for caching database pages and other server information. The size is the amount of memory. Select **Kilobytes**, **Megabytes**, **Gigabytes**. Select **%** to specify a percentage either of the physical system memory, or of the maximum non-AWE cache size, whichever is lower.
	Minimum cache size – Sets a minimum cache size as a lower limit to automatic cache resizing. The size is the amount of memory, in bytes. Select **Kilobytes**, **Megabytes**, **Gigabytes**. Select **%** to specify a percentage either of the physical system memory, or of the maximum non-AWE cache size, whichever is lower.
	Maximum cache size – Sets a maximum cache size, as a limit to automatic cache growth. The size is the amount of memory, in bytes. Select **Kilobytes**, **Megabytes**, **Gigabytes**. Select **%** to specify a percentage either of the physical system memory, or of the maximum non-AWE cache size, whichever is lower.
	Engine thread stack size – Sets server thread stack size. The size is the amount of memory, in bytes. Select **Kilobytes** or **Megabytes**.
	Number of engine threads – Sets the number of execution threads used for the catalog store and connectivity while running with multiple users. Recommended value is 1.5 times the maximum number of concurrent connections to the server; minimum of 25.
	Number of concurrent OS threads – Sets maximum number of physical processors to use (up to licensed maximum).
	External DLL thread stack size – Sets the stack size for threads running external functions, in bytes. The default is 32 KB.
	Maximum page size – Sets the maximum page size in KB.

Area	Description
Connections	**Shutdown after last database closes** – Automatically shut down after the last database is closed.
	Encrypt communication messages – Enable packet encryption on the network server.
	Enable client-server character translation – Character set translation is turned on by default. There is a performance cost associated with character set translation. If you can set up an environment such that no character set translation is required, then you do not have to pay this cost, and your setup is simpler to maintain.
	Maximum connections – Specifies the maximum number of concurrent user connections.
	Communication packet size – Specifies a maximum packet size. If you want to send small amounts of data over the network, keep the default network packet size small. The default is 512 bytes.
	Idle time before disconnect – Set the amount of client idle time before the connection is terminated. If a client runs for the idle timeout period without submitting a request, the connection is severed.
	Liveness timeout – A liveness packet is sent across a client/server to confirm that a connection is intact. If the client runs for the liveness timeout period without detecting a liveness packet, the communication will be severed. This parameter works only with network server and TCP/IP communications protocols. The default is 120 seconds.
	Quitting time – Lets you specify a time when the database server is to shut down.
	Broadcast level – Specifies how the server reacts to broadcasts. **(Ignore All)** causes the server not to start up any UDP broadcast listeners. **(dblocate)** causes the server to not respond to broadcasts from dblocate, while leaving connection logic unaffected.
Database	**Set database read-only** – Forces all databases that start on the database server to be read-only. No changes to the databases are allowed: the database server doesn't modify the database files or transaction log files.
	Truncate xact log after checkpoint – Causes the transaction log to be truncated after each checkpoint for all databases.

Area	Description
Debug	**Generate debug information** – Display debugging information. **Remember last statement on each connection** – Instruct the database server to capture the most recently prepared SQL statement for each connection to a database on the server. **Debug level** – Enables request-level logging of operations. **Debug output file name** – Redirects HTTP Web service client procedure debug log to a file. **Max output file size** – Specify maximum size of file for server request logging.

Area	Description
IQ	**Number of CPUs IQ can use** – Specifies the number of CPUs available to IQ, overriding the physical number of CPUs for resource planning purposes. The value defaults to the total number of CPUs, but the range of available values is 1 – 128.
	Number of threads – Sets the number of execution threads that will be used for the catalog store and connectivity while running with multiple users. This parameter applies to all operating systems and servers. Each connection uses a thread for each request, and when the request is completed, the thread is returned to the pool for use by other connections. As no connection can have more than one request in progress at one time, no connection uses more than one thread at a time.
	Thread stack size – Specifies the stack size, in KB, for server execution threads running either in the background or as part of a thread team assisting the main server connection thread. The default is 512KB on 64-bit platforms, and 200KB on 32-bit platforms.
	Wired memory pool size – Pool of wired memory on HP and Sun UNIX systems. This memory is locked down so it cannot be paged by the operating system. Specify the memory size, in MB. Use this switch only if you have enough memory to dedicate for this purpose. Otherwise, you may cause serious performance degradation.
	Main buffer cache size – Specifies the main IQ store cache size in MB. Always specify the value for the size, but no units of measurement; for example specify 32 instead of 32MB.
	Temporary buffer cache size – Specifies IQ temporary store cache size in MB. Always specify the value for the size, but no units of measurement; for example specify 32 instead of 32MB.
	Number of concurrent queries – The number of concurrent queries is not the same as the number of connections. This setting can help Sybase IQ optimize paging of buffer data out to disk and avoid overcommitting memory. The default value of this switch is equal to 2 times the number of CPUs on your machine, plus 10. You may find that another value, such as 2 times the number of CPUs plus 4, provides better throughput, especially when large numbers of users are connected.
	Number of partitions for buffer cache – Specifies the number of partitions in the IQ main and temp buffer caches. Must be a power of 2. Allowed values are: 0 (default), 1, 2, 4, 8, 16, 32, 64. By default,

Area	Description
	IQ computes the number of partitions automatically as number_of_cpus/8, rounded to the nearest power of 2, up to a maximum of 64. You may be able to improve performance by adjusting the number of cache partitions.
	Force recovery on database – Open database in forced recovery mode.
Misc	**Quiet mode** – Runs Interactive SQL in quiet mode.
	Run as a daemon – Using this option lets you run the server so that it continues running after the current user session ends.
	Use buffered disk I/O – Uses buffered disk I/O [Windows, UNIX].
	Syslog facility ID – Syslog facility ID. The default is user. Either none, user, daemon, local0,...,local7).
	Output message file name – Filename for copy of message window. File is truncated first.
	Output message file size – Appends .old to the log file name and starts a new file with the original name when log reaches the specified size.
	Touch temporary file timer – (UNIX servers) Causes the server to touch catalog store temporary files at intervals specified in minutes.
	User specified – Enables advanced users to specify configuration options not shown in the Config file editor window. Enter multiple parameters using the `space` character as a separator.

See also
- *Editing Server Properties* on page 218
- *Starting a Server* on page 221
- *Stopping a Server* on page 221
- *Generating Administration Scripts* on page 228

Generating Administration Scripts
Generate the configuration file and scripts for starting, stopping, and synchronizing servers.

Prerequisites
A Sybase Control Center agent is running on the Sybase IQ server.

Task

1. In the Perspective Resources view, select a resource, then select **Resource > Administration Console**.

2. In the left pane of the Administration Console, select either:

   - **IQ Servers**, or,
   - **IQ Servers > Multiplex Management > Multiplex Servers**

3. Select the single server or multiplex node from the right pane and either:

   - Click the drop-down arrow to the right of the name and select **Generate Administration Scripts**, or,
   - From the Administration Console menu bar select **Resource > Generate Administration Scripts**.

These scripts are generated:

- `params.cfg`
- `start_server.sh / start_server.bat`
- `stop_server.sh / stop_server.bat`
- `sync_server.sh (.bat)` – multiplex server only

See also
- *Editing Server Properties* on page 218
- *Starting a Server* on page 221
- *Stopping a Server* on page 221
- *Editing the Server Configuration File* on page 222

Manage a Database

Create a database and view database options and properties.

DBA authority is required for all database management tasks.

Note: A Sybase IQ server resource cannot have multiple databases. Attempting to manage, monitor, or administer a Sybase IQ server resource connected to multiple databases may cause unexpected results.

Viewing Database Properties
View editable and noneditable properties for the database.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA authority.

Task

1. In the Perspective Resources view, select the resource and select **Resource >
   Administration Console**.
2. In the left pane of the Administration Console, select **IQ Servers > Schema Objects >
   Databases**.
3. Select the database from the right pane and either:

   • Click the drop-down arrow to the right of the name and select **Properties**, or,
   • From the Administration Console menu bar, select **Resource > Properties**.

Area	Description
General	**Database** – Name of the database.
	ID – A unique number assigned by the server to each database that is started on it. The ID number lets you distinguish between databases running on the same server.
	Capability ID – The capability bits enabled for the database. You can view a list of the capabilities that are enabled for the database on the Extended Information tab of the Database property sheet.
	Java location – Path to the JRE executable file.
	Java main user ID – The main Java user ID.
	Page size – Page size of the database, in bytes.
	Database file – Root database file for the database.
	Log file – Name and location of the transaction log file for the database.
	Mirror log file – Name and location of the mirror log file for the database.
	Temporary file – Location of the temporary database file.
	Current user – User ID of the user connected to the database.
	Connection ID – Connection ID for the database connection from Sybase Control Center.
	Connection name – Connection name for the user connected to this database. Naming your connections allows multiple connections to the same database, or multiple connections to the same or different database server, to be easily identified.
	Communication link – Type of communications link used by the user's connection. If the connection is between a client and network server, the link type represents the network protocol being used.
	Total connections – Total number of current connections to the database from all users, including the Sybase Control Center connection.

Area	Description
Settings	**Encryption type** – Type of database encryption: Simple or AES. If encryption is not supported in the database, the value is None.
	Encryption scope – Scope of the encryption.
	Ignore trailing blanks – Database ignores trailing blanks in comparisons.
	CHAR collation sequence – Set of separators for the collation sequence of CHAR values.
	CHAR collation character set encoding – Character set used for CHAR values.
	CHAR case sensitivity – CHAR values in the database are case-sensitive. This property applies to the data in the database, and to passwords, but not to table names, column names, and other identifiers.
	NCHAR collation sequence – Separators for the collation sequence of NCHAR values.
	NCHAR collation character set encoding – Character set used for NCHAR values.
	NCHAR case sensitivity – NCHAR values in the database are case-sensitive. This property applies to the data in the database, and to passwords, but not to table names, column names, and other identifiers.
	Checkpoint urgency(%) – Time elapsed since the last checkpoint as a percentage of the checkpoint time setting of the database.
	Recovery urgency(%) – An estimate of the amount of time required to recover the database.
	Refresh –Update the checkpoint urgency and recovery urgency values.
	Collect information about deadlocks that occur in this database – Information is collected on deadlocked connections to the database.
	Include SQL statement information for deadlocked connections – SQL statement information is included with the deadlock connection information.
	Clear deadlock information now – Remove all deadlock information from the system.
Database Properties	A list of all database properties, including the property value and the description of the property.

See also
*   *Creating a Database* on page 232
*   *Setting Database Options* on page 235

Creating a Database

Use the Create Databases wizard to create a Sybase IQ database on the same host as the SCC agent.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- The Sybase Control Center agent is running on the host computer.

Task

Successful database creation results in a new, running, and registered Sybase IQ server visible in the resource perspective.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select **IQ Servers** and click the drop-down arrow that appears to the right.
3. Select **Create Databases**.
4. Click **Add**.
5. In the Add Database Definition dialog, specify Sybase Control Center agent and database information:

Area	Description
Host	The name of the host system where the new database is created.
SCC agent port	The port number that the Sybase Control Center agent on the given host is listening on. Default is **9999**.
SCC agent user	User ID for the Sybase Control Center agent. Default is **uafadmin**.
SCC agent password	Password for the Sybase Control Center agent user. Default is no password.
Utility server user	User ID for the utility database. Default is **DBA**.
Utility server password	Password for the utility database. Default is **sql**.
IQ Server Name	Name of the IQ server to be created.
IQ Server Port	Port number for the new IQ server.
IQ user	User ID for the new IQ server.
IQ password	User password for the IQ user.

Area	Description
Database path	Path to the database file. For example, **/<hostname>/sample/sample/my-testdb.db**.
IQ page size	Page size of the database, in kilobytes. If you do not make a selection, the application chooses a page size.
Catalog page size	Database page size for the catalog store, in kilobytes. If you do not make a selection, the application automatically chooses a page size.
IQ main dbspace path	Path to the IQ main store dbspace physical file on disk. For example, **/<hostname>/sample/sample/mytestdb.iq**.
IQ main dbspace size (MB)	(Disabled if Raw Device is selected.) Size, in megabytes, for the IQ main store dbspace. Specify at least **100**.
IQ main dbspace reserve (MB)	The amount of space, in megabytes, to reserve for future expansion in the IQ main store.
Raw device (IQ main dbspace)	Indicates a raw disk.
Local temporary dbspace path	Path to the temporary IQ store physical file on disk. For example: **/<hostname>/sample/sample/mytestdb.iqtmp**.
Local temporary dbspace size (MB)	(Disabled if Raw Device is selected.) Size, in megabytes, for the temporary IQ store.
Local temporary dbspace reserve (MB)	The amount of space, in megabytes, to reserve for future expansion in the temporary IQ store.
Raw device (local temp dbspace)	Indicates a raw disk.
Transaction log	File name of the IQ transaction log that records changes to the database. If you leave this field blank, the application automatically assigns a log name: `<database_name>.log`.
CHAR collation	Click **Select** to open the Select CHAR Collation dialog box. (**Select** is disabled until all required fields are filled.) Select the CHAR collation sequence used by the database to perform alphanumeric sorting. The default is **ISO_BINENG**.
NCHAR collation	Click **Select** to open the Select NCHAR Collation dialog box. (**Select** is disabled until all required fields are filled.) Select the NCHAR collation sequence used by the database to perform alphanumeric sorting on NCHAR data. The default is **UCA**.
Case insensitive for string comparison	Indicates that CHAR and NCHAR values in the database are case-insensitive. This property applies to the data in the database, and to passwords, but not to table names, column names, and other identifiers.

Area	Description
Enable simple encryption on entire database	Indicates simple database encryption is enabled.

6. Click **OK**.

    The Create Database wizard appears. If your database definition values are valid, a check mark appears in the left pane next to Database Definitions. If any input for the database definition is not valid, an X appears. Hover the mouse over the X to see error information.

7. Click **Next**, or, from the left pane of the wizard, select **Execution**.

8. Click **Execute** to start the database creation process.

9. (Optional) Import a list of database definitions from a CSV file.

    a) On the Database Definitions page of the wizard, click **Import** and specify:

Area	Description
File Name	File name of the CSV file containing delimited list of node definitions.
Browse	Opens a file explorer window to locate the CSV file.
Field Delimiter	The default field delimiter is a vertical bar ("|"). If your CSV file uses a different field delimiter, enter it in this field.

    b) Click **OK**.

10. (Optional) Export your existing database definitions to a CSV file for safekeeping or future import:

    a) On the Database Definitions page of the wizard, click **Export** and specify:

Area	Description
Field Delimiter	The default field delimiter is a vertical bar ("|"). You can specify a different field delimiter.

    b) Click **OK**.

11. Click **Finish** to exit the wizard.

See also

- *Viewing Database Properties* on page 229
- *Setting Database Options* on page 235
- *Troubleshooting Invalid Database Definitions* on page 326

Setting Database Options

Adjust the configurable settings of a database to change the way the database behaves or performs.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA authority.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select **Schema Objects > Databases**.
3. Select the database from the right pane and either:
   - Click the drop-down arrow to the right of the name and select **Options**, or,
   - From the Administration Console menu bar, select **Resource > Options**.
4. In the **Setting** column, change a database option value and click **Apply**.

See also
- *Viewing Database Properties* on page 229
- *Creating a Database* on page 232

Manage a Dbspace

Add, modify, view properties, or delete a dbspace.

Viewing Dbspace Properties

Display the properties for the selected dbspace.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. In the right pane, select a dbspace.
4. Click the drop-down arrow that appears next to the dbspace name and select **Properties**.

Property Name	Description
Name	Name of the dbspace.
Type	The dbspace type.
Mode	Mode of the dbspace: • ReadWrite • ReadOnly
Status	Status of the dbspace: • Online • Offline
Striping	Disk striping is enabled (On) or disabled (Off) for this dbspace.
Stripe size	The number of kilobytes (KB) to write to each file before the disk-striping algorithm moves to the next stripe for this dbspace.

See also

Adding a Dbspace

Add a new dbspace to the database using the Create Dbspace wizard.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. Click the drop-down arrow that appears next to **Dbspaces** and select **New**.
4. On the General Details page of the Create Dbspace wizard, specify:

Options	Description
Resource	The name of the resource on which to create the new dbspace.
Name of dbspace to create	The name of the new dbspace.
Store	The IQ store of the new dbspace: • Main store • Catalog store
Striping	Enables or disables disk striping for this dbspace.
Stripe size (Kb)	The number of kilobytes (KB) to write to each file before the disk-striping algorithm moves to the next stripe for this dbspace.

5. In the left pane of the Create Dbspace wizard, select **DB Files** to add files to the dbspace.

6. On the DB Files page, click **Add**.

7. Specify:

Options	Description
Logical name	User-defined name of the DB file.
Path to physical file on disk	Path to the physical file on disk. It is usually best to use an absolute path. A relative path works for simplex servers and for multiplexes that use the SCC shared-disk facility and have the SCC installation directory mounted in the same location on each multiplex server host. If you specify a relative path for a simplex that is later converted to a multiplex, the relative path might cause an error.
Raw device	A check mark indicates the file is on a raw device.
File size	Specify a file size value and units: kilobytes (KB), megabytes (MB), gigabytes (GB), or terabytes (TB). This option is active only if the file is not on a raw device.
Reserve size	Specify a reserve size value and units: kilobytes (KB), megabytes (MB), gigabytes (GB), or terabytes (TB).

8. Click **OK** to add the file.

9. Click **Back** to return to the General Details page to make further changes. When you are finished, click **Finish**.

See also
• *Viewing Dbspace Properties* on page 235

Modifying a Dbspace

Change the name, mode, status, and striping properties of a dbspace.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.

Task

You cannot change a dbspace's name if it is a system dbspace (with a name beginning with IQ_SYSTEM) or a system catalog dbspace.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. In the right pane, select a dbspace to modify.
4. Click the drop-down arrow that appears next to the dbspace name and select **Properties**.
5. On the General page of the properties dialog, specify:

Options	Description
Name	Name of the dbspace.
Type	(Not editable) The dbspace type.
Mode	(Optional) Mode of the dbspace: • ReadWrite • ReadOnly
Status	(Optional) Status of the dbspace: • Online • Offline
Striping	(Optional) Enables (On) or disables (Off) disk striping for this dbspace.
Stripe size	The number of kilobytes (KB) to write to each file before the disk-striping algorithm moves to the next stripe for this dbspace.

6. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also
- *Viewing Dbspace Properties* on page 235
- *Adding a Dbspace* on page 236
- *Changing a Dbspace to Read-Only* on page 239
- *Preallocating Space for a Dbspace* on page 240
- *Deleting a Dbspace* on page 240
- *Generating Dbspace DDL Commands* on page 241

Changing a Dbspace to Read-Only

Change a dbspace to read-only mode.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.

Task

You can add a new DB file to a read-only dbspace. You cannot delete a DB file from a read-only dbspace.

You can change paths of DB files in a read-only offline dbspace.

Note: Changing the file path here does not move the file to the new location. If a file is moved, changing the path here only informs the Sybase IQ server of the new file location.

IQ_SYSTEM_MAIN and IQ_SHARED_TEMP dbspaces cannot be read-only.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. In the right pane, select a dbspace to modify.
4. Click the drop-down arrow that appears next to the dbspace name and select **Properties**.
5. On the General page of the properties dialog, click **ReadOnly** mode.
6. Click Status **Online** or **Offline**.
7. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also
- *Viewing Dbspace Properties* on page 235
- *Adding a Dbspace* on page 236
- *Modifying a Dbspace* on page 238
- *Preallocating Space for a Dbspace* on page 240
- *Deleting a Dbspace* on page 240

• *Generating Dbspace DDL Commands* on page 241

Preallocating Space for a Dbspace
Preallocate space for a catalog dbspace.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA authority.

Task

This operation is valid only for catalog dbspaces.

Sybase IQ automatically increases the size of catalog DB files as additional space is needed. If your database has a high rate of change, you can pre-allocate disk space for dbspaces. Proper pre-allocation reduces fragmentation and improves performance.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. In the right pane, select a dbspace.
4. Click the drop-down arrow that appears next to the dbspace name and select **Pre-allocate Space**.
5. Specify a number and unit, for example, **2 megabytes**.
6. Click **OK**.

See also
• *Viewing Dbspace Properties* on page 235
• *Adding a Dbspace* on page 236
• *Modifying a Dbspace* on page 238
• *Changing a Dbspace to Read-Only* on page 239
• *Deleting a Dbspace* on page 240
• *Generating Dbspace DDL Commands* on page 241

Deleting a Dbspace
Delete a dbspace from the database.

Prerequisites

• Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.
• You must be the only user connected to the database in order to delete a catalog dbspace. You can become the only connected user by dropping all secondary nodes.

- You cannot drop the six initial dbspaces (SYSTEM, TEMPORARY, IQ_SHARED_TEMP, IQ_SYSTEM_MAIN, IQ_SYSTEM_TEMP, and IQ_SYSTEM_MSG). You may drop dbspaces from the IQ main store or catalog store, which may contain multiple dbspaces, as long as at least one dbspace in read-write mode remains.
- The dbspace must contain no user data or user-created objects.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. In the right pane, select a dbspace to delete.
4. Click the drop-down arrow that appears next to the dbspace name and select **Delete**.
5. Click **Yes** to confirm the deletion.

See also
- *Viewing Dbspace Properties* on page 235
- *Adding a Dbspace* on page 236
- *Modifying a Dbspace* on page 238
- *Changing a Dbspace to Read-Only* on page 239
- *Preallocating Space for a Dbspace* on page 240
- *Generating Dbspace DDL Commands* on page 241

Generating Dbspace DDL Commands
Display the data description language SQL code for creating a dbspace. The SQL code can be a useful reference and training tool.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **Dbspaces**.
3. In the right pane, select a dbspace.
4. Click the drop-down arrow that appears next to the dbspace name and select **Generate DDL**.
   The DDL window opens, showing the SQL code used to create the selected dbspace.

See also
- *Viewing Dbspace Properties* on page 235
- *Adding a Dbspace* on page 236

Manage DB Files

Add, modify, view properties, or delete a DB file.

Viewing DB File Properties

Display DB file properties for the selected DB file.

Prerequisites

No privileges required.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **DB Files**.
3. In the right pane, select a DB file to view.
4. Click the drop-down arrow that appears next to the DB file name and select **Properties**.

Property	Description
Name	Logical file name.
Path	Location of the physical file or raw partition.
Mode	Mode of the dbspace: • ReadWrite • ReadOnly
File size	Current size of the file or raw partition in kilobytes (KB) or megabytes (MB).
Reserved size	Reserved space that can be added to this file in the dbspace.
Total size	Current size of the file plus the size of the reserved space.
Modify file size	New value for file size.
Dbspace	Name of the dbspace.

Property	Description
Status	Status of the dbspace: • ONLINE • OFFLINE

See also
* *Adding a DB File* on page 243
* *Modifying a DB File* on page 245
* *Emptying a DB File* on page 246
* *Deleting a DB File* on page 247
* *Generating DB File DDL Commands* on page 248

Adding a DB File

Add a DB file to a dbspace.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.

Task

Follow these steps to add files to all shared dbspaces: IQ_SYSTEM_MAIN, IQ_SHARED_TEMP, and all user-defined dbspaces.

Note: Adding DB files to IQ_SYSTEM_MAIN synchronizes all running secondary nodes.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **DB Files**.
3. Click the drop-down arrow that appears next to **DB Files** and select **New**.
4. On the DB Files page of the wizard, select the resource and the dbspace to which you are adding the DB file.
5. Click **Add**.
6. On the DB File Details page, specify:

Options	Description
Logical name	User-defined name of the DB file.

Options	Description
Path to physical file on disk	Path to the physical file on disk. It is usually best to use an absolute path. A relative path works for simplex servers and for multiplexes that use the SCC shared-disk facility and have the SCC installation directory mounted in the same location on each multiplex server host. If you specify a relative path for a simplex that is later converted to a multiplex, the relative path might cause an error.
Raw device	A check mark indicates the file is on a raw device.
File size	Space you allocate to the DB file in kilobytes (KB), megabytes (MB), gigabytes (GB), or terabytes (TB). Unavailable on raw devices; required if the file is not on a raw device. Minimum file size is 8 MB.
Reserve size	(Optional; default is 0.) Extra space you allocate to the DB file in kilobytes (KB), megabytes (MB), gigabytes (GB), or terabytes (TB). You can use this space to increase the file size if the file grows. Guideline: reserve size should be about 30% of file size.
Mode	(Optional) • Read Only • Read/Write • Force Read/Write **Note:** This option appears only if you are adding DB files to the IQ_SHARED_TEMP dbspace in a multiplex setup.

7. Click **OK** to add the file, then click **Finish**.

Next
Create a DBSpace File Size in Use alert to warn you if the DB file exceeds a set percentage of the configured file size (70%, for example). If the alert is activated, you can increase the file size. To enable this alert, first schedule a data collection that includes its key performance indicator, such as the DBSpace File Statistics Collection or All Stats w/o Availability.

See also

Modifying a DB File

Modify the name, mode, and path properties of a DB file.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.

Task

The mode and status of the DB file determine which properties can be modified.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **DB Files**.
3. In the right pane, select a DB file to modify.
4. Click the drop-down arrow that appears next to the DB file name and select **Properties**.

   Note: If the DB file is read-only in a read-write dbspace, Empty File is enabled on the menu.

5. Change property values:

Property	Description
Name	Logical file name. If the dbspace is in Read Only mode and OFFLINE status, you cannot change the name of the DB file.
Path	Location of the physical file or raw partition. You can modify the path only when the dbspace's status is OFFLINE. It is usually best to use an absolute path. A relative path works for simplex servers and for multiplexes that use the SCC shared-disk facility and have the SCC installation directory mounted in the same location on each multiplex server host. If you specify a relative path for a simplex that is later converted to a multiplex, the relative path might cause an error.
Mode	Mode of the dbspace: • Read Only • Read/Write • Force Read/Write
File size	Current size of the file or raw partition. Use the **Modify file size** field to change the file size.
Reserved size	Reserved space in the dbspace that you can add to this DB file. Use the **Modify file size** field to transfer space between the reserved size and the file size.

Property	Description
Total size	Current size of the file plus the size of the reserved space.
Modify file size	Enter a new value for file size. You cannot increase the file size by more than the reserved size. For example, if the file size is 100 MB and the reserve is 30 MB, the maximum file size is 130 MB.
Dbspace	Name of the dbspace.
Status	Status of the dbspace: • ONLINE • OFFLINE

6. Click **OK**.

See also
- *Viewing DB File Properties* on page 242
- *Adding a DB File* on page 243
- *Emptying a DB File* on page 246
- *Deleting a DB File* on page 247
- *Generating DB File DDL Commands* on page 248

Emptying a DB File

Empty a DB file and move the objects in the file to another available read-write DB file.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- The dbspace must be in read-write mode.
- The DB file should be read-only.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **DB Files**.
3. In the right pane, select a DB file to empty.
4. Click the drop-down arrow that appears next to the DB file name and select **Empty File**.
5. Click **Yes** to continue.

Objects in the DB file are moved to another read-write DB file in the same dbspace.

See also
- *Viewing DB File Properties* on page 242
- *Adding a DB File* on page 243
- *Modifying a DB File* on page 245
- *Deleting a DB File* on page 247
- *Generating DB File DDL Commands* on page 248

Deleting a DB File

Delete a DB file from a dbspace. With the exception of the `IQ_SHARED_TEMP` dbspace, you cannot delete the last DB file in a dbspace.

Prerequisites
- Authenticate with Sybase IQ using an account that has DBA or Space Admin authority.
- The dbspace must be in read-write mode.
- The DB file must be empty.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **DB Files**.
3. In the right pane, select a DB file to delete.
4. Click the drop-down arrow that appears next to the DB file name and select **Delete**.
5. Click **Yes** to confirm the deletion.

See also
- *Viewing DB File Properties* on page 242
- *Adding a DB File* on page 243
- *Modifying a DB File* on page 245
- *Emptying a DB File* on page 246
- *Generating DB File DDL Commands* on page 248

Generating DB File DDL Commands

Display the data description language SQL code for adding a DB file to a dbspace. The SQL code can be a useful reference and training tool.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Space Management** and select **DB Files**.
3. In the right pane, select a DB file.
4. Click the drop-down arrow that appears next to the DB file name and select **Generate DDL**.
   The DDL tab shows the SQL code used to add the selected DB file to the dbspace.

See also

- *Viewing DB File Properties* on page 242
- *Adding a DB File* on page 243
- *Modifying a DB File* on page 245
- *Emptying a DB File* on page 246
- *Deleting a DB File* on page 247

Manage and Monitor a Sybase IQ Multiplex

Change a multiplex server configuration, manage the coordinator, secondary and failover nodes, and configure logical servers. Statistics let you monitor the availability and performance of a multiplex.

Viewing Multiplex Overview Statistics

Display the overall health of the Sybase IQ multiplex environment.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Overview**.
3. Select the tab for the required information.

Note: Click a column header to sort the data by that column.

To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description

Area	Description
Multiplex tab > Servers	**Server** – name of the server. **Host** – host name where the server is running. **Port** – port number where the server is running. **State** – current state of the server. Valid states include: • Unknown • Stopped • Running **Role** – role the server plays in the multiplex configuration. Roles include: • Coordinator • Reader • Writer **Status** – current status of the server in the multiplex. Valid states include: • Included • Excluded
Multiplex tab > CPU History	Percentage of total CPU usage over a period of time for each server. The legend below the chart identifies the colored line associated with each server.
Multiplex tab > IQ Memory	Allocation of the IQ memory between the main cache and temporary cache for each server in the multiplex.
Disk Usage tab	Available and used space for the main store and temporary store on each server in the multiplex.
Version Usage tab > Statistics	**Number of committed versions** – the number of table versions in the servers. **Total version space used (MB)** – total space consumed by all the table versions. **Oldest version ID** – the oldest table version identifier on the server. **Number of active versions** – total number of active write table versions on the servers. **Total active version space created (MB)** – amount of data created by active write transactions. **Total active version space to be destroyed (MB)** – amount of data destroyed by active write transactions. If these transactions commit, the destroyed data becomes an old version and is eventually dropped. If the transactions roll back, the created data is released.

Area	Description
Version Usage tab > Multiplex Version Usage	**Version ID** – the table version identifier.
	Server name – the name of the server where the table version exists.
	Connection ID – the connection ID using this table version.
	WasReported – indicates whether the server has received usage information for this table version.
	MinKBRelease – the minimum amount of space returned once this table version is no longer in use.
	MaxKBRelease – the maximum amount of space returned once this table version is no longer in use.

See also

- *Viewing Multiplex Topology Statistics* on page 250
- *Viewing Multiplex Server Statistics* on page 253
- *Viewing Multiplex Connection Statistics* on page 255
- *Viewing Multiplex Transaction Statistics* on page 256
- *Viewing Multiplex Dbspace Statistics* on page 257
- *Viewing Multiplex Cache Statistics* on page 257
- *Manage a Multiplex Server* on page 258
- *Viewing Overview Statistics* on page 204

Viewing Multiplex Topology Statistics

Display the topology view of the Sybase IQ multiplex.

The topology view represents the entire multiplex grid environment, which consists of nodes and links. A node represents a multiplex server, while a link represents the connection between two multiplex nodes. Only one coordinator node appears, and links exist only between the coordinator node and a secondary node. There are no links between two secondary nodes.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Topology**.
3. To modify the layout of the topology view, click **View Controls**.

   Layout options include:
   - **Autofit** – rearranges the nodes to ensure all nodes are visible.
   - **Zoom** – increases or decreases the size of the view.
   - **Reset** – restores the topology view to the default layout.

4. To view a list of the all nodes in the multiplex, and a list of the connections between the coordinator node and its secondary nodes within the multiplex, click **Details**.

5. To monitor a single node within the multiplex, right-click the node in the topology view and select **Monitor Node**.

   When monitoring a single node from the topology screen, the IQ Node Level Monitor is displayed, and the node is registered as a single node resource and appears in the Perspecive Resources view.

See also

Displaying the Properties of a Multiplex Node
View the server information for a single node in the multiplex environment.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.

2. In the left pane of the IQ Multiplex Level Monitor view, select **Topology**.

3. Right-click the node and select **Properties**.

The Server Properties dialog includes:

- **Server name** – name of the server.
- **Host** – host name where the server is running.
- **Port** – port number where the server is running.
- **State** – current state of the server. Valid states include:
  - Unknown
  - Stopped
  - Running
- **INC state** – state of the internode communication between the secondary node and the coordinator. Valid values include:
  - Active
  - Timed out
  - N/A (not available)
  - Unknown

> **Note:** When viewing the properties of the coordinator node, the INC state always displays N/A.

- **Role** – role of the server within the multiplex. Valid roles include:
  - Coordinator
  - Writer
  - Reader
- **Status** – current status of the server. Valid states include:
  - Included
  - Excluded
- **Database Path** – location of the database file on the server.

See also
- *Displaying Connection Properties* on page 252

Displaying Connection Properties
View the details of the connection between the coordinator node and secondary node.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Topology**.
3. Right-click the line between the coordinator node and the secondary node and select **Properties**.

The internode communication properties dialog includes:

- **Link** – name of the linked coordinator node and secondary node.
- **INC state** – state of the internode communication between the secondary node and the coordinator. Valid values include:
  - Active
  - Timed out
  - N/A (not available)
  - Unknown
- **Secondary server name** – name of the secondary server.
- **Heartbeat frequency** – number of seconds between polls to ensure the secondary server is connected.
- **Last successful heartbeat** – date and time the last successful heartbeat transmission was received.
- **Time not responding** – amount of time since the first failed heartbeat request.
- **Time until timeout** – amount of time until the server connection times out and the secondary node becomes inactive.
- **Liveness timeout** – amount of time before the connection is terminated.

- **Auto exclude timeout** – amount of time before the secondary node is automatically excluded.
- **Max connection pool size** – maximum number of connections to the server.
- **Current connection pool size** – current number of connections to the server.
- **Number of idle connections** – number of connections to the server without any activity.
- **Number of connections in use** – number of active connections to the server.

See also
- *Displaying the Properties of a Multiplex Node* on page 251

Viewing Multiplex Server Statistics

Display the statistics for the servers in the Sybase IQ multiplex.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Servers**.
3. In the Multiplex Servers area, select the server to display in the Server Details area.

Note: Click a column header to sort the data by that column.

To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Tab	Description
Multiplex Servers	**Server** – name of the server.
	Host – host name where the server is running.
	Port – port number where the server is running.
	State – current state of the server. Valid states include:
	• Unknown • Stopped • Running
	Role – role of the server within the multiplex. Valid roles include:
	• Coordinator • Writer • Reader
	Status – current status of the server. Valid states include:
	• Included • Excluded
Server Details	**Server name** – name of the server.
	State – current state of the server. Valid states include:
	• Unknown • Stopped • Running
	Database – name of the IQ database.
	Database path – location of the database file on the server.
	Server version – version of the IQ server.
	Platform – operating system running on the host of the server.

See also

Viewing Multiplex Connection Statistics

Display the connection statistics for all servers in a Sybase IQ multiplex.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Connections**.
3. To display a list of connections in the Multiplex Connections area, click the arrow beside the server name.
4. On the Active Connections tab, click the bar chart for a server to display the active connection details in a pie chart.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Connections tab > Multiplex Connections	Displays all users currently connected to each server, including: **Server** – the name of the server. **User ID** – the ID of the connected user. **Connection ID** – the ID of the connection. **Name** – the name of the connected user. **Connection create time** – the date and time the connection was established. **Client IP address** – the IP address of the client that made the connection. **Connection or cursor** – indicates a connection or an active cursor
Connections tab > User Connections/Disconnections Per Minute	Displays the number of user connections and disconnections per minute for each server.
Active Connections tab	Displays a chart of the number of user, internode incoming, and other connections to the selected server.

See also

- *Viewing Multiplex Overview Statistics* on page 248
- *Viewing Multiplex Topology Statistics* on page 250

Viewing Multiplex Transaction Statistics

Display transaction statistics for all servers in a Sybase IQ multiplex.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Transactions**.
3. To display a list of transactions, click the arrow beside the server name.
4. On the Transaction Statistics tab, click the bar chart for a server to display the transaction details in a pie chart.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Area	Description
Transactions tab	**Server** – the name of the server.
	Transaction ID – the unique identification number for the selected transaction.
	Version ID – the version identification number for the selected transaction.
	User ID – the user name of the user that started the selected transaction.
	State – the current state of the selected transaction. Possible states include:
	• **ACTIVE** – the transaction is being processed.
	• **COMMITTED** – the transaction has completed processing.
	Creation time – the date and time when the selected transaction was created.
Transaction Statistics tab	Displays a chart of the number of transactions for each server in the multiplex. Selecting a server in the chart displays a pie chart of user and internode communications (identified on the chart as INC) transactions.

See also

- *Viewing Multiplex Server Statistics* on page 253
- *Viewing Multiplex Connection Statistics* on page 255
- *Viewing Multiplex Dbspace Statistics* on page 257
- *Viewing Multiplex Cache Statistics* on page 257
- *Manage a Multiplex Server* on page 258
- *Viewing Transaction Statistics* on page 210

Viewing Multiplex Dbspace Statistics

View dbspace size and usage details for multiplex dbspaces, including shared temporary dbspaces.

1. In the Perspective Resources view, highlight the resource and select **Resource** > **Administration Console**.
2. In the left pane of the Administration Console view, select **IQ Servers** > **Multiplex Management** > **Multiplex Servers**.
3. In the Multiplex Servers tab, highlight a resource.
4. Hover the mouse pointer over the resource name, click the arrow, and select **Monitor**.
5. In the Launch Monitor view, click **Yes**.
6. In the left pane of the IQ Node Level Monitor view, select **Dbspaces**.
7. In the Dbspaces view, highlight the dbspace name.
   The Dbspace Details view displays DBspace size, and Percent available. The **Size Details** tab show a pie chart comparing Dbspace size in use to size available. The **DB Files** panel shows all the DB files that are part of highlighted DBspace.
8. Highlight the DB file to show its information in the **DB File Details** pane.

See also
- *Viewing Multiplex Overview Statistics* on page 248
- *Viewing Multiplex Topology Statistics* on page 250
- *Viewing Multiplex Server Statistics* on page 253
- *Viewing Multiplex Connection Statistics* on page 255
- *Viewing Multiplex Transaction Statistics* on page 256
- *Viewing Multiplex Cache Statistics* on page 257
- *Manage a Multiplex Server* on page 258

Viewing Multiplex Cache Statistics

Display the cache statistics for all servers in a Sybase IQ multiplex.

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Multiplex**.
2. In the left pane of the IQ Multiplex Level Monitor view, select **Caches**.

3. Select the tab for the cache type to show.

Each cache type tab provides a graph with colored lines for each selected server.

Note: To display the information in a chart or table in the full window, select the Maximize icon in the top right of the area.

Hover the mouse pointer over any line or bar graph to display information for that graph.

Chart	Description
Cache Reads	Number of cache reads per second for a period of time. The cache reads for the catalog, main, and temporary caches appear. The legend at the right of the chart identifies the colored line associated with each server.
Cache Size	The megabyte allocation for the selected cache type on each server in the multiplex, and the number of megabytes in use.

See also
- *Viewing Multiplex Overview Statistics* on page 248
- *Viewing Multiplex Topology Statistics* on page 250
- *Viewing Multiplex Server Statistics* on page 253
- *Viewing Multiplex Connection Statistics* on page 255
- *Viewing Multiplex Transaction Statistics* on page 256
- *Viewing Multiplex Dbspace Statistics* on page 257
- *Manage a Multiplex Server* on page 258
- *Viewing Cache Statistics* on page 213

Manage a Multiplex Server
Change a multiplex server configuration, manage the secondary and failover servers, and configure logical servers.

Note: If you do not have DBA authority, you must use a login that has access to all servers in the multiplex configuration when authenticating a multiplex resource. This access is governed by your user's login policy. See the *Introduction to Sybase IQ* manual for information on login policies. See the *Using Sybase IQ Multiplex* manual for information on logical server configuration.

See also
- *Viewing Multiplex Overview Statistics* on page 248
- *Viewing Multiplex Topology Statistics* on page 250
- *Viewing Multiplex Server Statistics* on page 253
- *Viewing Multiplex Connection Statistics* on page 255
- *Viewing Multiplex Transaction Statistics* on page 256

- *Viewing Multiplex Dbspace Statistics* on page 257
- *Viewing Multiplex Cache Statistics* on page 257
- *Authenticating a Login Account for a Managed Resource* on page 126
- *Registering and Authenticating a Sybase Control Center Agent* on page 123
- *Registering a Sybase IQ Server* on page 120

Changing a Multiplex Server Configuration

Modify general, configuration, and agent information for the selected multiplex server.

Prerequisites

- You have DBA or Multiplex Admin authority.
- The multiplex server is running.
- A Sybase Control Center agent is running and authenticated.

Task

1. In the Perspective Resources view, select a multiplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ > Multiplex Management > Multiplex Servers**.
4. From the right pane of the Administration Console, select a node.
5. Select a multiplex server from the right pane and either:
   - Click the arrow to the right of the name and select **Properties**, or
   - From the Administration Console menu bar, select **Resource > Properties**
6. Click **Configuration**.
7. In the **Server Name** field, enter a new name.
8. Modify the public hosts and ports and private hosts and ports:
   a) Double-click a public or private host name and enter a new host name.
   b) Double-click a public or private port and enter the new port number.
   c) If necessary, click **Add** to add a new port and enter the host name and port number.
   d) If necessary, select a public or private host and port and click **Drop** to drop the host and port.
9. In the **Database** file path field, modify the location of the .db file.
10. Click **OK** or **Apply**.

See also

- *Converting a Simplex Server to Multiplex* on page 260
- *Secondary Servers* on page 262
- *Logical Servers* on page 267

- *Failover* on page 274

Converting a Simplex Server to Multiplex

Add a secondary server to a simplex to convert the simplex into a multiplex.

Prerequisites

- You have DBA authority (or Multiplex Admin and Backup and Space Admin authority).
- The simplex server is running.
- The simplex server's SCC agent is registered, authenticated, and running.

Task

1. In the Perspective Resources view, select a simplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select either:

    - **IQ Servers**
    - **IQ Servers > Multiplex Management > Multiplex Servers**
4. Select the simplex from the right pane and either:

    - Click the arrow to the right of the name and select **Add Secondary Nodes**, or
    - From the Administration Console menu bar, select **Resource > Add Secondary Nodes**
5. Enter a name in the **Multiplex Name** field.
6. Cick **Add** and specify:

Area	Description
Server name	Secondary server name.
Host	The host name and port number for the new secondary server.
Database path	The path to the database file.
SCC agent port	Port number for the Sybase Control Center agent.
SCC agent user	User ID for the Sybase Control Center agent.
SCC agent password	Password for the Sybase Control Center agent.
Public host/port pairs	Host/port pairs in the format **host1:port2,host2:port2**, and so on.
Private host/port pairs	Host/port pairs in the format **host1:port2,host2:port2**, and so on.

Area	Description
Role	Reader or writer.
Status	Included or excluded.
Local temp dbspace path	Temporary store path.
Local temp dbspace size (MB)	Size, in megabytes, of the temporary IQ store. 0 if you select a raw device.
Local temp dbspace reserve (MB)	The amount of space, in megabytes, to reserve for future expansion in the temporary IQ store.
Raw device	Indicates a raw disk.

7. Click **OK**.

8. (Optional) Import server definitions from a CSV file.

a) On the Node Definitions page of the wizard, click **Import** and specify:

Area	Description
File Name	File name of the CSV file containing delimited list of server definitions.
Browse	Opens a window where you locate the CSV file.
Field Delimiter	The default field delimiter is a vertical bar "|". If your CSV file uses a different field delimiter, enter it in this field.

b) Click **OK**.

9. (Optional) Export your existing server definitions to a CSV file for safekeeping:

a) On the Node Definitions page of the wizard, click **Export** and specify:

Area	Description
Field Delimiter	The default field delimiter is a vertical bar "|". You can specify a different field delimiter.

b) Click **OK**.

10. Click **Execute**.

If Sybase Control Center continues to display the converted multiplex server as a simplex, reauthenticate the simplex resource to force SCC to update its display.

See also

- *Logical Servers* on page 267
- *Failover* on page 274

Secondary Servers

One or more secondary servers may participate in a Sybase IQ multiplex configuration.

One secondary server acts as a designated failover server, which is assigned the coordinator role if the current coordinator cannot continue.

Secondary servers can be either read-only servers (reader servers) or read-write servers (writer servers).

See also
- *Changing a Multiplex Server Configuration* on page 259
- *Converting a Simplex Server to Multiplex* on page 260
- *Logical Servers* on page 267
- *Failover* on page 274

Adding a Secondary Server
Add a secondary server to the multiplex server.

Prerequisites

- You have DBA authority (or Multiplex Admin and Backup and Space Admin authority).
- The multiplex server is running.

Task

1. In the Perspective Resources view, select a multiplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ > Multiplex Management > Multiplex Servers**.
4. Select a multiplex server from the right pane and either:

    - Click the arrow to the right of the name and select **Add Secondary Nodes**, or
    - From the Administration Console menu bar, select **Resource > Add Secondary Nodes**
5. On the Node Definitions page of the wizard, click **Add** and specify:

Area	Description
Server name	Secondary server name.

Area	Description
Host	The host name and port number for the new secondary server.
Database path	The path to the database file.
SCC agent port	Port number for the Sybase Control Center agent.
SCC agent user	User ID for the Sybase Control Center agent.
SCC agent password	Password for the Sybase Control Center agent.
Public host/port pairs	Host/port pairs in the format **host1:port2,host2:port2**, and so on.
Private host/port pairs	Host/port pairs in the format **host1:port2,host2:port2**, and so on.
Role	Reader or writer.
Status	Included or excluded.
Local temp dbspace path	Temporary store path.
Local temp dbspace size (MB)	Size, in megabytes, of the temporary IQ store. 0 if you select a raw device.
Local temp dbspace reserve (MB)	The amount of space, in megabytes, to reserve for future expansion in the temporary IQ store.
Raw device	Indicates a raw disk.

6. Click **OK**.
7. (Optional) Import server definitions from a CSV file:
   a) On the Node Definitions page of the wizard, click **Import** and specify:

Area	Description
File Name	File name of the CSV file containing delimited list of server definitions.
Browse	Opens a window where you locate the CSV file.
Field Delimiter	The default field delimiter is a vertical bar "|". If your CSV file uses a different field delimiter, enter it in this field.

   b) Click **OK**.
8. (Optional) Export your existing server definitions to a CSV file for safekeeping.
   a) On the Node Definitions page of the wizard, click **Export** and specify:

Area	Description
Field Delimiter	The default field delimiter is a vertical bar "|". You can specify a different field delimiter.

    b) Click **OK**.

9. On the Execute page of the wizard, click **Execute**.

See also

- *Dropping a Secondary Server* on page 264
- *Synchronizing a Secondary Server* on page 265
- *Including an Excluded Secondary Server* on page 265
- *Excluding a Secondary Server* on page 266

Dropping a Secondary Server

Remove a secondary server from a multiplex server. You can drop the designated failover server only if it is the only secondary server in the multiplex.

Prerequisites

- You have DBA authority.
- The SCC agent is running.

Task

1. In the Perspective Resources view, select a multiplex resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. From the left pane of the Administration Console, select **IQ > Multiplex Management > Multiplex Servers**.

4. Select a multiplex server from the right pane and either:

- Click the arrow to the right of the name and select **Drop Secondary Node**, or
- From the Administration Console menu bar, select **Resource > Drop Secondary Node**

5. (Optional) Select **Delete the database directory**. This cleans up the database directory by removing any files that belong to the server.

Warning! Use caution in selecting the option to delete files. If any of the files to be deleted are shared main files, data can be lost, and servers might not start correctly afterwards.

6. Click **Finish**.

See also

- *Adding a Secondary Server* on page 262

- *Synchronizing a Secondary Server* on page 265
- *Including an Excluded Secondary Server* on page 265
- *Excluding a Secondary Server* on page 266

Synchronizing a Secondary Server
Synchronization updates a secondary server with respect to the coordinator.

Prerequisites
You have DBA authority.

Task

1. In the Perspective Resources view, select a multiplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ > Multiplex Management > Multiplex Servers**.
4. From the right pane of the Administration Console, select one or more secondary servers and either:

   - Click the arrow to the right of the name and select **Synchronize Server**, or
   - From the Administration Console menu bar select **Resource > Synchronize Server**

   If you selected one secondary server, the secondary server is synchronized.
5. On the Synchronize Server window, click **Finish**.

See also
- *Adding a Secondary Server* on page 262
- *Dropping a Secondary Server* on page 264
- *Including an Excluded Secondary Server* on page 265
- *Excluding a Secondary Server* on page 266

Including an Excluded Secondary Server
Include a previously excluded secondary server when the excluded server's shutdown period is over.

Prerequisites

- You have DBA authority (or Multiplex Admin and Backup authority).
- The server you are including must have Excluded status.

Task

1. In the Perspective Resources view, select a multiplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ > Multiplex Management > Multiplex Servers**.
4. From the right pane of the Administration Console, select an excluded secondary server and either:

   - Click the arrow to the right of the name and select **Include Server**, or
   - From the Administration Console menu bar select **Resource > Include Server**
5. Click **Finish**.

See also
- *Adding a Secondary Server* on page 262
- *Dropping a Secondary Server* on page 264
- *Synchronizing a Secondary Server* on page 265
- *Excluding a Secondary Server* on page 266

Excluding a Secondary Server
Exclude a secondary server from the multiplex if it will be shut down for an extended period of time.

Prerequisites
You have DBA or Multiplex Admin authority.

Task

Excluding a secondary server allows the coordinator to ignore it during version cleanup. If you do not exclude a server, the coordinator preserves all old versions of IQ objects changed since the secondary server was shut down, which unnecessarily uses disk space.

Note: You cannot exclude the designated failover server, or the coordinator. Excluding a running server shuts it down.

1. In the Perspective Resources view, select a multiplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ > Multiplex Management > Multiplex Servers**.
4. From the right pane of the Administration Console, select a secondary node and either:

   - Click the arrow to the right of the name and select **Exclude Server**, or
   - From the Administration Console menu bar, select **Resource > Exclude Server**

5. Click **Finish**.

See also
- *Adding a Secondary Server* on page 262
- *Dropping a Secondary Server* on page 264
- *Synchronizing a Secondary Server* on page 265
- *Including an Excluded Secondary Server* on page 265

Logical Servers

A logical server allows you to group a subset of physical hardware resources together as a logical entity that appears as a single multiplex server, when it is actually one or more servers within the physical multiplex.

DBA or Multiplex Admin authority is required to work with logical servers.

In the Sybase IQ documentation, see *Using Sybase IQ Multiplex > Logical Servers* for logical server concepts and logical server policy concepts.

See also
- *Changing a Multiplex Server Configuration* on page 259
- *Converting a Simplex Server to Multiplex* on page 260
- *Secondary Servers* on page 262
- *Failover* on page 274

Changing Logical Server Properties

Edit the name of the logical server and the multiplex nodes associated with it.

Prerequisites

You have DBA or Multiplex Admin authority.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Multiplex Management > Logical Servers**.
4. Select the logical server from the right pane and either:
   - Click the arrow to the right of the name and select **Properties**, or
   - From the Administration Console menu bar, select **Resource > Properties**

Area	Description

Area	Description
General tab	**Name** – name of the logical server. **Mode** – membership status of the node (Shared). **Comment** – a text field for adding an optional comment about the logical server.
Multiplex Nodes tab	Shows the available multiplex nodes. Nodes included in the current logical server are selected.

See also

Creating a Logical Server
Create a logical server to group multiple physical multiplex servers into a single logical entity.

Prerequisites
You have DBA or Multiplex Admin authority.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **IQ Servers > Multiplex Management > Logical Servers**.
4. Click the arrow next to **Logical Servers** and select **New**.
5. On the Logical Server Name page of the wizard, specify:

Area	Description
Select the resource for which the logical server will be created.	List of multiplex servers.
What do you want to name the new logical server?	Unique name for the logical server.

Area	Description
What would you like the comment to be for this logical server?	Optional comment about this logical server.

6. Click **Next**.

Note: Using the **Next** button is optional. You can also navigate through the wizard by selecting wizard pages from the left pane.

---

7. On the Logical Server Node Membership page, select the multiplex nodes you want to add to the logical server. You can select the coordinator node, or secondary nodes.

Note: Select the FOR LOGICAL COORDINATOR node to add the current coordinator as a member. This ensures the coordinator is always available to the logical server as its member, regardless of which multiplex node plays the coordinator role. In the Sybase IQ documentation, see *Using Sybase IQ Multiplex > Logical Servers* for information on failover.

---

8. Click **Finish**.

See also
- *Changing Logical Server Properties* on page 267
- *Deleting a Logical Server* on page 269
- *Configuring Logical Server Node Membership* on page 270
- *Changing the Logical Server Policy* on page 271
- *Altering a Logical Server Assignment* on page 272
- *Generating DDL Commands for a Logical Server* on page 274

Deleting a Logical Server
Delete a logical server from the database. Deleting a logical server does not delete its component multiplex nodes from the database.

Prerequisites
You have DBA or Multiplex Admin authority.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Multiplex Management > Logical Servers**.
4. Select the logical server from the right pane and either:

- Click the arrow to the right of the name and select **Delete**, or
- From the Administration Console menu bar, select **Resource > Delete**

5. Click **Yes** when asked to confirm deletion.

See also
- *Changing Logical Server Properties* on page 267
- *Creating a Logical Server* on page 268
- *Configuring Logical Server Node Membership* on page 270
- *Changing the Logical Server Policy* on page 271
- *Altering a Logical Server Assignment* on page 272
- *Generating DDL Commands for a Logical Server* on page 274

Configuring Logical Server Node Membership
Add or remove multiplex nodes to or from the logical server.

Prerequisites

- You have DBA or Multiplex Admin authority.
- You are familiar with the restrictions and rules regarding which coordinator, reader, and writer nodes you can add to a logical server. In the Sybase IQ documentation, see *Using Sybase IQ Multiplex > Logical Servers*.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Multiplex Management > Logical Servers**.
4. Select the logical server from the right pane and either:

   - Click the arrow to the right of the name and select **Properties**, or
   - From the Administration Console menu bar, select **Resource > Properties**

5. Click the **Multiplex Nodes** tab.
6. Select or unselect multiplex nodes to modify the logical server configuration.

   Note: Select the FOR LOGICAL COORDINATOR node to add the current coordinator as a member. This ensures the coordinator is always available to the logical server as its member, regardless of which multiplex node plays the coordinator role. In the Sybase IQ documentation, see *Using Sybase IQ Multiplex > Logical Servers* for information on failover.

7. Click **Apply**.

8. Click **OK**.

See also
- *Changing Logical Server Properties* on page 267
- *Creating a Logical Server* on page 268
- *Deleting a Logical Server* on page 269
- *Changing the Logical Server Policy* on page 271
- *Altering a Logical Server Assignment* on page 272
- *Generating DDL Commands for a Logical Server* on page 274

Changing the Logical Server Policy
The only change you can make to the built-in root logical server policy to set the
ALLOW_COORDINATOR_AS_MEMBER option. You cannot drop the root logical server
policy or create new logical server policies.

Prerequisites
You have DBA or Multiplex Admin authority.

Task

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select **Multiplex Management > Logical Servers > Logical Server Policies**.

4. Select the logical server from the right pane and either:
- Click the arrow to the right of the name and select **Properties**, or
- From the Administration Console menu bar, select **Resource > Properties**

5. Change the value of **Allow Coordinator as Member**.

6. Click **Apply**.

7. Click **OK**.

Area	Description
Allow Coordinator As Member	Allows the coordinator node to be also available to one or more logical servers, as a member node, to process their normal workload. You can set this option to On (the default) or Off.

See also
- *Changing Logical Server Properties* on page 267
- *Creating a Logical Server* on page 268

- *Deleting a Logical Server* on page 269
- *Configuring Logical Server Node Membership* on page 270
- *Altering a Logical Server Assignment* on page 272
- *Generating DDL Commands for a Logical Server* on page 274

Altering a Logical Server Assignment

Add one or more user-defined logical servers to a login policy. A check ensures that no membership overlap exists among the logical servers assigned to the login policy.

Prerequisites

MPX ADMIN authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Multiplex Management > Logical Servers**.
4. Select the logical server from the right pane and either:

    - Click the arrow to the right of the name and select **Properties**, or
    - From the Administration Console menu bar, select **Resource > Properties**
5. On the General tab:

Area	Description
Options	**Password life time** – number of days the password is valid. The user must reset their password when the lifetime expires.
	Password grace time – number of days before password expiry that users receive warnings that their passwords are about to expire.
	Password expiry on next login – whether the user must reset the password at the next login.
	Locked – whether or not the user is locked out. When this option is on login attempts result in an error message.
	Maximum connections – number of times the same user can be logged in to the server.
	Maximum failed login attempts – number of failed login attempts before the account is locked.
	Maximum days since login – number of days allowed between logins before the account is locked.

Area	Description
Clear All Overridden Values	Clears all overrides of the default login policy option values.
Restore to IQ Default	Changes all option settings back to default.

6. (Optional) Enter a comment for this login policy.

7. Click **OK**.

8. On the Logical Server Assignment tab:

Area	Description
Which logical servers do you want to add to this login policy?	**CUSTOM** – Allows access to OPEN and user-defined logical servers. **DEFAULT** – Inherits the logical server assignment of the root login policy. **NONE** – Disallows access to all logical servers. **SERVER** – Allows access to any multiplex node. MPX authority is required to connect.
Modify logical server assignment for a login policy	Select the logical servers to add to the login policy.

9. Click **OK**.

10. On the Logical Server Level Option Overrides tab:

Area	Description
Select a logical server and specify option overrides	Specifies the value of the **Max Conn.** (maximum connection) parameter, which overrides the inherited value.

11. Click **OK**.

See also
- *Changing Logical Server Properties* on page 267
- *Creating a Logical Server* on page 268
- *Deleting a Logical Server* on page 269
- *Configuring Logical Server Node Membership* on page 270
- *Changing the Logical Server Policy* on page 271
- *Generating DDL Commands for a Logical Server* on page 274

Generating DDL Commands for a Logical Server
Display the SQL data description language for creating a new logical server. The SQL code can be a useful reference and training tool.

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Multiplex Management > Logical Servers**.
4. Select the logical server from the right pane and either:

   • Click the arrow to the right of the name and select **Generate DDL**, or
   • From the Administration Console menu bar, select **Resource > Generate DDL**

   The DDL tab opens, showing the SQL code used to create the selected logical server.

See also
• *Changing Logical Server Properties* on page 267
• *Creating a Logical Server* on page 268
• *Deleting a Logical Server* on page 269
• *Configuring Logical Server Node Membership* on page 270
• *Changing the Logical Server Policy* on page 271
• *Altering a Logical Server Assignment* on page 272

Failover
Manual failover promotes a server to act as the coordinator in the event of coordinator node failure or maintenance shutdown.

See also
• *Changing a Multiplex Server Configuration* on page 259
• *Converting a Simplex Server to Multiplex* on page 260
• *Secondary Servers* on page 262
• *Logical Servers* on page 267

Designating the Failover Node
Designate the secondary server as the failover node so it can take over as coordinator if the current coordinator stops running.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA authority.

Task

The designated failover node defaults to the first multiplex server added to the multiplex. If the default is not acceptable, designate a different failover node.

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ Servers > Multiplex Management > Multiplex Servers**.
4. Select a node from the right pane and either:
   - Click the arrow to the right of the name and select **Properties**, or
   - From the Administration Console menu bar, select **Resource > Properties**
5. From the left pane of the Properties window, select **General**.
6. Click **Change** next to the **Designated Failover Node** field.
7. Select a node from the list and click **Finish**.
8. Click **OK**.

See also
- *Performing Coordinator Node Failover* on page 275

Performing Coordinator Node Failover
Manually switch the coordinator role to the designated failover node.

Prerequisites
- Authenticate with Sybase IQ using an account that has DBA authority.
- Make sure that the former coordinator process is no longer running.

Warning! Performing failover while the former coordinator process is alive may cause database corruption.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. From the left pane of the Administration Console, select **IQ Servers > Multiplex Management > Multiplex Servers**.
4. Select a node from the right pane and either:
   - Click the arrow to the right of the name and select **Failover**, or
   - From the Administration Console menu bar, select **Resource > Failover**

5. In the Failover wizard, click **Finish**.

See also

* *Designating the Failover Node* on page 274

Manage Sybase IQ Security

Add, change, and delete users, or groups containing users or other groups, grant or revoke database authorities, and define login policies.

Manage Users and Groups

Add, change, and delete users, or groups containing users and/or other groups, as a prerequisite to managing user authorities.

Both users and groups are objects within the database. Groups are also containers for containing users and other groups. Think of a group as a user ID with special permissions, such as the ability to have members. You grant and revoke authorities for a group in exactly the same manner as you do for users.

For most databases, manage authorities by using groups, rather than by assigning authorities to individual users one at a time.

See also

* *Configuring Sybase IQ for Monitoring* on page 116

Viewing Properties for Sybase IQ Users or Groups

View the details of login parameters and authorities for users or groups.

Prerequisites

User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Users** or **Security > Groups**.
4. Select the user or group from the right pane and either:

   * Click the arrow to the right of the name and select **Properties**, or
   * From the Administration Console menu bar, select **Resource > Properties**

Area	Description
General tab	**Name** – name of the user or group.
	Enable Password – allows the user or group to connect to the database with password security. Clearing this option disables the Password and Confirm Password options.
	Password – the password for the user or group. Characters appear as asterisks.
	Confirm password – a field for confirming the password that you typed in the Password text box. The contents of the two fields must match exactly.
	Password creation time – date and time when the password was created.
	Change password on next login – force the user or group to change the password at the next login.
	Login policy – select the login policy that applies for this user or group.
	Last login time – last time the user successfully logged in.
	Failed login attempts – number of times the user has tried to log in with an incorrect password.
	Locked – displays **false** if the account is unlocked. Displays **true** if the user or group has exceeded the allowed number of failed login attempts.
	Unlock now – unlocks the account if **Locked** is true.
	Comment – a text field for adding an optional comment about the user or group.

Area	Description
Authori- ties tab	**Backup** - allows this user or group to back up the database. **DBA** - grants DBA authority to administer the database. **Multiplex Admin** - grants Multiplex Admin authority to the user or group; a user or group with Multiplex Admin authority can perform multiplex administration tasks. **Operator** - grants Operator authority to the user or group; a user or group with Operator authority can checkpoint databases, drop connections, backup databases, and monitor the system. **Perms Admin** - grants Perms Admin authority to the user or group; a user or group with Perms Admin authority can manage data permissions, groups, authorities, and passwords. **Profile** - allows this user or group to perform application and procedure profiling, and request log creation, and analysis. Profile authority is also required by the index consultant. **Read client file** - allows this user to read from a file on the client computer. **Read file** - grants Read file authority to the user or group, allowing the user or group to execute SELECT statements against a file using the OPENSTRING clause. **Remote DBA** - grants Remote DBA authority to any table the user or group can have access to. To ensure that actions are secure, run the SQL Remote Message Agent using a user ID with this type of authority. **Resource** - grants resource authority to the user or group; a user or group with resource authority can create database objects. **Space Admin** - grants Space Admin authority to the user or group; a user or group with Space Admin authority can manage dbspaces. **User Admin** - grants User Admin authority to the user or group; a user or group with User Admin authority can manage users, external logins, and login policies. **Validate** - allows this user or group to validate tables, materialized views, and indexes. **Note:** Materialized views are only supported for SQL Anywhere tables in the IQ catalog store. **Write client file** - allows this user to write to a file on the client computer.

See also
- *Adding a Sybase IQ User or Group* on page 279
- *Deleting a Sybase IQ User* on page 281
- *Deleting a Sybase IQ Group* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282

- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ User to a Group* on page 284
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Adding a Sybase IQ User or Group

Add a new user to the database using the Create User wizard. Add a new group to the database using the Create Group wizard. The two wizards are identical.

Prerequisites

User Admin authority is required.

Task

All new users are automatically added to the PUBLIC group.

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **IQ Servers > Security > Users** or **IQ Servers > Security > Groups**.
4. Select **Users** or **Groups**, hover the mouse pointer over the selection, click the arrow, and select **New**.
5. On the Welcome page of the wizard, select a server.
6. Enter a name for the new user or group, then click **Next**.

   Note: Using the **Next** button is optional. You can also navigate through the wizard by selecting wizard pages from the left pane.

7. On the Specify a Password page, specify:

Options	Description
Enable password	Allows the user or group to connect to the database with password security. Leave this option unselected to disable the password and confirm password options.
Password	The password for the user or group. Characters appear as asterisks.
Confirm password	Confirms the password. The contents of the two password fields must match exactly.
Change password on next login	Forces the user or group to change the password at the next login.
Login policy	List of available login policies.

8. Click **Next**.

The Select authorities page of the wizard appears.

9. Select the authorities you want to assign to the new user or group:

Options	Description
Backup	Allows this user or group to back up the database.
DBA	Grants DBA authority to administer the database.
Multiplex Admin	Grants Multiplex Admin authority to the user or group; Multiplex Admin authority allows a user or group to perform multiplex administration tasks.
Operator	Grants Operator authority to the user or group; Operator authority lets a user or group checkpoint databases, drop connections, backup databases, and monitor the system.
Perms Admin	Grants Perms Admin authority to the user or group; Perms Admin authority lets a user or group manage data permissions, groups, authorities, and passwords.
Profile	Allows this user or group to perform application and procedure profiling, and request log creation, and analysis. Profile authority is also required by the index consultant.
Read client file	Allows this user or group to read from a file on the client computer.
Read file	Grants Read file authority to the user or group, allowing the user or group to execute **SELECT** statements against a file using the **OPENSTRING** clause.
Remote DBA	Grants Remote DBA authority to any table the user or group can access. To ensure that actions are secure, run the SQL Remote Message agent using a user ID with this type of authority.
Resource	Grants resource authority to the user or group; resource authority lets a user or group create database objects.
Space Admin	Grants Space Admin authority to the user or group; Space Admin authority lets a user or group manage dbspaces.
User Admin	Grants User Admin authority to the user or group; User Admin authority lets a user or group manage users, external logins, and login policies.
Validate	Allows this user or group to validate tables, materialized views, and indexes.
Write client file	Allows this user or group to write to a file on the client computer.

10. Click **Next**.

11. (Optional) Enter a comment for this user or group.

12. Click **Finish**.

See also

- *Deleting a Sybase IQ Group* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282
- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ User to a Group* on page 284
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Deleting a Sybase IQ User

Delete a user from the database.

Prerequisites

User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Users**.
4. Select the user from the right pane and either:

   - Click the arrow to the right of the name and select **Delete**, or
   - From the Administration Console menu bar, select **Resource > Delete**

5. Click **Yes** when asked to confirm deletion.

See also
- *Viewing Properties for Sybase IQ Users or Groups* on page 276
- *Adding a Sybase IQ User or Group* on page 279
- *Deleting a Sybase IQ Group* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282
- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ User to a Group* on page 284
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Deleting a Sybase IQ Group

Delete a group from the database.

Prerequisites

User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Groups**.
4. Select the group from the right pane and either:

    - Click the arrow to the right of the name and select **Delete**, or
    - From the Administration Console menu bar, select **Resource > Delete**

5. Click **Yes** when asked to confirm deletion.

See also
- *Viewing Properties for Sybase IQ Users or Groups* on page 276
- *Adding a Sybase IQ User or Group* on page 279
- *Deleting a Sybase IQ User* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282
- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ User to a Group* on page 284
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Adding a Member to a Sybase IQ Group

Add a group member to the current group, or add the current group to a higher-level parent group.

Prerequisites

User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Groups**.
4. Select the group from the right pane and either:

    - Click the arrow to the right of the name, or
    - From the Administration Console menu bar, select **Resource**

5. Make a selection from the menu:

Menu Selection	Description
Manage members	Opens a window where you add users or groups to this group.
Manage parent groups	Opens a window where you add this group to a parent group.

6. Select a user or group from the left pane and click **Add.**
   The user or group is added to the group.
7. Click **OK**.

See also
- *Viewing Properties for Sybase IQ Users or Groups* on page 276
- *Adding a Sybase IQ User or Group* on page 279
- *Deleting a Sybase IQ User* on page 281
- *Deleting a Sybase IQ Group* on page 281
- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ User to a Group* on page 284
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Removing a Sybase IQ User from a Group
Remove a user or group from the current group.

Prerequisites
User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Groups**.
4. Select the group from the right pane and either:

   - Click the arrow to the right of the name and select **Manage Members**, or
   - From the Administration Console menu bar, select **Resource > Manage Members**
5. In the right pane, select the user or group you want to remove from the group and click **Remove**.
6. Click **OK**.

See also
- *Viewing Properties for Sybase IQ Users or Groups* on page 276
- *Adding a Sybase IQ User or Group* on page 279

- *Deleting a Sybase IQ User* on page 281
- *Deleting a Sybase IQ Group* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282
- *Changing a Sybase IQ User to a Group* on page 284
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Changing a Sybase IQ User to a Group

Change an existing user into a group that retains the user's authorities.

Prerequisites

User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Users**.
4. Select the user from the right pane and either:

   - Click the arrow to the right of the name and select **Change to group**, or
   - From the Administration Console menu bar, select **Resource > Change to group**

5. Click **Yes** to confirm the change.

See also
- *Viewing Properties for Sybase IQ Users or Groups* on page 276
- *Adding a Sybase IQ User or Group* on page 279
- *Deleting a Sybase IQ User* on page 281
- *Deleting a Sybase IQ Group* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282
- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ Group to a User* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Changing a Sybase IQ Group to a User

Change an existing group into a user.

Prerequisites

User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select **Security > Groups**.

4. Select the group from the right pane and either:

   - Click the arrow to the right of the name and select **Change to user**, or
   - From the Administration Console menu bar, select **Resource > Change to user**

5. Click **Yes** to confirm the change.

See also
- *Viewing Properties for Sybase IQ Users or Groups* on page 276
- *Adding a Sybase IQ User or Group* on page 279
- *Deleting a Sybase IQ User* on page 281
- *Deleting a Sybase IQ Group* on page 281
- *Adding a Member to a Sybase IQ Group* on page 282
- *Removing a Sybase IQ User from a Group* on page 283
- *Changing a Sybase IQ User to a Group* on page 284
- *Generating DDL Commands for Sybase IQ Users or Groups* on page 285

Generating DDL Commands for Sybase IQ Users or Groups
Display the SQL data description language for creating a user or group. The SQL code can be a useful reference and training tool.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA or User Admin authority.

Task

1. In the Perspective Resources view, select one or more resources and select **Resource > Administration Console**.

2. In the left pane of the Administration Console, select **Security > Users** or **Security > Groups**.

3. Select the user or group from the right pane and either:

   - Click the drop-down arrow to the right of the name and select **Generate DDL**
   - From the Administration Console menu bar, select **Resource > Generate DDL**

   The DDL tab opens, showing the SQL code used to create the selected user or group.

See also
*   *Viewing Properties for Sybase IQ Users or Groups* on page 276
*   *Adding a Sybase IQ User or Group* on page 279
*   *Deleting a Sybase IQ User* on page 281
*   *Deleting a Sybase IQ Group* on page 281
*   *Adding a Member to a Sybase IQ Group* on page 282
*   *Removing a Sybase IQ User from a Group* on page 283
*   *Changing a Sybase IQ User to a Group* on page 284
*   *Changing a Sybase IQ Group to a User* on page 284

Manage Authorities

Grant database authorities to a user or group, and revoke database authorities from a user or group.

For detailed information on using database authorities in Sybase IQ, see *System Administration Guide: Volume 1 > Managing User IDs and Permissions > Database Permissions and Authorities Overview*.

Members of a group can inherit the following authorities set for the group they belong to:

*   READCLIENTFILE
*   READFILE
*   WRITECLIENTFILE
*   DBA
*   RESOURCE
*   OPERATOR
*   MULTIPLEX ADMIN
*   PERMS ADMIN
*   SPACE ADMIN
*   USER ADMIN

Granting an Authority to a User or Group

Grant authorities to allow a user or group to perform specified database operations, such as the ability to manage dbspaces.

Prerequisites
User Admin authority is required.

Task

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select **Security > Users** or **Security > Groups**.
4. Select the user or group from the right pane and either:
   - Click the arrow to the right of the name and select **Properties**, or
   - From the Administration Console menu bar, select **Resource > Properties**
5. Click the **Authorities** tab.
6. Select the authorities you want to assign to the user or group:

Options	Description
Backup	Allows this user or group to back up the database.
DBA	Grants DBA authority to administer the database.
Multiplex Admin	Grants Multiplex Admin authority to the user or group; a user or group with Multiplex Admin authority can perform multiplex administration tasks.
Operator	Grants Operator authority to the user or group; a user or group with Operator authority can checkpoint databases, drop connections, backup databases, and monitor the system.
Perms Admin	Grants Perms Admin authority to the user or group; a user or group with Perms Admin authority can manage data permissions, groups, authorities, and passwords.
Profile	Allows this user or group to perform application and procedure profiling, and request log creation, and analysis. Profile authority is also required by the index consultant.
Read client file	Allows this user or group to read from a file on the client computer.
Read file	Grants Read file authority to the user or group, allowing the user or group to execute SELECT statements against a file using the OPENSTRING clause.
Remote DBA	Grants Remote DBA authority to any table the user or group can have access to. To ensure that actions are secure, run the SQL Remote Message Agent using a user ID with this type of authority.
Resource	Grants resource authority to the user or group; a user or group with resource authority can create database objects.
Space Admin	Grants Space Admin authority to the user or group; a user or group with Space Admin authority can manage dbspaces.
User Admin	Grants User Admin authority to the user or group; a user or group with User Admin authority can manage users, external logins, and login policies.
Validate	Allows this user or group to validate tables, materialized views, and indexes.
Write client file	Allows this user or group to write to a file on the client computer.

7. Click **Apply**.

8. Click **OK**.

See also
- *Revoking an Authority from a User or Group* on page 288
- *Configuring Sybase IQ for Administration* on page 117

Revoking an Authority from a User or Group

Revoke an authority to remove a user's or group's ability to perform specified database operations.

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select **Security > Users** or **Security > Groups**.

4. Select the user or group from the right pane and either:
- Click the arrow to the right of the name and select **Properties**, or
- From the Administration Console menu bar select **Resource > Properties**

5. Click the **Authorities** tab.

6. Clear the check box of the authority you want to revoke.

7. Click **Apply**.

8. Click **OK**.

See also
- *Granting an Authority to a User or Group* on page 286

Manage Login Policies

A login policy is a named object in a database that consists of a set of rules that are applied when you create a database connection for a user. All new databases include a root login policy.

You can modify the root login policy values, but you cannot delete the policy. Login policies govern only the rules for user login and are separate from authorities and permissions. Login policies are not inherited through group memberships.

Adding a Login Policy to a Simplex

Create a login policy to define password and login parameters for users.

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select **IQ Servers > Security > Login Policies**.

4. Hover the mouse pointer over **Login Policies**, click the arrow, and select **New**.
5. On the Login Policy Name page, specify:

Area	Description
Select the server for which the login policy will be created	Select a simplex server.
What do you want to name the new login policy?	Enter a unique name for the new login policy.

6. Click **Next**.

   Note: Using the **Next** button is optional. You can also navigate through the wizard by selecting wizard pages from the left pane.

7. On the Set the Options for this Login Policy page, specify:

Area	Description
Options	**Password life time** – Number of days the password is valid. The user must reset the password when the lifetime expires.
	Password grace time – Number of days before password expiry that users receive warnings that the password is about to expire.
	Password expiry on next login – Whether the user must reset the password at the next login.
	Locked – Whether the user account is locked when maximum number of failed login attempts is exceeded.
	Maximum connections – Number of times the same user can be logged in to the server.
	Maximum failed login attempts – Number of failed login attempts before the account is locked.
	Maximum days since login – Number of days allowed between logins before the account is locked.
Clear All Overridden Values	Clears all overrides.
Restore to IQ Default	Changes all option settings back to default.

8. Click **Next**.
9. (Optional) Enter a comment for this login policy.
10. Click **Finish**.

See also

* *Adding a Login Policy to a Multiplex* on page 290
* *Viewing Login Policy Properties* on page 292

Adding a Login Policy to a Multiplex

You can assign one or more logical servers to a login policy. Multiplex servers can only be accessed only by logical servers.

All users using the login policy can access only those multiplex servers that are effective members of the assigned logical servers.

1. In the Perspective Resources view, select a multiplex resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **Security > Login Policies**.
4. Click **Login Policies** and select **New** from the pull-down menu.
5. On the Login Policy Name page:

Area	Description
Select the resource for which the login policy will be created	Select a multiplex server from the list.
What do you want to name the new login policy?	Enter a unique name for the new login policy.

6. Click **Next**.

   Note: Using the **Next** button is optional. You can also navigate through the wizard by selecting wizard pages from the left pane.

7. On the Set the Options for this Login Policy page:

Area	Description
Options	**Password life time** – number of days the password is valid. The user must reset their password when the lifetime expires.
	Password grace time – number of days before password expiry that users receive warnings that their passwords are about to expire.
	Password expiry on next login – whether the user must reset the password at the next login.
	Locked – whether or not the user is locked out. When this option is on login attempts result in an error message.
	Maximum connections – number of times the same user can be logged in to the server.
	Maximum failed login attempts – number of failed login attempts before the account is locked.
	Maximum days since login – number of days allowed between logins before the account is locked.
Clear All Overridden Values	Clears all overrides of the default login policy option values.
Restore to IQ Default	Changes all option settings back to default.

8. Click **Next**.
9. On the Logical Server Assignment page:

Area	Description
Which logical servers do you want to add to this login policy?	**CUSTOM** – Allows access to OPEN and user-defined logical servers.
	DEFAULT – Inherits the logical server assignment of the root login policy.
	NONE – Disallows access to all logical servers.
	SERVER – Allows access to any multiplex node. MPX authority is required to connect.
Modify logical server assignment for a login policy	Select the logical servers to add to the login policy.

10. Click **Next**.
11. On the Logical Server Level Option Overrides page:

Area	Description
Select a logical server and specify option overrides	Specifies the value of the **Max Conn.** (maximum connection) parameter, which overrides the inherited value.

12. (Optional) Enter a comment for this login policy.

13. Click **Finish**.

See also

- *Adding a Login Policy to a Simplex* on page 288
- *Viewing Login Policy Properties* on page 292

Viewing Login Policy Properties

View the properties of a login policy for a simplex or multiplex.

1. In the Perspective Resources view, select a resource.

2. From the application menu bar, select **View > Open > Administration Console**.

3. In the left pane of the Administration Console, select **Security > Login Policies**.

4. Select the login policy from the right pane and either:

  - Click the arrow to the right of the name and select **Properties**, or
  - From the Administration Console menu bar, select **Resource > Properties**

Area	Description
General tab	Options: **Password life time** – Number of days the password is valid. The user must reset their password when the lifetime expires. **Password grace time** – Number of days before password expiry that users receive warnings that their password is about to expire. **Password expiry on next login** – Whether the user must reset the password at the next login. **Locked** – Whether or not the user is locked out. When this option is on, login attempts result in an error message. **Maximum connections** – Number of times the same user can be logged in to the server. **Maximum failed login attempts** – Number of failed login attempts before the account is locked. **Maximum days since login** – Number of days allowed between logins before the account is locked.
Logical Server Assignment tab	Assignment Type: **CUSTOM** – Allows access to OPEN and user-defined logical servers. **DEFAULT** – Inherits the logical server assignment of the root login policy. **NONE** – Disallows access to all logical servers. **SERVER** – Allows access to any multiplex node. MPX authority is required to connect. **Modify logical server assignment for a login policy** – Select the logical servers to add to the login policy.
Logical Server Option Overrides tab	**Modify the Logical Server level Login Policy Option Override** – Specifies the value of the **Max Conn.** (maximum connection) parameter, which overrides the inherited value.

See also

Manage Functions

Create, delete, display, and manage the properties of functions in Sybase IQ.

Sybase provides sample functions and procedures with Sybase IQ. Sample functions are visible in Sybase Control Center in the Administration Console (IQ Servers > Other Objects > Functions and IQ Servers > Other Objects > Procedures). The samples include explanatory comments and can be found in `IQ/IQ-X_X/samples/udf`, where `X_X` is the Sybase IQ release number (15.4, for example).

Listing Functions

Display a list of functions available on Sybase IQ servers in this perspective.

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Other Objects** and select **Functions**.
   A list of functions in this perspective appears in the right pane.
3. Click the name of any column to sort the display based on the values in that column.
4. (Optional) Enter a filter string in the field at the top of any column to display only functions that match the string. You can filter on two or more columns at once. For example, to find functions belonging to Fred and created in 2011, enter "fred" above the Owner column and "2011" above the Time Created column. (Filter strings are not case sensitive.)

See also

Creating a Watcom SQL or Transact-SQL Function

Set up a new Watcom SQL or Transact-SQL function on a Sybase IQ server.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA or Resource authority.

- If you have DBA authority, you can create any type of function and assign any user as the owner.

- If you have Resource authority, you can create only SQL functions (not external functions), and all functions you create are owned by you.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Other Objects**.
3. Select **Functions**, click the drop-down arrow that appears to the right, and select **New**.
4. In the Create Function wizard, select the Sybase IQ resource on which the function is to reside.
5. Select a user or group to own the function.
6. Enter a name for the function.
7. Select the language or SQL dialect for the function.
8. Check **Use SQL template** to insert a hardcoded SQL template and skip the wizard page for specifying a return type. You can also set a Watcom SQL function to be deterministic or nondeterministic.
9. If you chose not to use the SQL template, define a return variable on the Specify Return Type page.
10. On the SQL page, edit the code provided by the wizard.
    a) Define input parameters: name, type, and default value.
    b) Enter function statements.
    c) Make any other changes needed to complete your function.
11. (Optional) Enter a comment describing the function.

    Comments can include both HTML and Javadoc tags, so you can incorporate them into your generated database documentation.
12. Click **Finish** to save the function.

The new function appears in the Functions list in the right pane of the Administration Console.

See also

Creating an External C/C++ Scalar or Aggregate Function

Set up a new external C or C++ aggregate function on a Sybase IQ server.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- (Optional) Copy the library that contains the external function to a location accessible to the Sybase IQ server. (You can create the function if the library is not accessible to Sybase IQ, but Sybase IQ cannot execute the function.)

Task

For detailed information on scalar and aggregate functions, see *User-Defined Functions* in the Sybase IQ documentation set.

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Other Objects**.
3. Select **Functions**, click the drop-down arrow that appears to the right, and select **New**.
4. In the Create Function wizard, select the Sybase IQ server on which the function is to reside.
5. Select a user or group to own the function.
6. Enter a name for the function.
7. Select the C/C++ scalar or aggregate function type.
8. Check **Use SQL template** to insert a hardcoded SQL template and skip the wizard pages for defining the return variable, function attributes, and clauses.
9. If you chose not to use the SQL template, define a return variable on the Specify Return Type page.
10. If you chose not to use the SQL template, configure the attributes of the function, including the names of its shared library file and descriptor function, on the Function Attributes page.
11. If you chose not to use the SQL template, configure clauses that control windows on the Function Clauses page.
12. On the SQL page, edit the code provided by the wizard.
    a) Define input parameters: name, type, and default value.
    b) If you chose to use the SQL template, replace the placeholder external name at the end of the query with a real name of the form functionDescriptorName@libraryName.
    c) Make any other changes needed to complete your function.
13. (Optional) Enter a comment describing the function.

Comments can include both HTML and Javadoc tags, so you can incorporate them into your generated database documentation.

14. Click **Finish** to save the function.

The new function appears in the Functions list in the right pane of the Administration Console.

See also
- *Listing Functions* on page 294
- *Creating a Watcom SQL or Transact-SQL Function* on page 294
- *Creating an External Java Function* on page 297
- *Viewing and Modifying a Function* on page 298
- *Granting or Revoking Permissions on a Function* on page 299
- *Generating DDL for a Function* on page 300
- *Deleting a Function* on page 301

Creating an External Java Function
Set up a new external Java function on a Sybase IQ server.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- (Optional) Install the required Java classes and JAR files in the database. (You can create the function if the classes and JAR files are not installed, but Sybase IQ cannot execute the function.)

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.

2. In the Administration Console, expand **IQ Servers > Other Objects**.

3. Select **Functions**, click the drop-down arrow that appears to the right, and select **New**.

4. In the Create Function wizard, select the Sybase IQ server on which the function is to reside.

5. Select a user or group to own the function.

6. Enter a name for the function.

7. Select the External Java function type.

8. Check **Use SQL template** to insert a hardcoded SQL template and skip the wizard page for specifying the return type.

9. If you chose not to use the SQL template, define a return variable on the Specify Return Type page.

10. On the SQL page, edit the code provided by the wizard.

    a) Define input parameters: name, type, and default value.

    b) If you chose to use the SQL template, replace the placeholder external name at the end of the query with a real name that follows the format of the placeholder.

    c) Make any other changes needed to complete your function.

11. (Optional) Enter a comment describing the function.

Comments can include both HTML and Javadoc tags, so you can incorporate them into your generated database documentation.

12. Click **Finish** to save the function.

The new function appears in the Functions list in the right pane of the Administration Console.

See also

- *Listing Functions* on page 294
- *Creating a Watcom SQL or Transact-SQL Function* on page 294
- *Creating an External C/C++ Scalar or Aggregate Function* on page 296
- *Viewing and Modifying a Function* on page 298
- *Granting or Revoking Permissions on a Function* on page 299
- *Generating DDL for a Function* on page 300
- *Deleting a Function* on page 301
- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309

Viewing and Modifying a Function

Examine the SQL statement and properties of a function on a Sybase IQ server. You can edit the SQL statement and change some property values.

Prerequisites

- Create a function.
- Authenticate with Sybase IQ using an account that has DBA authority or owns the function. (If you do not have authority to modify properties, Sybase Control Center displays the properties dialog in read-only mode.)

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.

2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Functions**.

3. In the right pane, select a function.

4. Click the drop-down arrow that appears next to the function name and select
Properties.

5. View the contents of the properties dialog. You can modify some values:

Page in Function Properties	Modifiable Values
General	Comment
Parameters	None
SQL	SQL code
Permission	Permissions

6. Click **OK** to save changes or click **Cancel** to dismiss the dialog without saving.

See also

Granting or Revoking Permissions on a Function

Grant or revoke permission to execute a function on a Sybase IQ server.

Prerequisites

Authenticate with Sybase IQ using an account that:

- has DBA authority, or
- has Perms Admin authority, or
- owns the function.

If you do not have authority to modify permissions on this function, Sybase Control Center
displays the Permissions page in read-only mode.

Task

1. In the Perspective Resources view, select the resource and select **Resource >
Administration Console**.

2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and
select **Functions**.

3. In the right pane, select a function to modify.

4. Click the drop-down arrow that appears next to the function name and select **Properties**.

5. In the left pane of the properties dialog, click **Permission**.

6. To assign permissions:

  a) Click **Grant**.

  b) Select users and groups from the list and click **OK**.

7. To cancel permissions, select a name from the permissions list and click **Revoke**.

8. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also

Generating DDL for a Function
Generate the data description language for creating a function on a Sybase IQ server. The DDL code can be a useful reference and training tool.

Prerequisites
Create a function.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.

2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Functions**.

3. In the right pane, select one or more functions.

4. Select **Resource > Generate DDL** from the Administration Console's menu bar.
A dialog displays the DDL statements.

5. Select and copy from the DDL dialog as needed; click **Close** when you are finished.

See also

- *Creating an External C/C++ Scalar or Aggregate Function* on page 296
- *Creating an External Java Function* on page 297
- *Viewing and Modifying a Function* on page 298
- *Granting or Revoking Permissions on a Function* on page 299
- *Deleting a Function* on page 301

Deleting a Function
Remove a function from a Sybase IQ server.

Prerequisites

- Create a function.
- Authenticate with Sybase IQ using the account that owns the function or an account with DBA authority.

Task

1. In the Perspective Resources view, select one or more resources and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Functions**.
3. In the right pane, select one or more functions.
4. Select **Resource > Delete** from the Administration Console's menu bar.
5. Click **Yes** to confirm the deletion.

See also
- *Listing Functions* on page 294
- *Creating a Watcom SQL or Transact-SQL Function* on page 294
- *Creating an External C/C++ Scalar or Aggregate Function* on page 296
- *Creating an External Java Function* on page 297
- *Viewing and Modifying a Function* on page 298
- *Granting or Revoking Permissions on a Function* on page 299
- *Generating DDL for a Function* on page 300

Manage Procedures
Create, delete, display, and manage the properties of procedures in Sybase IQ, including table-valued user-defined functions and table parameterized functions.

Sybase provides sample procedures and functions with Sybase IQ. Sample procedures are visible in Sybase Control Center in the Administration Console (IQ Servers > Other Objects > Procedures). The samples include explanatory comments and can be found in `IQ/IQ-X_X/ samples/udf`, where `X_X` is the Sybase IQ release number (15.4, for example).

Listing Procedures

Display a list of procedures available on Sybase IQ servers in this perspective.

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Other Objects** and select **Procedures**.
   A list of procedures in this perspective appears in the right pane.
3. (Optional) Enter a filter string in the field at the top of any column to display only procedures that match the string. You can filter on two or more columns at once. For example, to find procedures belonging to Fred and created in 2011, enter "fred" above the Owner column and "2011" above the Time Created column. (Filter strings are not case sensitive.)

See also

- *Creating a Procedure* on page 302
- *Creating a Table UDF or Table Parameterized Function* on page 304
- *Viewing and Modifying a Procedure* on page 305
- *Granting or Revoking Permissions on a Procedure* on page 306
- *Generating DDL for a Procedure* on page 307
- *Deleting a Procedure* on page 308
- *Common Display Options* on page 7

Creating a Procedure

Set up a new procedure on a Sybase IQ server.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA or Resource authority.
  - If you have DBA authority, you can create any type of procedure and assign any user as the owner.
  - If you have Resource authority, you can create only Watcom SQL or Transact-SQL procedures (not external procedures), and all procedures you create are owned by you.
- (Optional) If you are creating a Java procedure, install the required Java classes and JAR files in the database. (You can create the procedure if the classes and JAR files are not installed, but Sybase IQ cannot execute the procedure.)
- (Optional) If you are creating a procedure that relies on a non-Java external environment, copy the library that contains the external procedure to a location accessible to the Sybase IQ server. (You can create the procedure if the library is not accessible to Sybase IQ, but Sybase IQ cannot execute the procedure.)

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Other Objects**.
3. Select **Procedures**, click the drop-down arrow that appears to the right, and select **New**.
4. In the Create Procedure wizard, select the Sybase IQ resource on which the procedure is to reside.
5. Select a user or group to own the procedure.
6. Enter a name for the procedure.
7. Select the language or SQL dialect for the procedure. If you choose **External environment**, use the drop-down menu to select an environment.
8. Check **Use SQL template** to insert a hardcoded SQL template. Leave the box unchecked to retrieve generated SQL from the back end.
9. On the SQL page, edit the code provided by the wizard.
   a) For Watcom SQL and Transact-SQL, define input parameters in the format provided.
   b) For Watcom SQL, define the result in the format provided.
   c) For external languages, fill in the external name in the format provided.
   d) Enter procedure statements.
   e) Make any other changes needed to complete your procedure.
10. (Optional) Enter a comment describing the procedure.

    Comments can include both HTML and Javadoc tags, so you can incorporate them into your generated database documentation.
11. Click **Finish** to save the procedure.

The new procedure appears in the Procedures list in the right pane of the Administration Console.

See also
- *Listing Procedures* on page 302
- *Creating a Table UDF or Table Parameterized Function* on page 304
- *Viewing and Modifying a Procedure* on page 305
- *Granting or Revoking Permissions on a Procedure* on page 306
- *Generating DDL for a Procedure* on page 307
- *Deleting a Procedure* on page 308
- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309

Creating a Table UDF or Table Parameterized Function

Set up an external table-valued user-defined function or table parameterized function in C/C
++ or an external table-valued user-defined function in Java on a Sybase IQ server.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- (Optional) If you are creating a Java procedure, install the required Java classes and JAR
  files in the database. (You can create the procedure if the classes and JAR files are not
  installed, but Sybase IQ cannot execute the procedure.)
- (Optional) If you are creating a C/C++ procedure, copy the library that contains the
  procedure to a location accessible to the Sybase IQ server. (You can create the procedure if
  the library is not accessible to Sybase IQ, but Sybase IQ cannot execute the procedure.)

Task

For detailed information on UDFs, see *User-Defined Functions* in the Sybase IQ
documentation set.

1. In the Perspective Resources view, select the resource and select **Resource >
   Administration Console**.
2. In the Administration Console, expand **IQ Servers > Other Objects**.
3. Select **Procedures**, click the drop-down arrow that appears to the right, and select **New
   Table UDF/TPF**.
4. In the Create Table UDF/TPF wizard, select the Sybase IQ server on which the procedure
   is to reside.
5. Select a user or group to own the procedure.
6. Enter a name for the procedure.
7. Select the type of procedure to create.
8. Check **Use SQL template** if you want the wizard to provide comments with syntax
   examples to help you define parameters and results. Leave the box unchecked if you want
   to use the Return Structure page to define the result.
9. If you chose not to use the SQL template, define the result on the Return Structure page.
10. On the SQL page, edit the code provided by the wizard.
    a) Define input parameters using the format provided.
    b) Replace the placeholder external name at the end of the query with a real name that
       follows the format of the placeholder.
    c) Make any other changes needed to complete your procedure.
11. (Optional) Enter a comment describing the procedure.

Comments can include both HTML and Javadoc tags, so you can incorporate them into your generated database documentation.

12. Click **Finish** to save the procedure.

The new procedure appears in the Procedures list in the right pane of the Administration Console.

See also
*   *Listing Procedures* on page 302
*   *Creating a Procedure* on page 302
*   *Viewing and Modifying a Procedure* on page 305
*   *Granting or Revoking Permissions on a Procedure* on page 306
*   *Generating DDL for a Procedure* on page 307
*   *Deleting a Procedure* on page 308
*   *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309

Viewing and Modifying a Procedure

Examine the SQL statement and properties of a procedure on a Sybase IQ server. You can edit the SQL statement and change some property values.

Prerequisites

*   Create a procedure.
*   Authenticate with Sybase IQ using an account that has DBA authority, Resource authority, or owns the function. (If you do not have authority to modify properties, Sybase Control Center displays the properties dialog in read-only mode.)

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Procedures**.
3. In the right pane, select a procedure.
4. Click the drop-down arrow that appears next to the procedure name and select **Properties**.
    The properties dialog opens.
5. View the contents of the properties dialog. You can modify some values:

Page in Procedure Properties	Modifiable Values
General	Comment

Page in Procedure Properties	Modifiable Values
Parameters	None
SQL	SQL code
Permission	Permissions

6. Click **OK** to save changes or click **Cancel** to dismiss the dialog without saving.

See also
- *Listing Procedures* on page 302
- *Creating a Procedure* on page 302
- *Creating a Table UDF or Table Parameterized Function* on page 304
- *Granting or Revoking Permissions on a Procedure* on page 306
- *Generating DDL for a Procedure* on page 307
- *Deleting a Procedure* on page 308

Granting or Revoking Permissions on a Procedure
Grant or revoke permission to execute a procedure on a Sybase IQ server.

Prerequisites
Authenticate with Sybase IQ using an account that:

- has DBA authority, or
- has Perms Admin authority, or
- owns the procedure.

If you do not have authority to modify permissions on this procedure, Sybase Control Center displays the Permissions page in read-only mode.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.

2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Procedures**.

3. In the right pane, select a procedure to modify.

4. Click the drop-down arrow that appears next to the procedure name and select **Properties**.

5. In the left pane of the properties dialog, click **Permission**.

6. To assign permissions:

   a) Click **Grant**.

b) Select users and groups from the list and click **OK**.

7. To cancel permissions, select a name from the permissions list and click **Revoke**.

8. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also
- *Listing Procedures* on page 302
- *Creating a Procedure* on page 302
- *Creating a Table UDF or Table Parameterized Function* on page 304
- *Viewing and Modifying a Procedure* on page 305
- *Generating DDL for a Procedure* on page 307
- *Deleting a Procedure* on page 308

Generating DDL for a Procedure
Generate the data description language for creating a procedure on a Sybase IQ server. The DDL code can be a useful reference and training tool.

Prerequisites
Create a procedure.

Task

1. In the Perspective Resources view, select the resource and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Procedures**.
3. In the right pane, select one or more procedures.
4. Select **Resource > Generate DDL** from the Administration Console's menu bar.
   A dialog displays the DDL statements.
5. Select and copy from the DDL dialog as needed; click **Close** when you are finished.

See also
- *Listing Procedures* on page 302
- *Creating a Procedure* on page 302
- *Creating a Table UDF or Table Parameterized Function* on page 304
- *Viewing and Modifying a Procedure* on page 305
- *Granting or Revoking Permissions on a Procedure* on page 306
- *Deleting a Procedure* on page 308

Deleting a Procedure
Remove a procedure from a Sybase IQ server.

Prerequisites

*   Create a procedure.
*   Authenticate with Sybase IQ using the account that owns the procedure or an account with DBA authority.

Task

1. In the Perspective Resources view, select one or more resources and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **IQ Servers > Other Objects** and select **Procedures**.
3. In the right pane, select one or more procedures.
4. Select **Resource > Delete** from the Administration Console's menu bar.
5. Click **Yes** to confirm the deletion.

See also
*   *Listing Procedures* on page 302
*   *Creating a Procedure* on page 302
*   *Creating a Table UDF or Table Parameterized Function* on page 304
*   *Viewing and Modifying a Procedure* on page 305
*   *Granting or Revoking Permissions on a Procedure* on page 306
*   *Generating DDL for a Procedure* on page 307

Manage External Environments
Install, delete, update, and modify the components of external development environments.

Sybase Control Center currently supports Java external environments.

Listing External Environments
Display a list of the external environments supported by servers in this perspective.

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers** and select **External Environments**. A list of external environments available in this perspective appears in the right pane.

See also
*   *Working in a Java External Environment* on page 309

Working in a Java External Environment

Install and manage Java classes and JAR files. Test a Java external environment and modify its properties.

See also

- *Listing External Environments* on page 308

Installing a Java Class or JAR File into a Sybase IQ Database

Install a Java class or JAR file to a Sybase IQ database. This enables the database to execute functions and procedures written in Java.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- Copy the Java class or JAR file to a location accessible to the Sybase IQ server.

Task

A Java external environment appears in the list of external environments in the Sybase Control Center Administration Console even when none of its classes or JAR files have been installed.

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > External Environments > Java**.
3. Select either:
   - **Classes** or
   - **JARs**

   Then click the drop-down arrow that appears to the right and select **New**.
4. In the installation dialog, select a Sybase IQ resource on which to install the class or JAR.
5. If you select a multiplex, the dialog displays a new menu. Use it to select the server in the multiplex on which to install the class or JAR.
6. If the SCC agent on the selected server is not registered, click **Register Agent** and enter the required registration information.

   The **Register Agent** button is not active if the server has been registered.
7. If the SCC agent on the selected server is not authenticated, click **Authenticate Agent**. Enter a user ID, password, and port for SCC to use when it logs in to the SCC agent.

   The **Authenticate Agent** button is not active if the server has been authenticated.
8. Enter the path to the Java class file. It must be on a file system accessible to the Sybase IQ resource.
9. If you are installing a JAR, enter its name.

10. (Optional) Enter a comment describing the class or JAR.

11. If you are installing a JAR, select classes to install.

12. Click **Finish** to install the class or JAR.

See also
- *Modifying the Properties of a Java Class or JAR File* on page 310
- *Updating a Java Class or JAR File* on page 311
- *Deleting a Java Class or JAR File* on page 311
- *Testing the Configuration of a Java External Environment* on page 312
- *Modifying the Properties of a Java External Environment* on page 313
- *Registering and Authenticating a Sybase Control Center Agent* on page 123

Modifying the Properties of a Java Class or JAR File
Display and change the properties of a Java class or JAR file associated with a Sybase IQ database.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA authority. (If you do not have authority to modify properties, Sybase Control Center displays the properties dialog in read-only mode.)

Task

1. From the application menu bar, select **View > Open > Administration Console**.

2. In the Administration Console, expand **IQ Servers > External Environments > Java** and select either:

   - **Classes** or
   - **JARs**

3. Select a class or JAR in the right pane, click the drop-down arrow that appears to the right, and select **Properties**.

4. In the properties dialog, edit the comments field. All other properties are read-only.

5. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also
- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309
- *Updating a Java Class or JAR File* on page 311
- *Deleting a Java Class or JAR File* on page 311
- *Testing the Configuration of a Java External Environment* on page 312
- *Modifying the Properties of a Java External Environment* on page 313

Updating a Java Class or JAR File

Update a Java class or a JAR file in a Sybase IQ database.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority.
- Put an updated version of the Java class or JAR file in a location accessible to the Sybase IQ server.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > External Environments > Java**.
3. Select either:
   - **Classes** or
   - **JARs**
4. Select a class or JAR in the right pane, click the drop-down arrow that appears to the right, and select **Update**.
5. If the **Register Agent** button is active in the update dialog, click it to register the SCC agent on the class or JAR's Sybase IQ server. Enter the required registration information.
6. If the **Authenticate Agent** button is active, click it to authenticate the SCC agent. Enter a user ID, password, and port for SCC to use when it logs in to the SCC agent.
7. Enter the path to the updated class or JAR file. It must be on a file system accessible to the Sybase IQ resource.
8. If you are updating a JAR, click **Next** and select classes to update.
9. Click **Finish** to update the class or JAR.

See also
- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309
- *Modifying the Properties of a Java Class or JAR File* on page 310
- *Deleting a Java Class or JAR File* on page 311
- *Testing the Configuration of a Java External Environment* on page 312
- *Modifying the Properties of a Java External Environment* on page 313
- *Registering and Authenticating a Sybase Control Center Agent* on page 123

Deleting a Java Class or JAR File

Remove a Java class or JAR file from a Sybase IQ database.

Prerequisites

Authenticate with Sybase IQ using an account that has DBA authority.

Task

You cannot use this procedure to remove a class from a JAR. Instead, update the JAR.

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > External Environments > Java**.
3. Select either:
   - **Classes** or
   - **JARs**
4. Select a class or jar in the right pane, click the drop-down arrow that appears to the right, and select **Delete**.
5. Review the confirmation dialog and click **Yes**.

See also
- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309
- *Modifying the Properties of a Java Class or JAR File* on page 310
- *Updating a Java Class or JAR File* on page 311
- *Testing the Configuration of a Java External Environment* on page 312
- *Modifying the Properties of a Java External Environment* on page 313

Testing the Configuration of a Java External Environment
Test the configuration of a Java external environment to make sure the JRE is accessible and responding.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA authority.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers** and select **External Environments**.
3. In the right pane, select a Java external environment and click the drop-down arrow that appears to the right.
4. Select **Test Configuration**.
   The Test Configuration dialog displays progress messages to report stopping and then starting the external environment.
5. If the test fails:
   - The **iq_java.sh** script in the Sybase IQ server's `bin` directory might have failed to execute. Run the script manually to check for errors.
   - The JRE might have been deleted or moved to a new location. To specify a new location:

1. In the Administration Console, select a Java external environment and click the drop-down arrow that appears to the right.
2. Select **Properties**.
3. In the **Location** field, enter the new path to the Java executable (`java.exe` or `java`).
4. From the **User** drop-down, select the Sybase IQ user or group that the external environment uses to log in to the database. The user or group must have DBA authority; the default is DBA.

See also
- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309
- *Modifying the Properties of a Java Class or JAR File* on page 310
- *Updating a Java Class or JAR File* on page 311
- *Deleting a Java Class or JAR File* on page 311
- *Modifying the Properties of a Java External Environment* on page 313

Modifying the Properties of a Java External Environment
Display and change the properties of a Java external environment, including the location of the executable and the associated user.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA authority. (If you do not have authority to modify properties, Sybase Control Center displays the properties dialog in read-only mode.)

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers** and select **External Environments**.
3. In the right pane, select a Java external environment and click the drop-down arrow that appears to the right.
4. Select **Properties**.
5. In the properties dialog, modify any of these:

Options	Description
Location	The path to the Java executable, `java.exe` or `java`.
User	The Sybase IQ user or group that the external environment uses to log in to the database. The user or group must have DBA authority; the default is DBA.
Comment	Notes describing the external environment.

6. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also

- *Installing a Java Class or JAR File into a Sybase IQ Database* on page 309
- *Modifying the Properties of a Java Class or JAR File* on page 310
- *Updating a Java Class or JAR File* on page 311
- *Deleting a Java Class or JAR File* on page 311
- *Testing the Configuration of a Java External Environment* on page 312

Manage Text Indexes

Create, delete, generate DDL, and manage the properties of text indexes and text configuration objects.

Text indexes significantly speed up full-text searching. You must configure a text index for a table before it is subjected to a full-text search.

Each text index requires a text configuration object. Sybase provides two default text configuration objects, one for CHAR collation, one for NCHAR collation. When you set up a text index, you can use one of the default text configuration objects or create a custom text configuration object.

For detailed information on text indexes, see *Unstructured Data Analytics in Sybase IQ*.

Creating a Text Index

Set up a text index for a table in Sybase IQ.

Prerequisites

- Authenticate with Sybase IQ using an account that has DBA authority or owns the table being indexed.
- (Optional) Create a text configuration object for your text index.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects**.
3. Select **Text Indexes**, click the drop-down arrow that appears to the right, and select **New**.
4. In the dialog, select the Sybase IQ server on which the table to be indexed resides.
5. Select the table.
6. Enter a name for the text index.
7. Select columns to include in the text index.

8. Select a text configuration object.
    If the index is on an IQ table in the column store, skip to step *10*. If the index is on a table in the catalog store, you see a page that lets you select a refresh type.

9. Select a refresh type (catalog store tables only). If you select Automatic, set the interval at which the text index refreshes.

    Note: If the index is on an IQ table in the column store, it refreshes automatically whenever the data in the underlying table changes.

10. Select the dbspace in which to store the text index.

11. (Optional) Enter a comment describing the text index.

12. Click **Finish** to create the text index.

See also
*   *Modifying the Properties of a Text Index* on page 315
*   *Generating DDL for a Text Index* on page 316
*   *Refreshing a Text Index* on page 317
*   *Truncating a Text Index* on page 318
*   *Deleting a Text Index* on page 319

Modifying the Properties of a Text Index
Display and change the properties of a text index for a Sybase IQ table (column store) or a catalog store table (row store).

Prerequisites
*   Create a text index.
*   Authenticate with Sybase IQ using an account that has DBA authority or owns the table being indexed. (If you do not have authority to modify properties, Sybase Control Center displays the properties dialog in read-only mode.)

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Indexes**.
3. In the right pane, select a text index and click the drop-down arrow that appears to the right.
4. Select **Properties**.
5. Select the pages of the Text Index Properties dialog (**General** and **Columns**) to display properties.

    You can change the name, comment, and refresh type (unless it is set to Immediate). If refresh type is set to Automatic, you can also set the refresh interval and units.

Some properties cannot be modified. These include, on the General page, the text index's type, table, dbspace, text configuration object, and (if it is set to Immediate) refresh type. You cannot change anything on the Columns page.

6. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also

Generating DDL for a Text Index
Generate data description language for one or more text indexes. The DDL code can be a useful reference and training tool.

Prerequisites
Create a text index.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Indexes**.
3. In the right pane, select one or more text indexes.
4. Select **Resource > Generate DDL** from the Administration Console's menu bar.
   A dialog displays the DDL statements.
5. Select and copy from the DDL dialog as needed; click **Close** when you are finished.

See also

Refreshing a Text Index

Refresh a text index for a catalog store table to bring it up to date when the underlying data has changed.

Prerequisites

- Create a text index for a catalog store table (row store).
- Authenticate with Sybase IQ using an account that has DBA authority or owns the table being indexed.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Indexes**.
3. In the right pane, select a text index and click the drop-down arrow that appears to the right.
4. Select **Properties**.
5. Select the **General** page of the properties dialog.
6. (Optional) Check the last refresh time displayed in the middle of the page.
7. Click **Refresh Now** to refresh the text index.

   If the data has not changed since the last refresh, the refresh date does not change.

   To force a refresh, truncate the text index, then refresh it.

   If the text index cannot be manually refreshed, the **Refresh Now** button is grayed out.
8. In the Text Index Refresh Data dialog, select the isolation level for the refresh:

Options	Description
Read uncommitted (level 0)	• Can read rows with or without write lock • Applies no read locks • Data may change during the refresh • Allows dirty reads, nonrepeatable reads, and phantom rows
Read committed (level 1)	• Can read only rows with no write lock • Read-locks only the current row and releases it immediately after reading • Data may change during the refresh • Prevents dirty reads • Allows nonrepeatable reads and phantom rows

Options	Description
Repeatable read (level 2)	• Can read only rows with no write lock • Read-locks each row as it is read; holds the lock until the refresh is done • Prevents dirty reads and nonrepeatable reads • Allows phantom rows
Serializable (level 3)	• Can read only rows with no write lock • Read-locks every row for the duration of the refresh operation • Prevents dirty reads, nonrepeatable reads, and phantom rows
Snapshot	• Applies no read locks • Can read any row • The database takes a snapshot of committed data when the refresh operation reads the first row
Share mode (the default)	• Allows other transactions to read the underlying table during the refresh operation • Uses shared table locks
Exclusive mode	• Does not change the isolation level • Locks the underlying table to ensure that the text index is updated to be consistent with committed data in the table • If an exclusive table lock cannot be obtained, the refresh fails

9. Click **OK** to refresh the text index.

See also

Truncating a Text Index
Truncate a text index for a catalog store table.

Prerequisites

- Create a text index for a catalog store table (row store).
- Authenticate with Sybase IQ using an account that has DBA authority or owns the table being indexed.

Task

Truncating lets you delete data from a text index without dropping the text index definition. For example, to modify the stoplist for a text index:

- Truncate the text index.
- Edit the stoplist in the text configuration object associated with the text index.
- Refresh the text index to bring in the new stoplist.

If a text index is set to immediate refresh, you cannot truncate it. Instead, drop it and create a new one.

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Indexes**.
3. In the right pane, select a text index and click the drop-down arrow that appears to the right.
4. Select **Properties**.
5. Select the **General** page of the properties dialog.
6. Click **Truncate Now** to truncate the text index.

   If the text index cannot be truncated, the **Truncate Now** button is grayed out.

See also
- *Creating a Text Index* on page 314
- *Modifying the Properties of a Text Index* on page 315
- *Generating DDL for a Text Index* on page 316
- *Refreshing a Text Index* on page 317
- *Deleting a Text Index* on page 319

Deleting a Text Index
Remove one or more text indexes.

Prerequisites

- Create a text index.
- Authenticate with Sybase IQ using an account that has DBA authority or owns the table being indexed.

Task

1. From the application menu bar, select **View > Open > Administration Console**.

2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Indexes**.

3. In the right pane, select one or more text indexes.

4. Select **Resource > Delete** from the Administration Console's menu bar.

5. Click **Yes** to confirm the deletion.

See also
- *Creating a Text Index* on page 314
- *Modifying the Properties of a Text Index* on page 315
- *Generating DDL for a Text Index* on page 316
- *Refreshing a Text Index* on page 317
- *Truncating a Text Index* on page 318

Manage Text Configuration Objects

Create, delete, generate DDL, and manage the properties of text configuration objects.

For detailed information on text configuration objects, see *Unstructured Data Analytics in Sybase IQ* in the Sybase IQ documentation set.

Creating a Text Configuration Object
Create a text configuration for use with a text index.

Prerequisites
Authenticate with Sybase IQ using an account that has DBA or Resource authority.

- If you have DBA authority, you can assign any user as the owner of text configuration objects you create.
- If you have Resource authority, all text configuration objects you create are owned by you.

Task

For detailed information on stoplists, see the Sybase IQ documentation: *Unstructured Data Analytics in Sybase IQ > TEXT Indexes and Text Configuration Objects > Text Configuration Objects > Text Configuration Object Settings > Stoplist Setting (STOPLIST).*

1. From the application menu bar, select **View > Open > Administration Console**.

2. In the Administration Console, expand **IQ Servers > Schema Objects**.

3. Select **Text Configuration Objects**, click the drop-down arrow that appears to the right, and select **New**.

4. In the wizard, select the Sybase IQ server on which to create the text configuration.

5. Select the Sybase IQ user who will own the text configuration.

6. Enter a name for the text configuration.
7. Select the database collation for this text configuration. Each database has a CHAR collation and an NCHAR collation, configured when the database was created. Sample collations:

   - ISO_BINENG – Binary ordering, English ISO/ASCII 7-bit letter case mappings
   - UCA – Standard default Unicode Collation Algorithm collation

8. (Optional) Enter a comment describing the text configuration.
9. Choose a term breaker algorithm and specify the minimum and maximum term lengths.
10. (Optional) If you chose the Generic term breaker algorithm, you can specify an external library function to break the text into terms. Use one of these formats:

    - function-name@library-file-name
    - Windows-function-name@library-file-name.dll
    - UNIX:UNIX-function-name@library-file-name.so

11. (Optional; available only for CHAR collations) Specify an external library to perform document filtering before term breaker processing.
12. (Optional) Create a stoplist by entering terms to omit from the text index. (Terms in the stoplist are also ignored in queries.) Separate terms with spaces.

    Many nonalphanumeric characters are ignored in stoplists; others (including spaces, apostrophes, and dashes) are interpreted as term delimiters. Consequently, including contractions and hyphenated terms may lead to undesirable results, even when you enclose the terms in quotes. How the stoplist is parsed depends on the term breaker and term lengths you specified in step *9* and step *10*.

13. Click **Finish** to create the text configuration.

Next
Create a text index that uses your new text configuration.

See also

Modifying the Properties of a Text Configuration Object
Display and change the properties of a text configuration object for a text index.

Prerequisites

- Create a text configuration object.

• Authenticate with Sybase IQ using an account that has DBA authority or owns the text configuration object. (If you do not have authority to modify properties, Sybase Control Center displays the properties dialog in read-only mode.)

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Configuration Objects**.
3. In the right pane, select a text configuration object and click the drop-down arrow that appears to the right.
4. Select **Properties**.
5. Select the pages of the properties dialog (**General**, **Settings**, **Stoplist**, and **Options**) to display the properties you want to modify.

   Some properties, including the text configuration's name, type, owner, and collation (on the General page) and the time and date formats (on the Options page), cannot be modified.
6. Click **Apply** to save your changes and keep editing, or click **OK** to save and close the properties dialog.

See also

Generating DDL for a Text Configuration Object

Generate data description language for one or more text configuration objects. The DDL code can be a useful reference and training tool.

Prerequisites

Create a text configuration object.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Configuration Objects**.
3. In the right pane, select one or more text configuration objects.
4. Select **Resource > Generate DDL** from the Administration Console's menu bar.
   A dialog displays the DDL statements.
5. Select and copy from the DDL dialog as needed; click **Close** when you are finished.

See also
- *Creating a Text Configuration Object* on page 320
- *Modifying the Properties of a Text Configuration Object* on page 321
- *Deleting a Text Configuration Object* on page 323

Deleting a Text Configuration Object
Remove one or more text configuration objects.

Prerequisites

- Create a text configuration object.
- Authenticate with Sybase IQ using an account that has DBA authority or owns the text configuration object.

Task

1. From the application menu bar, select **View > Open > Administration Console**.
2. In the Administration Console, expand **IQ Servers > Schema Objects** and select **Text Configuration Objects**.
3. In the right pane, select one or more text configuration objects.

   Note: You cannot delete the default text configuration objects, default_char and default_nchar.
4. Select **Resource > Delete** from the Administration Console's menu bar.
5. Click **Yes** to confirm the deletion.

See also
- *Creating a Text Configuration Object* on page 320
- *Modifying the Properties of a Text Configuration Object* on page 321
- *Generating DDL for a Text Configuration Object* on page 322

Executing a SQL Query
Execute an ad hoc SQL query or a stored procedure against one or more Sybase IQ servers.

You can use the Execute SQL view to execute any valid SQL statement, including queries and stored procedures. Anyone can launch a query; no permissions are required. However, if you do not have authority to perform the actions in the query, you will see an error.

1. In the Perspective Resources view, select one or more resources and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select **IQ Servers**.
3. In the right pane, select the Sybase IQ resources.

4. From the Administration Console menu bar, select **Resource > Execute SQL**.
The Execute SQL view opens.

5. Enter a query or the name of a stored procedure in the SQL Statements box and click
Execute.
The query runs on all the Sybase IQ servers you selected and results appear in the bottom
portion of the view. If you selected more than one server, the view includes a results tab for
each simplex or multiplex node. On the tabs,

- A green check indicates a successful query.
- A red X indicates an error.

Troubleshoot Sybase Control Center for Sybase IQ

Troubleshoot problems that occur in Sybase Control Center for Sybase IQ.

Authenticating a Chinese or Japanese Sybase IQ Server Fails

Problem: When authenticating a Chinese or Japanese Sybase IQ server, if the login name or password contain Chinese or Japanese characters, the login fails.

Solution:

1. In the Perspective Resources window, select the Chinese or Japanese server, click the arrow, and select **Properties**.
2. In the Resource Properties window, select **Connection**.
3. Enter the **Character set** used on the Sybase IQ server.
4. Click **OK**.

Disabled Features When You Lack Permissions

Sybase Control Center for Sybase IQ has several ways to indicate that you do not have the required permissions to perform administration tasks.

- Hidden/disabled Administration Console menu items – If you do not have the required permissions, the corresponding menu items in the Administration Console window will either not show, or are dimmed. For example, the **Start Server** menu item is dimmed if the Sybase Control Center agent is not registered, authenticated, and running.
- Warning messages in wizards – Creation wizards, such as the create logical server wizard, check if the you have the corresponding object creation permission when you select a resource on the first page of the wizard. If you do not, a message appears at the bottom of the wizard window, and the **Next** and **Finish** buttons are disabled.
- Popup error messages – Although comparatively rare, there are certain cases where the permission validation cannot be performed until you execute a task. When this occurs, a popup error message appears describing the lack of permissions.

To gain permissions, you must belong to the group SCC_MONITOR or have the appropriate authority for the task.

See also
- *Configuring Sybase IQ for Monitoring* on page 116

Troubleshooting Invalid Database Definitions

You cannot use the Create Databases wizard until you fix any invalid parameters.

If any of these errors exist, the Create Databases wizard flags the database definition with a red **x**:

Issue Causing Invalid Database Definition
Agent not running on port specified.
Database path points to a file that already exists.
Database path points to a location that is not writeable.
Specified IQ port is already in use.
Main dbspace path points to a file that already exists.
Main dbspace path points to a location that is not writeable.
Another database definition with the same name already exists in the wizard list of databases to be created.
Temp dbspace path points to a file that already exists.
Temp dbspace path points to a location that is not writeable.

See also
*   *Creating a Database* on page 232

SCC Fails to Control Multiplex Servers

Problem: On a Sybase IQ multiplex installed on a shared disk cluster, you are unable to perform control tasks like starting and stopping the nodes.

This problem occurs when multiple nodes in a multiplex run from the same installation of Sybase IQ. Sybase Control Center monitoring tasks are not affected.

Solution: Use the Sybase IQ installer to install the SCC agent on each machine in the multiplex environment. SCC needs a locally installed agent to perform control operations.

Problems with Basic Sybase Control Center Functionality

Troubleshoot problems that involve basic features like starting and stopping, authentication, alerts, and scheduling.

Cannot Log In

Problem: Cannot log in to Sybase Control Center Web console.

Solution: Make sure that Sybase Control Center has been configured:

- To allow logins through the operating system
- To grant appropriate roles to your login account

Ask the Sybase Control Center administrator to help you check.

See also
- *User Authorization* on page 106
- *Setting Up Security* on page 87

Sybase Control Center Fails to Start

Problem: The Sybase Control Center server does not start.

Solution 1: Port conflict
Solution: SCC might be using one or more ports that are also being used by another server or application on this machine. To check for port conflicts:

1. Execute this command:
   ```
   scc --info ports
   ```

   The command lists all the ports on which Sybase Control Center and its services listen, indicates whether each port is in use, and shows the service running on each port. If SCC is not running, any port shown to be in use represents a conflict.
2. If you discover a conflict, use **scc --port** to change the port used by the Sybase Control Center service.

Solution 2: Insufficient memory
You might see this error why you try to start: `Could not create the Java Virtual machine.` Increase the maximum memory setting.

See also
- *Configuring Ports* on page 102
- *Configuring Memory Usage* on page 78

Browser Refresh (F5) Causes Logout

Problem: Pressing the **F5** key to refresh your browser logs you out of Sybase Control Center.

Solution: Do not use **F5** when you are logged in to Sybase Control Center. Browser refresh does not refresh data inside Sybase Control Center, but refreshes the loaded application or pages in the browser—in this case, the Adobe Flash on which Sybase Control Center is built. Consequently, pressing **F5** logs you out of any servers you are currently logged in to, including Sybase Control Center.

Alerts Are Not Generated

Problem: Alerts are not being generated in Sybase Control Center.

Solution: Schedule a job to run the data collection that supports your alerts. See the data collections topic for your Sybase Control Center product module for information on which collections must be scheduled.

See also
- *Setting Up Statistics Collection* on page 128

Performance Statistics Do Not Cover Enough Time

Problem: I want to graph performance counters over a long period of time but the statistics chart displays only very recent data.

Solution: Ask your Sybase Control Center administrator to change the repository purging options to keep statistical data available for as long as you need it. By default, statistics are purged frequently to conserve disk space.

See also
- *Configuring Repository Purging* on page 194
- *Graphing Performance Counters* on page 160

Resetting the Online Help

Problem: Sybase Control Center online help is corrupted or cannot be found (404 error).

Solution: Clear online help files to force SCC to build new ones.

1. Shut down Sybase Control Center.
2. Remove this directory:
   ```
   <SCC-installation-directory>\SCC-3_2\services
   \EmbeddedWebContainer\container\Jetty-6.1.22\work
   \Jetty_0_0_0_0_8282_help.war__help__.smpe97
   ```

   Tip: In Windows, you might see a deletion error. Regardless of what the errors says, it might be caused by the length of the path. If deletion fails, rename the

Jetty_0_0_0_0_8282_help.war__help__.smpe97 folder to something very short, such as J. Then delete the renamed folder.

3. Remove these files:

```
<SCC-installation-directory>\SCC-3_2\services
\EmbeddedWebContainer\container\Jetty-6.1.22\contexts
\_help.xml
<SCC-installation-directory>\SCC-3_2\services
\SybaseControlCenter\help\com.sybase.infocenter.scc.zip
<SCC-installation-directory>\SCC-3_2\services
\SybaseControlCenter\help\help.war
<SCC-installation-directory>\SCC-3_2\services
\SybaseControlCenter\help\help_info.xml
```

4. Start SCC. After the server comes up it rebuilds the help, which takes a few minutes.

5. To display the help, go to `https://<your-SCC-host>:8283/help/index.jsp`.

Note: If you try to display the help too soon after restarting, you get a file not found error. Wait a minute or two and try again.

Data Collections Fail to Complete

Problem: A collection frequently times out or generates errors citing the REJECT_DUPLICATE_RESOURCE_AND_COLLECTION policy, but no problems with the monitored resources are evident.

The errors appear in the log and on the collection history screen.

Solution: Try to determine why the collection is taking so long. For example, are network delays slowing down traffic between Sybase Control Center and the monitored server?

In the case of network delays and other resource-related problems, the interval between collections might be shorter than the time needed to finish the collection. To fix this problem, increase the time between collections.

See also
- *Modifying the Data Collection Interval for a Job* on page 165

Memory Warnings at Startup

Problem: When Sybase Control Center starts, you see warnings about system memory or heap memory allocation.

Solution: Increase the maximum memory setting (*SCC_MEM_MAX* or `jvmopt=-Xmx`).

See also
- *Configuring Memory Usage* on page 78

OutOfMemory Errors

Problem: Sybase Control Center generates `OutOfMemory` errors.

Solution:

- If the `OutOfMemory` error says that Sybase Control Center is out of heap space, increase the maximum memory setting (*SCC_MEM_MAX* or `jvmopt=-Xmx`).
- If the `OutOfMemory` error says that Sybase Control Center is out of permanent generation space, increase the permanent memory setting (*SCC_MEM_PERM* or `jvmopt=-XX:MaxPermSize`).
- Repeated `OutOfMemory` errors may indicate a memory leak. `OutOfMemory` errors generate heap dumps:
  - When Sybase Control Center runs as a service in Windows:
    `C:/windows/system32`
  - When Sybase Control Center runs as a service in UNIX:
    `<SCC-install-directory>/SCC-3_2/bin`

Send the heap dump files to Sybase technical support for analysis.

See also
- *Configuring Memory Usage* on page 78

Glossary: Sybase Control Center for Sybase IQ

Glossary of Sybase Control Center terms related to Sybase IQ.

alert – a mechanism for notifying administrators when a managed resource experiences a status change, or when a performance metric passes a user-specified threshold.

alert notification – an indication that an alert has fired. Alert notifications appear in the Alert Monitor view. If e-mail notification is enabled, alert notifications are also delivered to the specified e-mail address.

alert storm – the result of issuing many redundant alerts associated with a common or root occurrence. See also alert storm suppression.

alert storm suppression – a Sybase Control Center feature that can be configured to prevent alert storms by suppressing repeat alert notifications for a specified period of time.

alert type – the basis on which an alert fires: state or threshold. State alerts are triggered by the state of their key performance indicator (for example, running or stopped), while threshold alerts are triggered when their KPI's numerical value passes a specified threshold.

authenticate – when SCC authenticates with a managed resource, it logs in to the resource with a user ID and password provided by you. SCC must log in to managed resources in order to gather performance statistics and perform management tasks. You can choose to have SCC use your current SCC login ID, or you can provide different credentials.

availability – indicates whether a resource is accessible and responsive.

catalog store – the portion of each Sybase IQ database that contains its metadata. (Metadata describes the layout of the Sybase IQ tables, columns, and indexes.) The catalog store contains the system dbspace and up to 12 additional catalog dbspaces. The default name for this file is *<dbname>.db*.

chart trend period – the period, in minutes, over which data is displayed in historical charts. Set the chart trend period on the Settings screen of the Sybase IQ Monitoring View. Contrast with screen refresh interval.

collection repeat interval – the period, in seconds, minutes, hours, or days, between successive repetitions of a statistics collection job. The collection repeat interval determines how often new data on historical monitoring screens is available to be refreshed. Set the collection repeat interval in the scheduler. See also screen refresh interval.

collection – a named, predefined set of key performance indicators for which values are collected from monitored servers at the same time. Collections supply the performance and availability data shown on Sybase Control Center screens and charts. Use the scheduler to

view a list of collections and to control which collections run, how often they run, and the length of time for which they run.

connection – a connection from a Sybase IQ server to a database.

database – a collection of tables that are related by primary and foreign keys. The tables hold the information in the database. The tables and keys together define the structure of the database. Sybase IQ databases are specially indexed to take advantage of the query speed of Sybase IQ.

dbspace – a named collection of DB files that provides space for data and can be administered as a logical subset of the total storage. The main store, catalog store, and temporary store consist of dbspaces.

event – an activity in the system, such as a user logging in, a service starting or stopping, or a condition changing. Use the alerts feature to detect and notify you about system events.

external environment – a development environment external to Sybase IQ (C/C++ or Java, for example) that you can use to create functions and procedures to run against Sybase IQ databases.

heat chart – a graphical view of resource availability and selected performance and status metrics for all the registered resources in the current perspective.

instance – an SCC agent or server run from a shared disk installation. See also shared-disk mode.

job – a task performed by the scheduler in Sybase Control Center.

key performance indicator (KPI) – a single metric used to evaluate the status or performance of a monitored resource. A KPI value can be a state (such as running, error, or stopped) or a numerical value. KPIs are grouped into collections (and also, for some product modules, into key performance areas, or KPAs). KPI values are collected by scheduled collection jobs and appear on monitoring screens and in the statistics and heat charts. Examples of KPIs are resource state and CPU usage.

key performance area (KPA) – a group of related key performance indicators.

main store – the Sybase IQ main store is the portion of each Sybase IQ database that contains persistent database structures, such as backup metadata and rollback data for committed transactions.

managed resource – a server, agent, or other entity monitored and administered by Sybase Control Center. Resources SCC can manage include Adaptive Server, Replication Server, Replication Agent, Mirror Replication Agent, and Sybase IQ.

multiplex – a powerful feature in Sybase IQ that provides application scalability through a clustered server configuration. A Sybase IQ multiplex is made up of several multiplex servers, or nodes. Each node is assigned a role: coordinator (one per multiplex), writer, or reader. Readers and writers can serve as secondary nodes, backing up the coordinator node in case of

failure. The multiplex feature allows concurrent data loads and queries via independent data processing nodes connected to a shared data source. Each multiplex server has its own catalog store and IQ temporary store; all the servers in the multiplex share a common IQ store. Contrast with simplex.

node – a topology object representing a server or other entity type, displayed in the form of an icon.

perspective – a named tab in Sybase Control Center that displays information related to a collection of managed resources (such as servers) and a set of views associated with those resources. The views in a perspective are chosen by users of the perspective. You can create as many perspectives as you need, and customize them to monitor and manage your resources. Perspectives allow you to group resources in ways that make sense in your environment—for example by location, department, or project.

repository – a database in Sybase Control Center that stores information related to managed resources, along with user preference data, operational data, and performance statistics.

resource – a unique Sybase product component (such as a server) or a subcomponent.

SCC-enabled login account – a user account that has been granted privileges in Sybase Control Center by mapping appropriate Sybase Control Center roles. (Roles are typically mapped to a group to which the account belongs rather than to the account itself.) The user account and group can be native to Sybase Control Center or created in the operating system or the LDAP directory service to which Sybase Control Center authentication is delegated. You must use an SCC-enabled account to log in to Sybase Control Center.

SCC agent – a Sybase Control Center agent that runs on a managed server and enables Sybase Control Center to manage it. The SCC agent is installed automatically as part of the Sybase server.

schedule – the definition of a task (such as the collection of a set of statistics) and the time interval at which Sybase Control Center executes the task.

screen refresh interval – the period in seconds between refreshes of screens in the monitor views views (IQ Node Level Monitor and IQ Multiplex Level Monitor). Refreshing a screen redraws it with the most recent available data. Set the screen refresh interval on the Settings screen of either monitor view. See also collection repeat interval.

shared-disk mode – a feature that enables multiple instances of Sybase Control Center to execute from a single installation on a shared disk. Instances can be SCC servers, agents, or a mixture of the two.

simplex – a Sybase IQ implementation consisting of a single server that is not part of a multiplex. Contrast with multiplex.

singleton installation – a Sybase Control Center installation that runs a single SCC agent or server. Contrast with instance; see also shared-disk mode.

store – a store is one or more dbspaces that store persistent or temporary data for a special purpose. See catalog store, main store, or temporary store.

table version – the unit of versioning is the table. Table-level versioning structures aggregate data for columns at the table level. With table-level versioning, Sybase IQ can control access to the data at the level where write operations occur, and where query results are focused.

temporary store – the Sybase IQ temporary store is the portion of each Sybase IQ database that stores temporary tables and temporary scratch space data structures. The Sybase IQ server uses temporary data structures to sort and process data. Data in these tables persists only as long as you are connected to the database.

text configuration object – associated with a text index to define the way the index breaks text into terms, or searchable strings. You can create your own text configuration object or use one of the defaults provided by Sybase. See also text index.

text index – a mechanism to speed up full-text searching of tables. You must create a text index for a table before full-text searches can be performed. Each text index requires a text configuration object. See also text configuration object.

topology – a graphical representation of how the servers in a multiplex environment are connected to each other. Found in the IQ Multiplex Level Monitor, it is a network diagram that provides a visual map of the availability of the Sybase IQ server environment.

transaction – a set of related SQL statements that are treated as a single unit of work. To ensure consistency, if all the statements in the set cannot be executed, the changes made by the query are rolled back. The tables queried during the transaction are locked until a transaction is completed.

transaction log – the Sybase IQ transaction log records changes to the database. The transaction log includes version information, free space, and other information you can use to recover from a system failure. By default, the transaction log is created in the same directory as the catalog store. The default name for this dbfile is *<dbname>.log*.

trend period – See chart trend period.

view – a window in a perspective that displays information about one or more managed resources. Some views also let you interact with managed resources or with Sybase Control Center itself. For example, the Perspective Resources view lists all the resources managed by the current perspective. Other views allow you to configure alerts, view the topology of a replication environment, and graph performance statistics.

Index

W

WARN logging level 198
Watcom SQL functions in Sybase IQ 294

Windows
    configuring authentication 43, 89
    starting, stopping Sybase Control Center or
        SCC agent 31, 72

Index