SYBASE®

An **SAP** Company

Sybase Control Center for Adaptive Server® Enterprise

3.2.5

Contents

Contents

About Sybase Control Center for Adaptive Server

Sybase® Control Center for Adaptive Server® is a Web-based tool for monitoring the status and availability of Adaptive Servers.

Sybase Control Center supports Adaptive Server Enterprise version 15.0.2 and later. It supports clustered configurations on Adaptive Server Cluster Edition version 15.0.3 and later.

The Sybase Control Center client/server architecture allows multiple clients to monitor and control all Adaptive Servers in an enterprise using one or more Sybase Control Center servers. Sybase Control Center for Adaptive Server provides availability monitoring, historical performance monitoring, and administration capabilities in a scalable Web application that is integrated with management modules for other Sybase products. It offers shared, consolidated management of heterogeneous resources from any location, alerts that provide state- and threshold-based notifications about availability and performance in real time, and intelligent tools for spotting performance and usage trends, all via a thin-client, rich Internet application (RIA) delivered through your Web browser.

Use Sybase Control Center for Adaptive Server to track a variety of performance metrics, gathering statistics that over time will give you powerful insight into patterns of use and the behavior of databases, devices, caches, and processes on your servers. You can display collected data as tables or graphs. By plotting results over any period of time you choose, from a minute to a year, you can both see the big picture and focus on the particulars. Detailed knowledge of how your servers have performed in the past helps you ensure that Adaptive Server meets your needs in the future.

New Features in Sybase Control Center for Adaptive Server Enterprise

Descriptions of new and enhanced features in Sybase Control Center for Adaptive Server Enterprise and links to associated topics.

Table 1. New and enhanced Sybase Control Center for Adaptive Server features

Feature	Topics
Adaptive Server Cluster Edition management – create and manage an ASE shared-disk cluster. Add, display and manage instances, logical clusters, Backupserver and XP Server for a cluster. Manage availability, load balancing and failover rules.	*Creating a Cluster* on page 227 *Manage Logical Clusters* on page 235 *Manage Cluster Instances* on page 242 *Manage Load Profiles* on page 247 *Manage Routes* on page 250

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
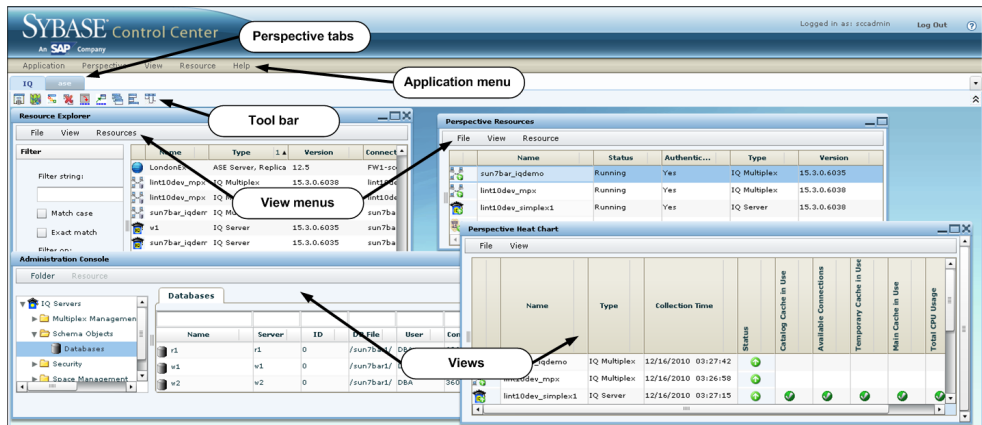
Toolbar Icons

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 Open 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.

Status Icons

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

Resource Status Icons in the Perspective Resources View and Heat Chart
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.

In the heat chart, hover the mouse over an icon in the Status column to display the status in English text.

Table 3. Resource Status Icons

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

KPI Status Icons in the Heat Chart
The heat chart uses KPI status icons to indicate the health of the KPIs it displays.

Hover the mouse over a KPI icon in any column to the right of the Status column to display the value of that KPI.

Table 4. KPI Status Icons

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

Display and Copy Options in Adaptive Server monitors

Options for collecting or displaying data on the user interface for Adaptive Server monitoring.

An **Options** drop-down menu enables you to effectively collect or display just the data that you need while monitoring the Adaptive Server. You can:

- **Choose columns** – Choose just the columns that you want displayed. By default all the columns are selected.
- **Copy selected row** – Cut data from a selected row and paste it into another application. The data is formatted like a row, and cells are separated by spaces.
- **Copy table** – Cut data from an entire table and paste it into another application. The data is formatted like a table, with rows and columns separated by spaces.

While monitoring Adaptive Server cluster configurations, you can also:

- **Expand all nodes** – Display table information for every instance of the cluster. With one click, the entire cluster information is displayed.
- **Collapse all nodes** – Hide instance-level information in a table, and only display cluster-level information.

On some windows of the Adaptive Server monitor, controls **Save All** and **Reset All** allow you to save several configured values, or reset all of them, with one click.

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.

Sorting Option	Description
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

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.

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
`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.

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.

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*.

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 Adaptive Server for Monitoring*

   On each server you plan to monitor, grant mon_role to the user account used to log in to the Adaptive Server and set monitoring options in the configuration file.
5. *Registering an Adaptive Server*

   Register a resource (for example, a server that can be monitored) to make Sybase Control Center aware of it and its connection information.
6. *Authenticating a Login Account for a Managed Resource*

   Specify the login account and password 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. *Displaying the Performance Overview*

The Overview screen shows Adaptive Server performance status.

See also

- *Get Started in a Production Environment* on page 20

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 Adaptive Server for Monitoring

On each server you plan to monitor, grant mon_role to the user account used to log in to the Adaptive Server and set monitoring options in the configuration file.

The Adaptive Server component of Sybase Control Center needs a user account to log in to Adaptive Server. To gather monitoring data, that account needs the role mon_role.

You can enable monitoring options using the **sp_configure** stored procedure or by editing the configuration file.

1. Create or select a login account for Sybase Control Center to use when it connects to Adaptive Server.

2. Use the **sp_role** stored procedure to grant **mon_role** to the login account, which in this example is called scc:

```
sp_role "grant", mon_role, scc
```

3. To enable monitoring, the values of monitor specific parameters must be set. In SCC 3.2.5 and higher, you are prompted with a Configuration Parameters Validation screen that displays the current values and the values needed to enable monitoring. Change the current values to the required value or, in the case of non-binary values, use values that are consistent with the the level of activity on the monitored Adaptive Server. You must have **sa_role** to change parameter values. In earlier versions, use one of these methods to set monitoring configuration options.

- Option 1: Use **sp_configure** to set the monitoring options to the values shown in the example below. (For information on using **sp_configure**, see the chapter on setting configuration parameters in the Adaptive Server *System Administration Guide*, Volume 1.)
- Option 2: Edit the Adaptive Server configuration file manually:

a) Shut down Adaptive Server.

b) In a text editor, open the server's configuration file, found at:

Windows: `%SYBASE%\<Adaptive-Server-name>.cfg`

UNIX: `$SYBASE/<Adaptive-Server-name>.cfg`

c) Save a backup copy of the configuration file.

d) Search for the Monitoring section of the file.

e) Set the monitoring options to the values shown in the example below.

f) Save the file and exit.

g) Start Adaptive Server.

Example

Note: After adding **mon_role**, if a resource is open for monitoring before you make the suggested parameter changes, a configuration dialog will open allowing you to modify the parameter values.

This example shows the monitoring section of the configuration file. Set these options either using **sp_configure**, or by manually editing the file. The Adaptive Server Monitor uses all these parameters, and notifies you if any of these options is not enabled. You may have to increase the values of **sql text pipe max messages** and **errorlog pipe max messages** depending on the level of activity on the monitored Adaptive Server.

```
[Monitoring]
        enable monitoring = 1
        sql text pipe active = 1
        sql text pipe max messages = 2000
        plan text pipe active = DEFAULT
        plan text pipe max messages = DEFAULT
        statement pipe active = 1
        statement pipe max messages = 2000
        errorlog pipe active = DEFAULT
        errorlog pipe max messages = DEFAULT
        deadlock pipe active = 1
        deadlock pipe max messages = 200
        wait event timing = 1
        process wait events = 1
        object lockwait timing = 1
        SQL batch capture = 1
        statement statistics active = 1
        per object statistics active = 1
        max SQL text monitored = 4096
        performance monitoring option = DEFAULT
```

Next
Register your Adaptive Server with Sybase Control Center and add it to a perspective—see Registering a Resource. Then continue with the Adaptive Server set-up tasks in this section.

Registering an Adaptive Server

Register a resource (for example, a server that can be monitored) to make Sybase Control Center aware of it and its connection information.

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, using uppercase and lowercase letters. If the name registered in Sybase Control Center does not exactly match the server name, some monitoring functions, including the topology view, do not work.
Resource Type	Select a resource type: • ASE Server (15.0.2.0) – monitor Adaptive Server 15.0.2.0 or later. Choose this type for full Adaptive Server monitoring capabilities. • ASE Server, Replication Only (12.5.0.0) – monitor only the RepAgent threads for an Adaptive Server that is older than version 15.0.2.0. Choose this type for an Adaptive Server that is part of a replication environment.
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
Server Host Name/Host Name	Local host name
Port Number	Local host port number
Character Set	Character set configured on Adaptive Server
	Note: If the Adaptive Server is configured to use a language that requires a multibyte character set such as Chinese, make sure to specify the correct character set in the connection profile.
Language	Language configured on Adaptive Server

5. Click **Next**.

6. (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).

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.)

9. Click **Finish**.

Authenticating a Login Account for a Managed Resource

Specify the login account and password 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 resource 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 information about the status represented by an icon in the chart, hover the mouse over the icon.

   • Status column – icon tooltips describe the status of the resource (Running or Stopped, for example).

- All columns to the right of Status – icon tooltips give the value of the KPI listed at the top of the column.

3. (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.

4. (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).

5. (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.

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

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

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

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

10. (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.

11. (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.

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

Displaying the Performance Overview

The Overview screen shows Adaptive Server performance status.

Check the Overview window to find out whether the server is running, and details about memory usage, CPU utilization, recent alerts, and so on. Other windows in the Adaptive Server monitor display more detailed information about the status of individual server resources such as engines, databases, caches, and processes. In Adaptive Server cluster configurations, this window allows you to check whether a particular cluster is running, how many instances of the cluster are down, and so on.

Note: The **Overview** screen is called **Cluster Overview** in Adaptive Server cluster configurations.

1. In the Perspective Resources view, select the server, click the drop-down arrow, and select **Monitor**. Alternately, in the Administration Console, select the server, click the drop-down arrow, and select **Monitor**.

The Adaptive Server monitor opens and displays the Overview screen. Check the server information in the upper left corner of the screen for the server's name, software product and version, its hardware platform, and an indication of whether the server is running. For cluster configurations, you also see the status of instances of each cluster, and the number of blocked process.

Note: When Sybase Control Center shows a server status of "Stopped," it means that the server is unreachable over the network.

2. (Optional) If data collections are running, mouse over the **Engine CPU Utilization** graph to display precise figures (values, times, and dates) for points on the curve.

   The graph shows the aggregate CPU utilization for all engines on the server. For cluster configurations, the graph shows the aggregate CPU utilization for each instance of the cluster.

3. (Optional – not in Adaptive Server cluster configurations) Move your mouse over the **Device IO/Sec** graph to display precise figures (values, times, and dates) for points on the curve.

   The graph shows device I/O aggregated across all devices on the server.

4. (Optional – not in Adaptive Server cluster configurations) Look at the Processes chart (far right) to see the number of configured and currently running processes and the highest number of concurrent processes since the server started, as well as the number of blocked processes.

5. (Optional – not in Adaptive Server cluster configurations) Look at the Memory chart to see statistics on caches and on physical, logical, and unused memory.

6. (Optional) Click a tab to see information about the resource you want to monitor:

   • **Details** - displays the version, edition, platform, number of deadlocks, platform, page size, device size and counters cleared for the Adaptive Server.

   • **Configured Resources** - displays, in tabular form, the configurable resources for each server or cluster instance. Each configuration option is displayed along with its currently configured value, run value that is currently used by the server, percentage of the resource that is currently in use, and the high water mark, which is the maximum amount of resource that has been used since the server was booted. Any column can be used to sort the table.

   The configured value for a resource is an editable field, denoted as such by a "pencil" **Edit** icon. Input a new numerical value for one or more resources, then choose either:

   • **Save All** to update the server with the new values. Sybase Control Center displays the new values. If the Adaptive Server encounters an error while applying the new value for a resource, Sybase Control Center displays the error below the table, and also next to the changed field in the row that causes the server error.

   • **Reset All** to restore the original value for the resource.

Note: Resources for each server or cluster instance may also be configured on the Server Configurations window.

- **Wait Events** - displays a list of server-wide wait events that can be very useful in performance tuning. Information about the wait-events includes the number of waits, wait time, average wait time, and wait description. For clustered servers, this information is displayed per instance.
- **Licenses** - displays a list of licenses that are currently checked out by the server or cluster instances. There is also information about the number, type, status (expirable, permanent and so on), and expiration date of each license.
- **Alerts** - displays a list of all fired alerts configured at server, cluster or cluster instance levels. For each alert, there is information about the time at which the alert was fired, severity, current statistic and threshold.

7. (Optional - Adaptive Server Cluster configurations) Mouse over the **Cluster Instances** graph to display precise information for points on the bar graph.

8. (Optional - Adaptive Server Cluster configurations) If data collections are running, mouse over the **Logical Cluster** graph to display precise information for points on the bar graph.

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 login accounts, statistics collection, alerts, and other Adaptive Server monitoring options.

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
```

On UNIX systems, omit the **-service** option.

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 28
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71

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>`.

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]
[-host host-name]
[-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. Sybase does not recommend using **-force** to remove or refresh a running instance in a Windows environment.
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-host** *host-name* – specify the host for this instance. Use with **-create**; required only when the instance name does not match the name of the host on which this instance will run. (The instance name defaults to the name of the current host unless you use **-instance** to specify another name.)

- **-instance [*instance-name*]** – 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.

  sccinstance assumes that the host name is the same as the instance name unless you use **-host** to specify a different host name.
- **-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
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

- **-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 on the current host, 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
```

- **Deploy two instances on the same host** – creates two agent instances on the host fireball. The first command does not need the **-host** option because the instance name is the same as the host name.

```
sccinstance -create -agent -instance fireball -portconfig rmi=9991
sccinstance -create -agent -instance fireball2 -host fireball
-portconfig rmi=9992
```

Note: In a production environment, Sybase recommends that you deploy no more than one SCC instance of each type (one server and one agent) on the same host.

- **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 3.2.5"

The Sybase Control Center 3.2.5 service is starting......
The Sybase Control Center 3.2.5 service was started
successfully.
```

       For an instance, include the instance name in parentheses:

```
net start "sybase control center 3.2.5 (Boston-1)"

The Sybase Control Center 3.2.5 (Boston-1) service is
starting......
The Sybase Control Center 3.2.5 (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 3.2.5"

The Sybase Control Center 3.2.5 service is stopping.....
The Sybase Control Center 3.2.5 service was stopped
successfully.
```

       For an instance, include the instance name in parentheses:

```
net stop "sybase control center 3.2.5 (Boston-1)"

The Sybase Control Center 3.2.5 (Boston-1) service is
stopping.....
The Sybase Control Center 3.2.5 (Boston-1) service was
stopped successfully.
```

See also

- *Deploying an Instance from a Shared Disk Installation* on page 21

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 commands 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 commands 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 34 or step *5.d* on page 34.
  - 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:MaxPermSize` – 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 37

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 36
- *Changing a Memory Option for an SCC UNIX Service* on page 37

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 36
- *Changing a Memory Option for an SCC UNIX Service* on page 37

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 36
- *Changing a Memory Option for an SCC Windows Service* on page 36

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 34

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 55

<u>Security</u>

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.

See also
- *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 86

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
* *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 86
* *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 51

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 45

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 43

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 42

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	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

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 38

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 111

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 57
• *Adding a Login Account to a Group* on page 58
• *Logins, Roles, and Groups* on page 59

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 57
- *Adding a Login Account to a Group* on page 58
- *Logins, Roles, and Groups* on page 59

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 57
- *Adding a Group* on page 57
- *Logins, Roles, and Groups* on page 59

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 Control 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 Replication Server Data Assurance Option.

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
   ```

   On UNIX systems, omit the **-service** option.

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 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *Instances* on page 169

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 62
*   *sccinstance Command* on page 63
*   *Instances* on page 169

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 61
*   *sccinstance Command* on page 63
*   *Instances* on page 169

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]
[-host host-name]
[-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. Sybase does not recommend using **-force** to remove or refresh a running instance in a Windows environment.
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-host** *host-name* – specify the host for this instance. Use with **-create**; required only when the instance name does not match the name of the host on which this instance will run. (The

instance name defaults to the name of the current host unless you use **-instance** to specify another name.)

- **-instance [***instance-name***]** – 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.

  sccinstance assumes that the host name is the same as the instance name unless you use **-host** to specify a different host name.

- **-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

Port Name	Description	Service Names	Property Names	Default Port
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 on the current host, 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
```

- **Deploy two instances on the same host** – creates two agent instances on the host fireball. The first command does not need the **-host** option because the instance name is the same as the host name.

```
sccinstance -create -agent -instance fireball -portconfig rmi=9991
sccinstance -create -agent -instance fireball2 -host fireball
-portconfig rmi=9992
```

Note: In a production environment, Sybase recommends that you deploy no more than one SCC instance of each type (one server and one agent) on the same host.

- **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 61
- *Shared-Disk Mode* on page 62

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 189

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 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *Configuring Memory Usage* on page 75
- *scc Command* on page 78

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 3.2.5"

The Sybase Control Center 3.2.5 service is starting......
The Sybase Control Center 3.2.5 service was started
successfully.
```

  For an instance, include the instance name in parentheses:
```
net start "sybase control center 3.2.5 (Boston-1)"

The Sybase Control Center 3.2.5 (Boston-1) service is
starting......
The Sybase Control Center 3.2.5 (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 3.2.5"

The Sybase Control Center 3.2.5 service is stopping.....
```

```
The Sybase Control Center 3.2.5 service was stopped
successfully.
```

For an instance, include the instance name in parentheses:

```
net stop "sybase control center 3.2.5 (Boston-1)"

The Sybase Control Center 3.2.5 (Boston-1) service is
stopping.....
The Sybase Control Center 3.2.5 (Boston-1) service was
stopped successfully.
```

See also
- *Registering the ODBC Driver in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *Configuring Memory Usage* on page 75
- *scc Command* on page 78

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 commands 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 commands 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 74 or step *5.d* on page 74.

- 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 68
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Configuring Memory Usage* on page 75
- *scc Command* on page 78

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.)

Modify this value	When	Guidelines
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

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 77
- *Changing a Memory Option for an SCC UNIX Service* on page 77
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *scc Command* on page 78

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 76
- *Changing a Memory Option for an SCC UNIX Service* on page 77
- *Starting and Stopping Sybase Control Center in Windows* on page 68

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 76
- *Changing a Memory Option for an SCC Windows Service* on page 77
- *Starting and Stopping Sybase Control Center in UNIX* on page 71

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 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
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
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 68
*   *Starting and Stopping Sybase Control Center in Windows* on page 68
*   *Starting and Stopping Sybase Control Center in UNIX* on page 71
*   *Configuring Memory Usage* on page 75
*   *Configuring Ports* on page 98
*   *Logging or Message Levels* on page 187

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.

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.

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 86
* *Configuring an LDAP Authentication Module* on page 87
* *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 96

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 87
- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 96
- *Adding a Login Account to the System* on page 106

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
- *Configuring Authentication for Windows* on page 40
- *Configuring an LDAP Authentication Module* on page 42
- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 96
- *Adding a Login Account to the System* on page 106

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 96

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 90
• *Encrypting a Password* on page 97

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 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 89

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 87
- *Configuring Authentication for Windows* on page 85
- *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 86
- *Assigning a Role to a Login or a Group* on page 102
- *User Authorization* on page 102

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 89

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

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
- *Alert-Triggered Scripts* on page 143
- *Alerts* on page 155
- *Adaptive Server Data Collections* on page 122
- *Adaptive Server Alerts* on page 137
- *Substitution Parameters for Scripts* on page 144
- *Key Performance Indicators for Adaptive Server* on page 125
- *Launching Sybase Control Center* on page 67
- *Logging in to Sybase Control Center* on page 82

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 67
- *Logging in to Sybase Control Center* on page 82

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 119

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

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 102

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 104
- *Adding a Login Account to a Group* on page 105
- *Removing a Login Account from a Group* on page 105

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 104
- *Adding a Login Account to a Group* on page 105
- *Removing a Login Account from a Group* on page 105

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 104
- *Removing a Group* on page 104
- *Removing a Login Account from a Group* on page 105

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 104
- *Removing a Group* on page 104
- *Adding a Login Account to a Group* on page 105

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
- *Configuring Authentication for Windows* on page 85
- *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 86
- *Configuring an LDAP Authentication Module* on page 87
- *Removing a Login Account from the System* on page 107
- *Modifying a User Profile* on page 108

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

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 Control 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.

See also
- *Role Assignment in Sybase Control Center for Adaptive Server* on page 119
- *Assigning a Role to a Login or a Group* on page 102

Get Started

Configure

Configure login accounts, statistics collection, alerts, and other Adaptive Server monitoring options.

1. *Configuring Adaptive Server for Monitoring*

   On each server you plan to monitor, grant mon_role to the user account used to log in to the Adaptive Server and set monitoring options in the configuration file.

2. *Registering an Adaptive Server*

   Register a resource (for example, a server that can be monitored) to make Sybase Control Center aware of it and its connection information.

3. *Importing Resources for Batch Registration*

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

4. *Registering the Agent for an Adaptive Server*

   Use the Adaptive Server Administration Console to register the agent.

5. *Creating a Perspective*

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

6. *Adding a Resource to a Perspective*

   Add one or more resources to the current perspective.

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

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

8. *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.

9. *Setting Display Options for Adaptive Server Performance Data*

   Change the screen refresh interval, chart trend period, alert list size, historical SQLs size, and historical SQLs trend period for an Adaptive Server.

10. *Setting Adaptive Server Parameters in the Configuration File*

    Set options that control the behavior of Sybase Control Center for Adaptive Server.

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 56
*   *Logins, Roles, and Groups* on page 109
*   *Setting Up Security* on page 83
*   *Assigning a Role to a Login or a Group* on page 102

Configuring Adaptive Server for Monitoring

On each server you plan to monitor, grant mon_role to the user account used to log in to the Adaptive Server and set monitoring options in the configuration file.

The Adaptive Server component of Sybase Control Center needs a user account to log in to Adaptive Server. To gather monitoring data, that account needs the role mon_role.

You can enable monitoring options using the **sp_configure** stored procedure or by editing the configuration file.

1. Create or select a login account for Sybase Control Center to use when it connects to Adaptive Server.
2. Use the **sp_role** stored procedure to grant **mon_role** to the login account, which in this example is called scc:

    ```
    sp_role "grant", mon_role, scc
    ```

3. To enable monitoring, the values of monitor specific parameters must be set. In SCC 3.2.5 and higher, you are prompted with a Configuration Parameters Validation screen that displays the current values and the values needed to enable monitoring. Change the current values to the required value or, in the case of non-binary values, use values that are consistent with the the level of activity on the monitored Adaptive Server. You must have **sa_role** to change parameter values. In earlier versions, use one of these methods to set monitoring configuration options.

    *   Option 1: Use **sp_configure** to set the monitoring options to the values shown in the example below. (For information on using **sp_configure**, see the chapter on setting configuration parameters in the Adaptive Server *System Administration Guide*, Volume 1.)
    *   Option 2: Edit the Adaptive Server configuration file manually:

    a) Shut down Adaptive Server.
    b) In a text editor, open the server's configuration file, found at:

        Windows: `%SYBASE%\<Adaptive-Server-name>.cfg`

        UNIX: `$SYBASE/<Adaptive-Server-name>.cfg`
    c) Save a backup copy of the configuration file.

d) Search for the Monitoring section of the file.
e) Set the monitoring options to the values shown in the example below.
f) Save the file and exit.
g) Start Adaptive Server.

Example

Note: After adding **mon_role**, if a resource is open for monitoring before you make the suggested parameter changes, a configuration dialog will open allowing you to modify the parameter values.

This example shows the monitoring section of the configuration file. Set these options either using **sp_configure**, or by manually editing the file. The Adaptive Server Monitor uses all these parameters, and notifies you if any of these options is not enabled. You may have to increase the values of **sql text pipe max messages** and **errorlog pipe max messages** depending on the level of activity on the monitored Adaptive Server.

```
[Monitoring]
        enable monitoring = 1
        sql text pipe active = 1
        sql text pipe max messages = 2000
        plan text pipe active = DEFAULT
        plan text pipe max messages = DEFAULT
        statement pipe active = 1
        statement pipe max messages = 2000
        errorlog pipe active = DEFAULT
        errorlog pipe max messages = DEFAULT
        deadlock pipe active = 1
        deadlock pipe max messages = 200
        wait event timing = 1
        process wait events = 1
        object lockwait timing = 1
        SQL batch capture = 1
        statement statistics active = 1
        per object statistics active = 1
        max SQL text monitored = 4096
        performance monitoring option = DEFAULT
```

Next

Register your Adaptive Server with Sybase Control Center and add it to a perspective—see Registering a Resource. Then continue with the Adaptive Server set-up tasks in this section.

See also

Registering an Adaptive Server

Register a resource (for example, a server that can be monitored) to make Sybase Control Center aware of it and its connection information.

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, using uppercase and lowercase letters. If the name registered in Sybase Control Center does not exactly match the server name, some monitoring functions, including the topology view, do not work.
Resource Type	Select a resource type: • ASE Server (15.0.2.0) – monitor Adaptive Server 15.0.2.0 or later. Choose this type for full Adaptive Server monitoring capabilities. • ASE Server, Replication Only (12.5.0.0) – monitor only the RepAgent threads for an Adaptive Server that is older than version 15.0.2.0. Choose this type for an Adaptive Server that is part of a replication environment.
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
Server Host Name/Host Name	Local host name
Port Number	Local host port number

Field	Description
Character Set	Character set configured on Adaptive Server
	Note: If the Adaptive Server is configured to use a language that requires a multibyte character set such as Chinese, make sure to specify the correct character set in the connection profile.
Language	Language configured on Adaptive Server

5. Click **Next**.

6. (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).

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.)

9. Click **Finish**.

See also
* *Common Display Options* on page 5
* *Resources* on page 163
* *Unregistering a Resource* on page 163

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 163
*   *Unregistering a Resource* on page 163

Registering the Agent for an Adaptive Server

Use the Adaptive Server Administration Console to register the agent.

Note: For Adaptive Server versions 15.7 and earlier, the agent is the Unified Agent. For versions 15.7 ESD1 and higher, the agent is the SCC agent. Follow the steps provided here to register either type of agent.

The agent is installed and set up as a component of the Adaptive Server installation. For information, see the Adaptive Server installation guide for your platform.

You must register and authenticate the agent to use Sybase Control Center to perform certain administrative tasks such as starting the Adaptive Server or viewing the Adaptive Server error log.

You must register an agent for each Adaptive Server you have configured. The agent is configured on the same host as the Adaptive Server that it manages. When you register the agent, you are updating Sybase Control Center with information on the machine and port number on which the agent is configured.

For information about configuration options for the Unified Agent, see *Unified Agent and Agent Management Console Users Guide > Installing and Configuring Unified Agent and Agent Management Console* in the Adaptive Server documentation set.

For information on security features for the Unified Agent, see *Unified Agent and Agent Management Console Users Guide > Security* in the Adaptive Server documentation set.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.
2. Click **ASE Servers**.
    You see a list of monitored servers.
3. Select the server you want to manage and click the drop-down arrow that appears next to the name.
4. Click **Register Agent**.

You see the Server Properties screen.

5. Enter the port number for the agent and click **Register**.

> **Note:** After the agent is registered, you can authenticate the agent, or clear the registration.

6. (Optional) Enter the login name and password for the agent, and click **Authenticate**.

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.

Prerequisites
Register the resources.

Task
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

Authenticating a Login Account for a Managed Resource

Specify the login account and password 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 resource is a Replication Server, also enter the RSSD user name and password.
7. Click **OK** to save and exit the dialog.

See also
- *Setting Up Statistics Collection* on page 120
- *Configuring Adaptive Server for Monitoring* on page 112
- *Setting Adaptive Server Parameters in the Configuration File* on page 131
- *User Authorization* on page 102

Role Assignment in Sybase Control Center for Adaptive Server

Sybase Control Center automatically grants administrative or monitoring privileges for Adaptive Server.

Sybase Control Center automatically assigns monitoring privileges to users who have **mon_role** privileges on an Adaptive Server, which lets them perform monitoring tasks for that server.

Sybase Control Center automatically assigns administration privileges to users who have **sa_role** privileges on an Adaptive Server, which lets them perform certain administrative tasks for that server.

SCC checks for role validation every 30 minutes. The **Monitor** option in the Resource menu of the Perspective Resources view is greyed out until the resource is authenticated. If you

authenticate a resource without having **mon_role** privileges, SCC warns you that you cannot perform monitoring tasks or configure alerts on the server.

Note: If a role is revoked on the Adaptive Server from outside of Sybase Control Center, SCC does not register the change till the next role-check occurs. However, as the monitoring or administrative role has been revoked on the Adaptive Server, the user cannot execute such a task through Sybase Control Center. If a role is revoked on Adaptive Server from within SCC, SCC registers the change immediately.

See also

Encrypted Authentication for Adaptive Server

Sybase Control Center uses encrypted passwords to connect to Adaptive Servers that are configured for network password encryption.

If an Adaptive Server is configured to use network password encryption by setting **net password encryption reqd**, Sybase Control Center establishes a connection to an Adaptive Server using a password that is encrypted during network transmission.

The Adaptive Server must be configured to use network password encryption to transmit encrypted passwords. See *System Administration Guide: Volume 1 > Setting Configuration Parameters*.

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
- *Authenticating a Login Account for a Managed Resource* on page 119

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.)

Adaptive Server Data Collections

Collection of Adaptive Server data may be scheduled through either a default data collection or preconfigured statistics collections.

When an Adaptive Server is first authenticated, Sybase Control Center schedules a default data collection called **collection_ase_availability**.

The user who first authenticates and monitors an Adaptive Server resource owns its default collection. You can begin to schedule a default collection in one of these ways:

- Authenticate an Adaptive Server after registering it. This is the default method of scheduling a default collection.
- Create a scheduled job that is initiated when the Sybase Control Center server starts.

You can set up jobs in the scheduler to collect preconfigured statistics collections using the collections **collection_ase_all_client_kpis**, **collection_ase_histmon**, or **collection_ase_rat**.

Listed below are the Adaptive Server data collections and corresponding key performance indicators (KPIs).

collection_ase_availability
The default collection of data that is set up when an Adaptive Server is first authenticated. The data in the default collection is gathered in 60 second intervals.

Note: The default collection contains the same key performance indicators (KPIs) as the Perspective Heatchart, and therefore these KPIs need not be scheduled in additional collections.

The KPIs used in the default collection include:

Server Percent CPU Utilization	Number of Suspended Processes
Number of Blocked Processes	Server Availability State

collection_ase_all_client_kpis
Collects the data for historical charts in the Adaptive Server component, including those on the Overview, Devices, Engines, and Segments screens in the Adaptive Server monitor. Schedule this collection to see real time charting for these resources.

The KPIs used in **collection_ase_all_client_kpis** include:

Active Connections in Cluster Instance	Device IO Response Time
Device IO Rate in Cluster Instance	Engine CPU Utilization
Engine CPU Utilization in Cluster Instance	Segment Space Usage
CPU Busy Value for Logical Cluster Workload	Server CPU Utilization
IO Load Value for Logical Cluster Workload	Server Device IO Rate
Load Score for Logical Cluster Workload	Cluster Temp DB IO Rate
Run Queue Length for Logical Cluster Workload	Thread User CPU Utilization
Device APF Reads	Thread System CPU Utilization

Device IO Rate	Thread Total CPU Utilization

collection_ase_histmon

Collects historical statistics for an Adaptive Server. These statistics are not displayed in the Adaptive Server monitor. To view these statistics, launch the 'Statistics Chart' window from the context menu. Schedule this collection to activate alerts and view them in a statistics chart.

The KPIs used in **collection_ase_histmon** include:

Workload Load Score in Cluster Instance	Number of Bytes Received in Network IO
Number of Bytes Received in Cluster Instance	Number of Bytes Sent in Network IO
Number of Bytes Sent in Cluster Instance	Number of Deadlocks
Number of CIPC Messages Received in Cluster Instance	Number of Locks
Number of CIPC Messages Sent in Cluster Instance	Number of Packets Received in Network IO
Number of Committed Transactions in Cluster Instance	Number of Packets Sent in Network IO
Number of Packets Received in Cluster Instance	Number of Page Locks
Number of Packets Sent in Cluster Instance	Number of Row Locks
Cache Hit Ratio	Number of Table Locks
Number of Cache Misses	Number of Transactions
Number of Cache Searches	Number of User Connections
Device Free Space	Procedure Cache Hit Ratio
Device Space Usage	Server Device IO Rate
Active Connections in Logical Cluster	sp_who Response Time
Number of Failover Instances in Logical Cluster	Statement Cache Hit Ratio
Segment Free Space	Server tempdb Free Space
Average Blocked Process Wait Time	Server tempdb Space Used
Number of Address Locks	

collection_ase_rat

Collects RepAgent Threads metrics for the charts on the Adaptive Server component's Replication Agent screen. Include this collection only if you are planning to monitor replication from this primary database.

See also

Key Performance Indicators for Adaptive Server

Lists and describes the key performance indicators (KPIs) that provide the statistics displayed on Adaptive Server screens and charts in Sybase Control Center.

Adaptive Server Object	KPI Name	Description	Data Collection Name
Cluster Instances	Active Connections in Cluster Instance	Number of active connections to a cluster instance. This KPI is collected separately for each cluster instance.	collection_ase_all_client_kpis
	Workload Load Score in Cluster Instance	Load score in each cluster instance. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Bytes Received in Cluster Instance	Number of bytes received during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Bytes Sent in Cluster Instance	Number of bytes sent during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of CIPC Messages Received in Cluster Instance	Number of CIPC messages received during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon

Configure

Adaptive Server Object	KPI Name	Description	Data Collection Name
	Number of CIPC Messages Sent in Cluster Instance	Number of CIPC messages sent during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Committed Transactions in Cluster Instance	Number of committed transactions in the cluster instance during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Packets Received in Cluster Instance	Number of packets received during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Packets Sent in Cluster Instance	Number of packets sent during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Device IO Rate in Cluster Instance	This KPI is collected separately for each cluster instance and is used to generate the Instance connection chart on the Cluster Instance screen.	collection_ase_all_client_kpis
	Engine CPU Utilization in Cluster Instance	Engine CPU Utilization in a cluster instance. This KPI is collected separately for each cluster instance	collection_ase_all_client_kpis
Cluster Workloads	CPU Busy Value for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the CPU busy chart on the Workload screen.	collection_ase_all_client_kpis
	IO Load Value for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the IO load chart on the Workload screen.	collection_ase_all_client_kpis

Adaptive Server Object	KPI Name	Description	Data Collection Name
	Load Score for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the load score chart on the Workload screen.	collection_ase_all_client_kpis
	Run Queue Length for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the run queue length chart on the Workload screen.	collection_ase_all_client_kpis
Data Caches	Cache Hit Ratio	Hit ratio in the data cache during the current collection cycle.	collection_ase_histmon
	Number of Cache Misses	Number of cache misses (or reads into the cache from the disk) during the current collection cycle.	collection_ase_histmon
	Number of Cache Searches	Number of cache searches during the current collection cycle.	collection_ase_histmon
Devices	Device APF Reads	Rate per second of asynchronous pre-fetch read operations on this device.	collection_ase_all_client_kpis
	Device Free Space	Total amount of free space, in megabytes, on this device.	collection_ase_histmon
	Device Space Usage	Total amount of space on this device, in megabytes, that is used by processes.	collection_ase_histmon
	Device IO Rate	Rate of I/O operations per second on this device.	collection_ase_all_client_kpis
	Device IO Response Time	Response time, in milliseconds, for I/O operations performed on this device.	collection_ase_all_client_kpis
Engines	Engine CPU Utilization	Percentage of CPU cycles used by this Adaptive Server engine.	collection_ase_all_client_kpis
Logical Clusters	Active Connections in Logical Cluster	Number of active connections using the Logical Cluster at the time of collection. This KPI is collected separately for each Logical Cluster.	collection_ase_histmon

Adaptive Server Object	KPI Name	Description	Data Collection Name
	Number of Failover Instances in Logical Cluster	Number of instance that are failed-over in the currently active Logical Cluster. This KPI is collected separately for each logical cluster.	collection_ase_histmon
Segments	Segment Free Space	Amount, in megabytes, of free space in the segment. This KPI is collected separately for each segment.	collection_ase_histmon
	Segment Space Usage	Change, in megabytes, in the amount of space used by this segment since the last refresh.	collection_ase_all_client_kpis
Server	Average Blocked Process Wait Time	Average time, in milliseconds, that the current blocked processes have waited.	collection_ase_histmon
	Number of Address Locks	Number of address-level locks server-wide.	collection_ase_histmon
	Number of Blocked Processes	Number of currently blocked processes that have been blocked for more than 5 seconds. (The Heat Chart uses this metric to display server status.)	collection_ase_availability
	Number of Bytes Received in Network IO	Number of bytes received during the current collection cycle.	collection_ase_histmon
	Number of Bytes Sent in Network IO	Number of bytes sent during the current collection cycle.	collection_ase_histmon
	Number of Deadlocks	Number of deadlocks on the server since the most recent execution of the collection.	collection_ase_histmon
	Number of Locks	Total number of active locks of all types on the server.	collection_ase_histmon
	Number of Packets Received in Network IO	Number of packets received during the current collection cycle.	collection_ase_histmon
	Number of Packets Sent in Network IO	Number of packets sent during the current collection cycle.	collection_ase_histmon

Adap-tive Server Object	KPI Name	Description	Data Collection Name
	Number of Page Locks	Number of page-level locks server-wide.	collection_ase_histmon
	Number of Row Locks	Number of row-level locks server-wide.	collection_ase_histmon
	Number of Suspended Processes	Number of processes that are currently suspended. (The Heat Chart uses this metric to display server status.)	collection_ase_availability
	Number of Table Locks	Number of table-level locks server-wide.	collection_ase_histmon
	Number of Transac-tions	Total number of transactions during the current collection cycle. **Note:** This KPI is available only on Adaptive Server versions 15.0.3 Cluster Edition and 15.0.3 ESD #3, and later.	collection_ase_histmon
	Number of User Con-nections	Current number of user connections on the server.	collection_ase_histmon
	Procedure Cache Hit Ratio	Hit ratio in the procedure cache.	collection_ase_histmon
	Resource State	Status of the Adaptive Server. Values of most interest are STOPPED and RUN-NING.	collection_ase_availability
	Server CPU Utilization	Average CPU utilization percentage across all active Adaptive Server en-gines on the server.	collection_ase_availabili-ty, collection_ase_all_cli-ent_kpis
	Server Device IO Rate	Total number of I/O operations per-formed by all devices on the server during the current collection cycle.	collection_ase_histmon, collection_ase_all_cli-ent_kpis

Adaptive Server Object	KPI Name	Description	Data Collection Name
	sp_who Response Time	Time in milliseconds the sp_who stored procedure takes to return a response. sp_who is called each time collection_ase_histmon is executed to collect Adaptive Server performance statistics.	collection_ase_histmon
	Statement Cache Hit Ratio	Hit ratio in the statement cache during the current collection cycle.	collection_ase_histmon
	Server tempdb Free Space	Amount, in megabytes, of free space in the **tempdb** database.	collection_ase_histmon
	Server tempdb Space Used	Amount, in megabytes, of space used in the **tempdb** database.	collection_ase_histmon
TempDBs Activity	Cluster Temp DB IO Rate	The rate of **tempdb** IO activity, per second, in a cluster.	collection_ase_all_client_kpis
Threads	Thread User CPU Utilization	CPU utilization percentage in handling user committed queries for each thread.	collection_ase_all_client_kpis
	Thread System CPU Utilization	CPU utilization percentage in handling system level operations for each thread.	collection_ase_all_client_kpis
	Thread Total CPU Utilization	Total CPU utilization obtained by adding Thread User CPU Utilization and Thread System CPU Utilization.	collection_ase_all_client_kpis

See also
- *Adaptive Server Alerts* on page 137

Setting Display Options for Adaptive Server Performance Data

Change the screen refresh interval, chart trend period, alert list size, historical SQLs size, and historical SQLs trend period for an Adaptive Server.

Follow these steps to set the options on the Settings screen:

1. Select the server you want to configure in the Perspective Resources view, click the drop-down arrow, and select **Monitor**.
2. Select **Settings** from the left panel.
3. (Optional) Enter a new value in the **Screen Refresh Interval** field. Refreshing a screen redraws it with the most recent available data. The screen refresh interval is the period between refreshes, with a default of 30 seconds.
4. (Optional) Enter a new value in the **Chart Trend Period** field. The chart trend period is the amount of time covered by historical charts in the Adaptive Server component. The default is 15 minutes.
5. (Optional) Enter a new value in the **Alert list size** field. The Alerts table in the Overview window contains a list of all alerts configured for a server. The maximum number of rows in the Alerts table is controlled by the alert list size, with a default of 100.
6. (Optional) Enter a new value in the **Historical SQLs size** field. The active SQLs table in the SQL Activity window contains a list of active SQL statements. The maximum number of statements in this table is controlled by the historical SQLs size, with a default of 500.
7. (Optional) Enter a new value in the **Historical SQLs trend period** field. The list of active SQL statements in the active SQLs table is displayed for a maximum period of time controlled by the historical SQLs trend period. The default value for this setting is 5 minutes.
8. Click **Apply Settings**.

See also
* *Setting Up Statistics Collection* on page 120

Setting Adaptive Server Parameters in the Configuration File

Set options that control the behavior of Sybase Control Center for Adaptive Server.

The configuration file for the Adaptive Server component of Sybase Control Center resides here:

Windows: `%SYBASE%\SCC-3_2\plugins\ASEMap\config.properties`
UNIX: `$SYBASE/SCC-3_2/plugins/ASEMap/config.properties`

1. Open the `config.properties` file with a text editor.
2. Save a backup copy of the file.
3. (Optional) To change the number of times Sybase Control Center tries to reopen a broken JDBC connection to Adaptive Server, edit the value of **attempts_reopen_con**.
4. (Optional) To change the interval (in seconds) between successive attempts to reopen a broken JDBC connection, edit the value of **time_between_reattempts**.

5. (Optional) To change the interval (in hours) between refreshes of the list of monitored objects in the Sybase Control Center repository, edit the value of **revalidation_frequency**.

6. (Optional) To change the time interval (in seconds) for JDBC internal queries to execute before they time out, edit the value of **jdbc_internal_query_timeout**.

7. Save and exit the file.

8. To make the new settings take effect, restart Sybase Control Center.

Note: Sybase Control Center does not declare a server stopped (down) until all the retries called for by **attempts_reopen_con** have failed.

See also
- *Creating an Alert* on page 133

Configuration Parameters for Adaptive Server

Lists configuration parameters for Adaptive Server, including default values in Sybase Control Center.

Sybase Control Center provides these configuration parameters on Adaptive Server:

- **attempts_reopen_con** – Changes the number of times Sybase Control Center tries to reopen a broken JDBC connection to Adaptive Server. The default value is 5. This parameter only tries to re-establish broken connections—it does not control retries when a first attempt to connect to the server fails.
- **time_between_reattempts** – Changes the interval (in seconds) between successive attempts to reopen a broken JDBC connection. The default value is 6. This parameter only tries to re-establish broken connections—it does not control retries when a first attempt to connect to the server fails.
- **revalidation_frequency** – Changes the interval (in hours) between refreshes of the list of monitored objects in the Sybase Control Center repository.
  The repository stores a list of the caches, devices, engines, and segments for all monitored Adaptive Servers. Sybase Control Center does not collect statistics for objects unless they are listed in the repository. The revalidation_frequency parameter controls how often Sybase Control Center refreshes (revalidates) the list of objects associated with each monitored server. The default value is 24.
- **jdbc_internal_query_timeout** – Changes the time interval (in seconds) for JDBC internal queries to execute before they time out. The default value is 15. This prevents server connection attempts and other operations from hanging when resources on the server have been exhausted. You may want to modify this if your ASE server is up, but performance is so slow that the ASEMAP agent fails to successfully establish a connection to the server. This failure usually indicates a configuration problem on the ASE server and not in the SCC ASEMAP component. Sybase recommends that you contact Sybase Customer Support to determine the cause of connection problems before modifying this parameter.

- **query_timeout** – The default query timeout for any statements created on this connection. The default value is 30 seconds.
- **query_timeout_spaceused** – The query timeout for **sp_spaceused**, which is used to determine the space usage parameters for a database. The default value is 180 secs,

Note: If the parameter value is changed, the Sybase Control Center server must be shutdown and restarted.

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 16. 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 Adaptive Server Parameters in the Configuration File* on page 131
- *Optional Configuration Steps* on page 145
- *Alert-Triggered Scripts* on page 143
- *Adaptive Server Data Collections* on page 122
- *Adaptive Server Alerts* on page 137

- *Substitution Parameters for Scripts* on page 144
- *Key Performance Indicators for Adaptive Server* on page 125
- *Alert Types, Severities, and States for Adaptive Server* on page 142
- *Assigning a Role to a Login or a Group* on page 102
- *Configuring the E-mail Server* on page 100
- *Alerts* on page 155
- *Testing an Alert-Triggered Script* on page 158

Adaptive Server Alerts

Lists and describes alerts you can configure for Adaptive Server.

The alerts are based on the same key performance indicators (KPIs) that are collected for the performance and availability monitor displays and for the Statistics Chart.

All alerts are of type "Threshold", except for **Resource State** and **RepAgent Thread State Change** which are of type "State".

Note: When an alert is raised, the status bar displays an alert icon. To view the alert, click the **Alerts** tab on the Overview screen.

Adaptive Server Object	Alert	Description	Data Collection Name
Cluster Instances	Active Connections in Cluster Instance	Number of active connections to a cluster instance.	collection_ase_all_client_kpis
	Engine CPU Utilization in Cluster Instance	Percentage CPU utilization of an instance in a shared disk cluster.	collection_ase_all_client_kpis
	Number of Bytes Received in Cluster Instance	Number of bytes received during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Bytes Sent in Cluster Instance	Number of bytes sent during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of CIPC Messages Received in Cluster Instance	Number of messages received during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon

Adaptive Server Object	Alert	Description	Data Collection Name
	Number of CIPC Messages Sent in Cluster Instance	Number of messages sent during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Packets Received in Cluster Instance	Number of packets received during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Packets Sent in Cluster Instance	Number of packets sent during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Number of Committed Transactions in Cluster Instance	Number of committed transactions in the cluster instance during the current collection cycle. This KPI is collected separately for each cluster instance.	collection_ase_histmon
	Workload Load Score in Cluster Instance	Load score in each cluster instance. This KPI is collected separately for each cluster instance.	collection_ase_histmon
Cluster Workload	CPU Busy Value for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the CPU busy chart on the Workload screen.	collection_ase_all_client_kpis
	IO Load Value for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the IO load chart on the Workload screen.	collection_ase_all_client_kpis
	Load Score for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the load score chart on the Workload screen.	collection_ase_all_client_kpis

Adap-tive Server Object	Alert	Description	Data Collection Name
	Run Queue Length for Logical Cluster Work-load	This KPI is collected separately for each cluster workload and is used to generate the run queue length chart on the Workload screen.	collection_ase_all_cli-ent_kpis
Data Ca-ches	Cache Hit Ratio	Hit ratio in the data cache during the current collection cycle.	collection_ase_histmon
Devices	Device IO Rate	Rate of I/O operations per second on this device.	collection_ase_all_cli-ent_kpis
	Device IO Response Time	Response time, in milliseconds, for I/O operations performed on this device.	collection_ase_all_cli-ent_kpis
	Device Free Space	Total amount of free space, in mega-bytes, on this device.	collection_ase_histmon
Engines	Engine CPU Utilization	Percentage of CPU cycles used by this Adaptive Server engine.	collection_ase_all_cli-ent_kpis
Logical Clusters	Active Connections in Logical Cluster	Number of active connections using the Logical Cluster at the time of col-lection. This KPI is collected separate-ly for each Logical Cluster.	collection_ase_histmon
	Number of Failover In-stances in Logical Clus-ter	Number of instance that are failed-over in the currently active Logical Cluster. This KPI is collected sepa-rately for each logical cluster.	collection_ase_histmon
Replica-tion Agent	RepAgent Thread State Change	The alert to send when an Adaptive Server RepAgent Thread changes state.	collection_ase_rat
	Transaction Log Size	The size of an Adaptive Server Rep-Agent Thread's transaction log.	collection_ase_rat
	Number of Log Opera-tions Scanned per Sec-ond	The number of operations scanned by an Adaptive Server RepAgent thread.	collection_ase_rat

Adap-tive Server Object	Alert	Description	Data Collection Name
	Number of Log Opera-tions Processed per Second	The number of operations processed by an Adaptive Server RepAgent thread.	collection_ase_rat
Segments	Segment Free Space	Amount of free space in the segment. This KPI is collected separately for each segment.	collection_ase_histmon
Server	Average Blocked Proc-ess Wait Time	Average time, in milliseconds, that the current blocked processes have wai-ted.	collection_ase_histmon
	Number of Address Locks	Number of address-level locks server-wide.	collection_ase_histmon
	Number of Blocked Processes	Number of currently blocked process-es that have been blocked for more than 5 seconds. The Heat Chart uses this metric to display server status.	collection_ase_availability
	Number of Bytes Re-ceived in Network IO	Number of bytes received during the current collection cycle.	collection_ase_histmon
	Number of Bytes Sent in Network IO	Number of bytes sent during the cur-rent collection cycle.	collection_ase_histmon
	Number of Deadlocks	Number of deadlocks on the server since the most recent execution of the collection.	collection_ase_histmon
	Number of Locks	Total number of active locks of all types on the server.	collection_ase_histmon
	Number of Packets re-ceived in Network IO	Number of packets received during the current collection cycle.	collection_ase_histmon
	Number of Packets sent in Network IO	Number of packets sent during the current collection cycle.	collection_ase_histmon
	Number of Page Locks	Number of page-level locks server-wide.	collection_ase_histmon
	Number of Row Locks	Number of row-level locks server-wide.	collection_ase_histmon

Adaptive Server Object	Alert	Description	Data Collection Name
	Number of Suspended Processes	Number of processes that are currently suspended. The Heat Chart uses this metric to display server status.	collection_ase_availability
	Number of Table Locks	Number of table-level locks server-wide.	collection_ase_histmon
	Number of Transactions	Total number of transactions during the current collection cycle.	collection_ase_histmon
	Number of User Connections	Current number of user connections on the server.	collection_ase_histmon
	Procedure Cache Hit Ratio	Hit ratio in the procedure cache.	collection_ase_histmon
	Resource State	Status of the Adaptive Server. Values of most interest are STOPPED and RUNNING.	collection_ase_availability
	Server CPU Utilization	Average CPU utilization percentage across all active Adaptive Server engines on the server.	collection_ase_availability, collection_ase_all_client_kpis
	Server Device IO Rate	Total number of I/O operations performed by all devices on the server during the current collection cycle.	collection_ase_availability, collection_ase_all_client_kpis
	Server tempdb Free Space	Amount of free space in the **tempdb** database, in megabytes.	collection_ase_histmon
	Server tempdb Space Used	Amount of space used by the **tempdb** database, in megabytes.	collection_ase_histmon
	sp_who Response Time	Time in milliseconds the sp_who stored procedure takes to return a response. sp_who is called each time collection_ase_histmon is executed to collect Adaptive Server performance statistics.	collection_ase_histmon
	Statement Cache Hit Ratio	Hit ratio in the statement cache during the current collection cycle.	collection_ase_histmon

Adaptive Server Object	Alert	Description	Data Collection Name
Threads	Thread CPU Utilization	Total CPU utilization by a thread, including user and system CPU utilization.	collection_ase_all_client_kpis

See also
* *Key Performance Indicators for Adaptive Server* on page 125

Alert Types, Severities, and States for Adaptive Server

Learn about the properties that define and control alerts.

An alert's type determines what causes it to fire.

Table 17. 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 18. 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.

Table 19. Adaptive Server States

Alert/KPI	State	Description
Resource State	Running	Resource or component is operating normally. This state is associated with a severity of Normal.
	Stopped	The resource or component tracked by this metric is unreachable. This state is associated with a severity of Critical.

See also
* *Alert-Triggered Scripts* on page 143
* *Alerts* on page 155
* *Assigning a Role to a Login or a Group* on page 102
* *Creating an Alert* on page 133
* *Configuring the E-mail Server* on page 100

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
- *Testing an Alert-Triggered Script* on page 158
- *Alerts* on page 155

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 20. 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
%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 21. 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 22. 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 102
Authentication	Add authentication modules to allow Windows, UNIX, and LDAP users to log in to Sybase Control Center.	*Setting up Security* on page 83
Alerts	Modify alert thresholds and subscriptions and delete alerts.	*Alerts* on page 155
Data collection	Modify collection intervals and schedules, suspend and resume the schedule, and delete collection jobs.	*Job Scheduling* on page 150
Resources	Unregister resources, add them to perspectives, or remove them.	*Resources* on page 163
Perspectives	Create, remove, and rename perspectives.	*Perspectives* on page 165
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 169
Repository	Set purging options and schedule backups of the repository database.	*Repository* on page 178

See also
• *Creating an Alert* on page 133

Manage and Monitor

Manage and monitor the Adaptive Server.

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 information about the status represented by an icon in the chart, hover the mouse over the icon.
   - Status column – icon tooltips describe the status of the resource (Running or Stopped, for example).
   - All columns to the right of Status – icon tooltips give the value of the KPI listed at the top of the column.
3. (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.
4. (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).
5. (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.

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

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

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

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

10. (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.

11. (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.

12. (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

Manage Sybase Control Center

Manage Sybase Control Center for Adaptive Server using monitoring statistics and the Sybase Control Center log for Adaptive Server.

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
• *Setting Up Statistics Collection* on page 120
• *Adaptive Server Data Collections* on page 122

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 151
• *Resuming and Suspending a Data Collection Job* on page 152
• *Adding a New Schedule to a Job* on page 152
• *Modifying the Data Collection Interval for a Job* on page 153
• *Resuming and Suspending the Scheduler* on page 154
• *Viewing the Job Execution History* on page 154

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 151
• *Resuming and Suspending a Data Collection Job* on page 152
• *Adding a New Schedule to a Job* on page 152
• *Modifying the Data Collection Interval for a Job* on page 153
• *Resuming and Suspending the Scheduler* on page 154
• *Viewing the Job Execution History* on page 154

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

*   *Executing and Stopping a Data Collection Job* on page 151
*   *Deleting a Data Collection Job* on page 151
*   *Adding a New Schedule to a Job* on page 152
*   *Modifying the Data Collection Interval for a Job* on page 153
*   *Resuming and Suspending the Scheduler* on page 154
*   *Viewing the Job Execution History* on page 154

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 151
- *Deleting a Data Collection Job* on page 151
- *Resuming and Suspending a Data Collection Job* on page 152
- *Modifying the Data Collection Interval for a Job* on page 153
- *Resuming and Suspending the Scheduler* on page 154
- *Viewing the Job Execution History* on page 154

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 151
- *Deleting a Data Collection Job* on page 151
- *Resuming and Suspending a Data Collection Job* on page 152
- *Adding a New Schedule to a Job* on page 152
- *Resuming and Suspending the Scheduler* on page 154
- *Viewing the Job Execution History* on page 154

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

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

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 143
- *Creating an Alert* on page 133
- *Assigning a Role to a Login or a Group* on page 102
- *Configuring the E-mail Server* on page 100

Types, Severities, and States

Learn about the properties that define and control alerts.

An alert's type determines what causes it to fire.

Table 23. 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 24. 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* on page 156
- *Modifying an Alert* on page 157
- *Testing an Alert-Triggered Script* on page 158
- *Deleting an Alert* on page 159
- *Alert Subscriptions* on page 159
- *Alert Notifications* on page 161

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

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 Suppression 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

<u>**Testing an Alert-Triggered Script**</u>
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 155
- *Viewing Alerts* on page 156
- *Modifying an Alert* on page 157
- *Deleting an Alert* on page 159
- *Alert Subscriptions* on page 159
- *Alert Notifications* on page 161
- *Alert-Triggered Scripts* on page 143
- *Substitution Parameters for Scripts* on page 144
- *Creating an Alert* on page 133

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

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

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 25. 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 161

- *Enabling and Disabling Alert Subscription* on page 161

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 160
- *Enabling and Disabling Alert Subscription* on page 161

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 160
- *Unsubscribing from an Alert* on page 161

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 155
- *Viewing Alerts* on page 156
- *Modifying an Alert* on page 157
- *Testing an Alert-Triggered Script* on page 158
- *Deleting an Alert* on page 159
- *Alert Subscriptions* on page 159

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 162

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 162

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 an Adaptive Server* on page 114
* *Importing Resources for Batch Registration* on page 115

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 164
- *Removing a Resource from a Perspective* on page 164
- *Searching for Resources in the Resource Explorer* on page 165
- *Registering an Adaptive Server* on page 114
- *Importing Resources for Batch Registration* on page 115

Adding a Resource to a Perspective
Add one or more resources to the current perspective.

Prerequisites
Register the resources.

Task
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 163
- *Removing a Resource from a Perspective* on page 164
- *Searching for Resources in the Resource Explorer* on page 165
- *Registering an Adaptive Server* on page 114
- *Importing Resources for Batch Registration* on page 115

Removing a Resource from a Perspective
Remove one or more resources from the current perspective.

Removing a resource from a perspective does not unregister the resource; it remains in any other perspectives to which it has been added, and remains accessible in the Resource Explorer.

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 163
* *Adding a Resource to a Perspective* on page 164
* *Searching for Resources in the Resource Explorer* on page 165
* *Adding a Resource to a Perspective* on page 118

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
* *Unregistering a Resource* on page 163
* *Adding a Resource to a Perspective* on page 164
* *Removing a Resource from a Perspective* on page 164

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 118

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 166
* *Renaming a Perspective* on page 166

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 166
* *Renaming a Perspective* on page 166

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 166
* *Removing a Perspective* on page 166

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 perspective.

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.

Task	Action
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.
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 168

Arranging View Layout in a Perspective

Use the view layout options to manage your perspective space.

Click one of these icons in the Sybase Control Center toolbar:

Icon	Action
	Close All Open Views
	Minimize All Open Views
	Restore All Minimized Views
	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 167

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 169
- *Refreshing or Converting an Instance* on page 171
- *Removing an Instance* on page 172
- *Shared-Disk Mode* on page 173
- *sccinstance Command* on page 174

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
    ```

    On UNIX systems, omit the **-service** option.

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 169
- *Refreshing or Converting an Instance* on page 171
- *Removing an Instance* on page 172
- *Shared-Disk Mode* on page 173
- *sccinstance Command* on page 174
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71

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 169
- *Deploying an Instance from a Shared Disk Installation* on page 169
- *Removing an Instance* on page 172
- *Shared-Disk Mode* on page 173
- *sccinstance Command* on page 174

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 169
- *Deploying an Instance from a Shared Disk Installation* on page 169
- *Refreshing or Converting an Instance* on page 171
- *Shared-Disk Mode* on page 173
- *sccinstance Command* on page 174

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]
[-host host-name]
[-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. Sybase does not recommend using **-force** to remove or refresh a running instance in a Windows environment.
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-host** *host-name* – specify the host for this instance. Use with **-create**; required only when the instance name does not match server name of the host on which this instance will run. (The instance name defaults to the name of the current host unless you use **-instance** to specify another name.)
- **-instance [***instance-name***]** – 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.

  sccinstance assumes that the host name is the same as the instance name unless you use **-host** to specify a different host name.
- **-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

Port Name	Description	Service Names	Property Names	Default Port
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
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 on the current host, 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
```

- **Deploy two instances on the same host** – creates two agent instances on the host fireball. The first command does not need the **-host** option because the instance name is the same as the host name.

```
sccinstance -create -agent -instance fireball -portconfig rmi=9991
sccinstance -create -agent -instance fireball2 -host fireball
-portconfig rmi=9992
```

Note: In a production environment, Sybase recommends that you deploy no more than one SCC instance of each type (one server and one agent) on the same host.

- **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 169
• *Deploying an Instance from a Shared Disk Installation* on page 169
• *Refreshing or Converting an Instance* on page 171
• *Removing an Instance* on page 172
• *Shared-Disk Mode* on page 173
• *Instances* on page 169

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 180
- *Forcing an Immediate Backup* on page 181
- *Restoring the Repository from Backups* on page 181
- *Configuring Repository Purging* on page 183

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 179
- *Forcing an Immediate Backup* on page 181
- *Restoring the Repository from Backups* on page 181
- *Configuring Repository Purging* on page 183

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 179
- *Modifying the Backup Schedule* on page 180
- *Restoring the Repository from Backups* on page 181
- *Configuring Repository Purging* on page 183

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 182.

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

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

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 Adaptive Server Component Log

View event logs for Sybase Control Center for Adaptive Server.

The Sybase Control Center for Adaptive Server log files are located at:

- Windows – `%SCC_HOME%\SCC-3_2\plugins\ASEMAP\log\ASEMAP.log`
- UNIX – `$SCC_HOME/SCC-3_2/plugins/ASEMAP/log/ASEMAP.log`

See also

- *Viewing the Sybase Control Center Client Log* on page 185
- *Changing the Logging Level* on page 186
- *Logging or Message Levels* on page 187
- *Changing Logging Configuration* on page 187

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 Adaptive Server Component Log* on page 184
- *Viewing the Sybase Control Center Client Log* on page 185
- *Changing the Logging Level* on page 186
- *Logging or Message Levels* on page 187
- *Changing Logging Configuration* on page 187

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 Adaptive Server Component Log* on page 184

- *Viewing Sybase Control Center Server Logs* on page 185
- *Changing the Logging Level* on page 186
- *Logging or Message Levels* on page 187
- *Changing Logging Configuration* on page 187

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 Adaptive Server Component Log* on page 184
- *Viewing Sybase Control Center Server Logs* on page 185
- *Viewing the Sybase Control Center Client Log* on page 185
- *Logging or Message Levels* on page 187
- *Changing Logging Configuration* on page 187
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71

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 Adaptive Server Component Log* on page 184
* *Viewing Sybase Control Center Server Logs* on page 185
* *Viewing the Sybase Control Center Client Log* on page 185
* *Changing the Logging Level* on page 186
* *Changing Logging Configuration* on page 187
* *scc Command* on page 78

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

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

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

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 189

- *shutdown command* on page 191
- *status Command* on page 191

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 189
- *info Command* on page 190
- *status Command* on page 191

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

Settings

Learn about Adaptive Server monitoring controls on the Settings screen.

Table 26. Controls on the Settings screen

Control	Description	Default
Screen Refresh Interval (seconds)	The period between refreshes of screens in the Adaptive Server monitor. Refreshing a screen redraws it with the most recent available data.	30 seconds
Chart Trend Period (minutes)	The period of time over which data is displayed in historical charts on the Overview, Devices, Engines, and Segments screens, and on the Statistics Chart.	12 minutes
Alert List Size	The maximum number of alert notifications that can appear in the Alerts table on the Overview screen. When the Alerts table is full, the addition of a new alert notification causes the oldest notification to be removed.	100 alerts
Historical SQLs Size	The maximum number of active SQL statements that can appear in the Active SQLS table of the SQL Activity window. When the Active SQLS table is full, adding new SQL statements causes the oldest statement to be deleted.	500 statements
Historical SQLs Trend Period	The period of time in which SQL statements are displayed in the Active SQLS table of the SQL Activity window.	5 minutes

Manage and Monitor the Adaptive Server Environment

Monitor the performance, processes, databases, and other aspects of Adaptive Servers.

Managing an Adaptive Server

Sybase Control Center allows you to register and authenticate a running an agent, start and stop your server, and view the server error log.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.
2. Click **ASE Servers**.
   You see a list of monitored servers.
3. Select the server you want to manage and click the drop-down arrow that appears next to the name.
4. (Optional) To register the agent, click **Register Agent**. To clear agent registration, click **Clear Agent Registration**.
5. (Optional) To authenticate the agent, click **Authenticate Agent**. To clear agent registration, click **Clear Agent Authentication**.
6. (Optional) To start the Adaptive Server, click **Start Server**.
7. (Optional) To stop a running Adaptive Server, click **Stop Server**.
8. (Optional) To view the error log of an Adaptive Server, click **View Log**. You can further choose to filter the error log.
9. (Optional) To view the properties of the Adaptive Server and agent, click **Properties**.

Executing SQL Statements

Use Sybase Control Center to execute SQL statements on one or more Adaptive Servers.

You can use the Execute SQL view to run 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, SCC displays an error.

1. In the Administration Console, select one or more servers, click the drop-down arrow, and select **Execute SQL**.
2. Enter the SQL statements.
3. Click **Execute**.

   The query runs on all the servers you selected and results appear in the bottom portion of the view. The view includes a results tab for each server. On the tabs,
   • A green check indicates a successful query.

- A red X indicates an error. A tab with a red X also displays an error message.

Registering the Agent for an Adaptive Server

Use the Adaptive Server Administration Console to register the agent.

Note: For Adaptive Server versions 15.7 and earlier, the agent is the Unified Agent. For versions 15.7 ESD1 and higher, the agent is the SCC agent. Follow the steps provided here to register either type of agent.

The agent is installed and set up as a component of the Adaptive Server installation. For information, see the Adaptive Server installation guide for your platform.

You must register and authenticate the agent to use Sybase Control Center to perform certain administrative tasks such as starting the Adaptive Server or viewing the Adaptive Server error log.

You must register an agent for each Adaptive Server you have configured. The agent is configured on the same host as the Adaptive Server that it manages. When you register the agent, you are updating Sybase Control Center with information on the machine and port number on which the agent is configured.

For information about configuration options for the Unified Agent, see *Unified Agent and Agent Management Console Users Guide > Installing and Configuring Unified Agent and Agent Management Console* in the Adaptive Server documentation set.

For information on security features for the Unified Agent, see *Unified Agent and Agent Management Console Users Guide > Security* in the Adaptive Server documentation set.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.
2. Click **ASE Servers**.
   You see a list of monitored servers.
3. Select the server you want to manage and click the drop-down arrow that appears next to the name.
4. Click **Register Agent**.
   You see the Server Properties screen.
5. Enter the port number for the agent and click **Register**.

   Note: After the agent is registered, you can authenticate the agent, or clear the registration.

6. (Optional) Enter the login name and password for the agent, and click **Authenticate**.

See also

Authenticating the Unified Agent

For Adaptive Server versions 15.7 and earlier. Authenticate the Unified Agent using the Authenticate Agent option in the Adaptive Server Administration Console.

Note: While executing this command, you may experience performance degradation if there is a firewall between the Sybase Control Center server and the monitored Adaptive Server host machines. Communication with the agent uses the Java RMI network protocol, which requires a network connection between the Sybase Control Center server and agent. If the firewall prevents these connections from being established, then performance may degrade.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.
2. Click **ASE Servers**.
   You see a list of monitored servers.
3. Select the server you want to manage and click the drop-down arrow that appears next to the name.
4. Click **Authenticate Agent**.
   You see the Server Properties screen.
5. Enter the agent user name and password (optional) for the Unified Agent.

   You must have configured the Unified Agent to allow the Sybase Control Center user to authenticate. The Sybase Control Center user must also have the necessary permissions to manage the Adaptive Server using the Unified Agent.

   Note: The authentication credentials for the Unified Agent are different from those used to authenticate the Adaptive Server resource.

6. (Optional) Click **Clear Authentication**.

See also
- *Registering the Agent for an Adaptive Server* on page 194
- *Starting an Adaptive Server* on page 196
- *Stopping an Adaptive Server* on page 197
- *Viewing and Filtering the Adaptive Server Error Log* on page 197

Authenticating the SCC Agent

For Adaptive Server versions 15.7 ESD1 and later. Authenticate the SCC Agent using the Authenticate Agent option in the Adaptive Server Administration Console.

Note: While executing this command, you may experience performance degradation if there is a firewall between the Sybase Control Center server and the monitored Adaptive Server host machines. Communication with the SCC agent uses the Java RMI network protocol, which requires a network connection between the Sybase Control Center server and SCC agent. If the firewall prevents these connections from being established, then performance may degrade.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.

2. Click **ASE Servers**.
   You see a list of monitored servers.

3. Select the server you want to manage and click the drop-down arrow that appears next to the name.

4. Click **Authenticate Agent**.
   You see the Server Properties screen.

5. Enter the agent user name and password for the SCC Agent.

   You must have configured the SCC agent to allow the Sybase Control Center user to authenticate. The Sybase Control Center user must also have the necessary permissions to manage the Adaptive Server using the SCC agent.

   Note: The authentication credentials for the SCC agent are different from those used to authenticate the Adaptive Server resource.

6. (Optional) Click **Clear Authentication**.

Starting an Adaptive Server

Start an Adaptive Server using the Start Server option in the Administration Console.

Note: While executing this command, you may experience performance degradation if there is a firewall between the Sybase Control Center server and the monitored Adaptive Server host machines. Communication with the agent uses the Java RMI network protocol, which requires a network connection between the Sybase Control Center server and agent. If the firewall prevents these connections from being established, then performance may degrade.

Sybase Control Center uses the RUN server script to start the Adaptive Server.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.

2. Click **ASE Servers**.
   You see a list of monitored servers.

3. Select the server you want to manage and click the drop-down arrow that appears next to the name.

4. Click **Start Server**.

   Start-up logs and messages are displayed.

   Note: If the associated agent is not running on the host , or the agent is not registered or authenticated, the **Start Server** option is disabled.

5. Click **Yes**.

See also
- *Registering the Agent for an Adaptive Server* on page 194
- *Authenticating the Unified Agent* on page 195
- *Stopping an Adaptive Server* on page 197
- *Viewing and Filtering the Adaptive Server Error Log* on page 197

Stopping an Adaptive Server

Shut down an Adaptive Server using the Stop Server option in the Administration Console.

Note: While executing this command, you may experience performance degradation if there is a firewall between the Sybase Control Center server and the monitored Adaptive Server host machines. Communication with the agent uses the Java RMI network protocol, which requires a network connection between the Sybase Control Center server and agent. If the firewall prevents these connections from being established, then performance may degrade.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.
2. Click **ASE Servers**.
   You see a list of monitored servers.
3. Select the server you want to manage and click the drop-down arrow that appears next to the name.
4. Click **Stop Server**.
   You see the table of current server processes.

   Note: To see the running processes in the table, you must first have the resource authenticated.

5. (Optional) Select the option to shut down the server immediately.

See also
- *Registering the Agent for an Adaptive Server* on page 194
- *Authenticating the Unified Agent* on page 195
- *Starting an Adaptive Server* on page 196
- *Viewing and Filtering the Adaptive Server Error Log* on page 197

Viewing and Filtering the Adaptive Server Error Log

Apply filters and view the Adaptive Server error log.

Prerequisites

Authenticate the agent to enable the View Log command.

Task

Note: While executing this command, you may experience performance degradation if there is a firewall between the Sybase Control Center server and the monitored Adaptive Server host

machines. Communication with the agent uses the Java RMI network protocol, which requires a network connection between the Sybase Control Center server and agent. If the firewall prevents these connections from being established, then performance may degrade.

1. In the Perspective Resources view, select a resource, click the drop-down arrow that appears next to the name, and **Administration Console**.

2. Click **ASE Servers**.
   You see a list of monitored servers.

3. Select the server you want to manage and click the drop-down arrow that appears next to the name.

4. Click **View Log**.
   You see the error log of the Adaptive Server.

5. (Optional) In the **Show messages matching** field, enter a word or phrase to isolate all the lines of the log that contain it. For example, enter "memory" to find all log output pertaining to memory usage.

6. (Optional) Select the **Settings** option.
   You see the Server Log window.

   Note: User-defined regular expressions are valid only in the session in which they are defined. When the Sybase Control Server is restarted, you must redefine the regular expressions.

7. Specify how to retrieve log entries from the error log. You can retrieve all log entries, a specified number of log entries last accessed, all log entries from a specified time period, or all log entries matching a regular expression.

   The predefined regular expressions provided by the dialog are generally sufficient for filtering. They contain the most commonly accessed keywords in the error log such as **Error**, **deadlock**, **warning**, and so on.

   However, to define your own custom regular expression, click **add**, and specify a descriptive name and the regular expression clause. For example, to view lines of the log that refer to configuration, specify "All lines with configuration" for **Name** and "configuration" for **Regular Expression**.

8. Click **OK**.

See also
- *Registering the Agent for an Adaptive Server* on page 194
- *Authenticating the Unified Agent* on page 195
- *Starting an Adaptive Server* on page 196
- *Stopping an Adaptive Server* on page 197

Server Properties

Use the server Properties wizard to view server and agent information, to configure the agent, and to stop or restart the server.

Click **Properties** on your server name to initiate the Properties wizard.

Wizard Option	Server Properties
General	View server information such as: • Name and type • Version, build date, and build options • Edition – click **Details** for licensing information. • Character set, sort order, and language • Host name, port number, platform and operating system • Page size • Status and log file
Agent	View agent properties such as port number, user name, and version information. You can also authenticate or clear the authentication of your agent. View server properties such as version, server name, status, and start-up file. You can also start or stop your server using the agent.

Displaying the Performance Overview

The Overview screen shows Adaptive Server performance status.

Check the Overview window to find out whether the server is running, and details about memory usage, CPU utilization, recent alerts, and so on. Other windows in the Adaptive Server monitor display more detailed information about the status of individual server resources such as engines, databases, caches, and processes. In Adaptive Server cluster configurations, this window allows you to check whether a particular cluster is running, how many instances of the cluster are down, and so on.

Note: The **Overview** screen is called **Cluster Overview** in Adaptive Server cluster configurations.

1. In the Perspective Resources view, select the server, click the drop-down arrow, and select **Monitor**. Alternately, in the Administration Console, select the server, click the drop-down arrow, and select **Monitor**.
   The Adaptive Server monitor opens and displays the Overview screen. Check the server information in the upper left corner of the screen for the server's name, software product and version, its hardware platform, and an indication of whether the server is running. For cluster configurations, you also see the status of instances of each cluster, and the number of blocked process.

> **Note:** When Sybase Control Center shows a server status of "Stopped," it means that the server is unreachable over the network.

2. (Optional) If data collections are running, mouse over the **Engine CPU Utilization** graph to display precise figures (values, times, and dates) for points on the curve.

   The graph shows the aggregate CPU utilization for all engines on the server. For cluster configurations, the graph shows the aggregate CPU utilization for each instance of the cluster.

3. (Optional – not in Adaptive Server cluster configurations) Move your mouse over the **Device IO/Sec** graph to display precise figures (values, times, and dates) for points on the curve.

   The graph shows device I/O aggregated across all devices on the server.

4. (Optional – not in Adaptive Server cluster configurations) Look at the Processes chart (far right) to see the number of configured and currently running processes and the highest number of concurrent processes since the server started, as well as the number of blocked processes.

5. (Optional – not in Adaptive Server cluster configurations) Look at the Memory chart to see statistics on caches and on physical, logical, and unused memory.

6. (Optional) Click a tab to see information about the resource you want to monitor:

   • **Details** - displays the version, edition, platform, number of deadlocks, platform, page size, device size and counters cleared for the Adaptive Server.

   • **Configured Resources** - displays, in tabular form, the configurable resources for each server or cluster instance. Each configuration option is displayed along with its currently configured value, run value that is currently used by the server, percentage of the resource that is currently in use, and the high water mark, which is the maximum amount of resource that has been used since the server was booted. Any column can be used to sort the table.

     The configured value for a resource is an editable field, denoted as such by a "pencil" **Edit** icon. Input a new numerical value for one or more resources, then choose either:

     • **Save All** to update the server with the new values. Sybase Control Center displays the new values. If the Adaptive Server encounters an error while applying the new value for a resource, Sybase Control Center displays the error below the table, and also next to the changed field in the row that causes the server error.

     • **Reset All** to restore the original value for the resource.

     > **Note:** Resources for each server or cluster instance may also be configured on the Server Configurations window.

   • **Wait Events** - displays a list of server-wide wait events that can be very useful in performance tuning. Information about the wait-events includes the number of waits, wait time, average wait time, and wait description. For clustered servers, this information is displayed per instance.

- **Licenses** - displays a list of licenses that are currently checked out by the server or cluster instances. There is also information about the number, type, status (expirable, permanent and so on), and expiration date of each license.
- **Alerts** - displays a list of all fired alerts configured at server, cluster or cluster instance levels. For each alert, there is information about the time at which the alert was fired, severity, current statistic and threshold.

7. (Optional - Adaptive Server Cluster configurations) Mouse over the **Cluster Instances** graph to display precise information for points on the bar graph.

8. (Optional - Adaptive Server Cluster configurations) If data collections are running, mouse over the **Logical Cluster** graph to display precise information for points on the bar graph.

See also
- *Displaying the Cluster Overview* on page 219
- *Interpreting Statistics* on page 371

Performance Overview Statistics and Details
Learn about the information presented on the Overview screen.

The Overview screen displays high-level information about this Adaptive Server. The tables and charts are populated by data from the collection_ase_all_client_kpis, covering the current chart trend period.

Note: The **Overview** screen is called **Cluster Overview** in Adaptive Server cluster configurations.

Engine CPU Utilization	Displays aggregate CPU utilization for all engines on this server. For information on individual Adaptive Server engines, see the Engines screen. (Because all I/O for a process goes through one engine, CPU usage is not always evenly distributed across engines.) For cluster configurations, the graph shows the aggregate CPU utilization for each instance of the cluster.
Device IO/Sec	(For Adaptive Server cluster configurations, this information is on the Cluster Instances window) Displays device I/O per second aggregated across all devices on this server. For information on individual devices, see the Devices screen.

Memory	(For Adaptive Server cluster configurations, this information is on the Cluster Instances window) Displays memory usage statistics, including: • The amounts of physical and logical memory in use • The amount of unused memory • The size of the procedure, statement, and data caches
Processes	Displays process statistics, including: • Max User Processes – number of processes this server is configured for • High Water Mark – highest number of processes that ran concurrently since this server started • Active – processes running now • (Adaptive Server cluster configurations only) Blocked Processes – processes that are waiting for a resource or for another process to finish For more on processes, see the Processes screen.
Details tab	Displays information about this server, including number of days it has been running, number of deadlocks, data cache hit rate, procedure cache stalls, page and device sizes, maximum number of online engines, number of open databases, and the dates and times of the server's most recent restart and of the clearing of these counters. For cluster configurations, the Details tab also displays information about software product and version, hardware platform, server edition , number of dead locks, page and device sizes, and when the server was last restarted, and the counters last cleared.

Configured Resources tab	Displays usage statistics for many of the configured resources for this server or cluster instance:
	• Current – amount of this resource the server is using now • Run value – configured maximum for this resource • Percentage – percentage of the configured maximum represented by the current use of this resource • High Water Mark – maximum amount of this resource that has been used since this server started
	Use the Percentage and High Water Mark columns to identify resources that might be over-configured or under-configured.
	For information on configured resources, see the Adaptive Server Enterprise documentation.
Wait Events	Displays a list of server-wide wait events that can be very useful in performance tuning. Information about the wait-events includes the number of waits, wait time, average wait time, and wait description. For clustered servers, this information is displayed per instance.
Licenses tab	Displays information about software licenses for Adaptive Server Enterprise on this server or cluster instance.
Alerts tab	Displays, for this server, cluster, or cluster instance, all alert notifications that have occurred since the Adaptive Server monitor was opened. If any alerts have occurred since you last looked at the Alerts tab, a yellow warning icon appears on the tab.
	You can control the number of alerts displayed using the Alert List Size property on the Settings screen.

See also

Settings

Learn about Adaptive Server monitoring controls on the Settings screen.

Table 27. Controls on the Settings screen

Control	Description	Default
Screen Refresh Interval (seconds)	The period between refreshes of screens in the Adaptive Server monitor. Refreshing a screen redraws it with the most recent available data.	30 seconds
Chart Trend Period (minutes)	The period of time over which data is displayed in historical charts on the Overview, Devices, Engines, and Segments screens, and on the Statistics Chart.	12 minutes
Alert List Size	The maximum number of alert notifications that can appear in the Alerts table on the Overview screen. When the Alerts table is full, the addition of a new alert notification causes the oldest notification to be removed.	100 alerts
Historical SQLs Size	The maximum number of active SQL statements that can appear in the Active SQLS table of the SQL Activity window. When the Active SQLS table is full, adding new SQL statements causes the oldest statement to be deleted.	500 statements
Historical SQLs Trend Period	The period of time in which SQL statements are displayed in the Active SQLS table of the SQL Activity window.	5 minutes

Caches

Monitor Adaptive Server data, procedure, statement caches, and in-memory storage.

Monitor Caches

From the Perspective Resources view, use the Monitor option on your Adaptive Server to display cache statistics and monitor cache activity.

Monitoring Adaptive Server Data Caches

Display information about data caches.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Caches**.

3. Click the **Data Caches** tab.

4. Select a cache in the **Data Caches** table.
   The tabs at the bottom of the screen are populated with information about the selected cache.

5. (Optional for Adaptive Server cluster configurations) Select **Global** for information about global data caches. Select **Local** for information on data caches in each cluster instance.

6. (Optional for Adaptive Server cluster configurations) Select a cache in the **Global Data Caches** or **Local Data Caches** table.
   The tabs at the bottom of the screen are populated with information about the selected cache.

7. Click the tabs to display **Pool Information** or **Cached Objects** details.

8. (Optional for Adaptive Server cluster configurations) Click the tab to display **Distribution** details for a selected global data cache.

For more information on data caches, see *Adaptive Server System Administration Guide Volume 2*.

See also

Modifying Data Cache Sizes
Sybase Control Center allows you to modify the size of the data cache, and specify the number of partitions in the data cache.

You need **sa_role** to modify data caches. For more details on administering data caches, see the *Adaptive Server System Administration Guide*.

1. In the Adaptive Server monitor, select **Caches**.

2. From the Data Caches table, select the cache to configure.

3. Right-click the selected cache to display the **Resize** option that allows you to modify the size of the data cache.
   You see the current size of the data cache, available space, and partition information.

4. Enter the new size of the data cache.

An increase in the data cache size is immediately effected in the server. A decrease in the data cache size requires the server to be re-started for the change to take effect.

5. (Optional) Enter a new value for the partitions in the data cache.

6. (Optional) Use **Calculate Overhead** to calculate the amount of memory required to resize the data cache with the new input size.

7. Click **Save**.
   The dialog box closes if the operation succeeds, else Sybase Control Center displays an error and the dialog box stays open.

Adding Data Cache Buffer Pools

Sybase Control Center allows you to change the configuration of your data cache buffer pools.

1. In the Adaptive Server monitor, select **Caches**.

2. From the Data Caches table, select the cache to configure.

3. Right-click the selected cache to display the **Add Buffer Pool** option that allows you to modify the configuration of the data cache. Alternately, click **Add Buffer Pool** in the Pool Information table.

4. To configure the new buffer pool, enter values for I/O buffer size, amount in pool, wash size, and a prefetch limit. You may also select the pool from which space will be allocated for the changed configuration.

5. Click **Save**.
   The dialog box closes if the operation succeeds, else Sybase Control Center displays an error and the dialog box stays open.

Data Cache Statistics and Details

Interpret the Data Cache window for Adaptive Server.

The Data Cache window displays information about Adaptive Server data caches. In Adaptive Server cluster configurations, the **Global Datacaches** window provides information about the global Adaptive Server data caches and the **Local Datacaches** screen provides information about local data caches grouped by cluster instances.

The Data Caches table shows the size and level of activity in each data cache, including hit rate (the percentage of database requests that can be answered from the cache), volatility, number of partitions in this cache, relaxed replacement, and physical reads and writes. Select a cache in the table to populate the tabs at the bottom of the screen with details about that cache.

Right-click on a selected cache to use **Resize** to modify the size of the data cache, or **Add Buffer Pool** to change the configuration of your data cache buffer pool.

Table 28. Tabs on the Data Cache screen

Pool Information	Shows information about the pools of different sizes that are used to optimize I/O in the selected cache. Details include size, usage, reread ratio, physical and dirty reads, pages touched, buffers to MRU (most recently used), and buffers to LRU (least recently used). The Buffers to MRU and Buffers to LRU columns show buffers added to the ends of the buffer list. The oldest buffers (the least recently used) are flushed first.
Cached Objects	Lists the tables in the selected data cache and their size, in kilobytes. Click on the **Cached Size** column heading to sort the table by size.
Distribution	Shows information about the distribution metrics for each instance of a cluster, including hit rate, volatility, cache partitions, relaxed replacement, and physical reads and writes.

See also
- *Procedure Cache Statistics and Details* on page 208
- *Statement Cache Statistics and Details* on page 210
- *In-memory Storage Statistics and Details* on page 211
- *Monitoring Adaptive Server Data Caches* on page 204
- *Monitoring the Adaptive Server Procedure Cache* on page 207
- *Monitoring the Adaptive Server Statement Cache* on page 208
- *Monitoring Adaptive Server In-memory Storage* on page 210

Monitoring the Adaptive Server Procedure Cache
Display information about the procedure cache.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Caches**.
3. Click **Procedure Cache**.
4. (Optional for Adaptive Server cluster configurations) Select a specific cluster instance. The Cached Procedures table shows all the procedures for the selected instance.

For more information on the procedure cache, see *Adaptive Server Performance and Tuning Series: Basics.*

See also

- *Monitoring Adaptive Server Data Caches* on page 204
- *Monitoring the Adaptive Server Statement Cache* on page 208
- *Monitoring Adaptive Server In-memory Storage* on page 210
- *Data Cache Statistics and Details* on page 206
- *Procedure Cache Statistics and Details* on page 208
- *Statement Cache Statistics and Details* on page 210
- *In-memory Storage Statistics and Details* on page 211

Procedure Cache Statistics and Details

The Procedure Cache screen displays information about the contents of the Adaptive Server procedure cache which is a memory pool used for stored procedures and a variety of other objects.

The Adaptive Server functions that use the procedure cache are called modules—there are over 20 modules in the system. The bar chart on this screen, Top 10 Procedure Cache Module Users, shows the modules that use the cache most heavily. The Procedural Objects module contains stored procedures; there is also a module for the statement cache. Use the bar chart to see which parts of the system are using the procedure cache.

The Cached Procedures table lists the stored procedures in the cache (in the Procedural Objects Module). For each stored procedure, it gives the name, database name, cached size, owner's name, compile date, and plan ID.

For Adaptive Server cluster configurations, information for each selected instance of the cluster is depicted in the Procedure Cache Summary, and includes the hit ratio, number of stalls or dirty reads, amount of procedure cache memory that is currently in use, and the total procedure cache memory allocated.

See also

- *Data Cache Statistics and Details* on page 206
- *Statement Cache Statistics and Details* on page 210
- *In-memory Storage Statistics and Details* on page 211
- *Monitoring Adaptive Server Data Caches* on page 204
- *Monitoring the Adaptive Server Procedure Cache* on page 207
- *Monitoring the Adaptive Server Statement Cache* on page 208
- *Monitoring Adaptive Server In-memory Storage* on page 210

Monitoring the Adaptive Server Statement Cache

Display information about the statement cache, including cached SQL queries.

Note: To enable the Statement Cache tab, the statement cache must be configured in Adaptive Server and the **enable stmt cache monitoring** option must be turned on in Adaptive Server.

Turning on the **enable stmt cache monitoring** option is not necessary to enable monitoring of an Adaptive Server.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Caches**.
3. Click the **Statement Cache** tab.

    The Statement Cache tab is disabled if statement cache monitoring is not enabled. Statement cache monitoring is controlled by the two Adaptive Server configuration options **enable stmt cache monitoring** and **statement cache size**.
4. (Optional for Adaptive Server cluster configurations) Select a specific cluster instance. The Cached Statements table shows all the statements for the selected instance.
5. Select a statement in the **Cached Statements** table.
    The SQL query appears at the bottom of the screen.
6. Click the tabs to display pool information or details about cached objects.

For more information on the statement cache, see *Adaptive Server System Administration Guide, Volume 2*.

See also
- *Monitoring Adaptive Server Data Caches* on page 204
- *Monitoring the Adaptive Server Procedure Cache* on page 207
- *Monitoring Adaptive Server In-memory Storage* on page 210
- *Data Cache Statistics and Details* on page 206
- *Procedure Cache Statistics and Details* on page 208
- *Statement Cache Statistics and Details* on page 210
- *In-memory Storage Statistics and Details* on page 211

Setting the Statement Cache Size
You can use one of two methods to set the statement cache size.

If you do not configure the statement cache in the monitored Adaptive Server, Sybase Control Center displays an error when you open the **Statement Cache** tab on the Adaptive Server Monitor View **Caches** screen:

```
[error#=12052] Collection of monitoring data for table
'monCachedStatement' requires that the 'statement cache size'
configuration option(s) be enabled. To set the necessary
configuration, contact a user who has the System Administrator
(SA) role.
```

Use either method to enable the statement cache:

- In the Sybase Control Center, open the Adaptive Server Monitor View Server or Administration Console Configuration screen and change the value of *statement cache size* to a non-zero value.
- Log in to the Adaptive Server and execute the **sp_configure** system procedure to set the value of the statement cache size parameter to a non-zero value.

Statement Cache Statistics and Details

The Statement Cache window displays information about SQL queries and query plans stored in the Adaptive Server statement cache.

The Statement Cache Summary gives details about the size, hit count, and traffic in the cache. The Cached Statements table lists SQL statements by statement ID (SSQLID), and gives the owner name, use count, CPU time, elapsed time to execute, and logical and physical I/O figures for each query.

The Cached Statement Text window displays the query selected in the Cached Statements table.

Note: A cached query and query plan can be reused only by the user who first entered the query. Thus, if two or more users enter the same query, it appears in the cache several times.

See also
- *Data Cache Statistics and Details* on page 206
- *Procedure Cache Statistics and Details* on page 208
- *In-memory Storage Statistics and Details* on page 211
- *Monitoring Adaptive Server Data Caches* on page 204
- *Monitoring the Adaptive Server Procedure Cache* on page 207
- *Monitoring the Adaptive Server Statement Cache* on page 208
- *Monitoring Adaptive Server In-memory Storage* on page 210

Monitoring Adaptive Server In-memory Storage

Sybase Control Center provides detailed information about in-memory storage.

Follow this procedure only if you have configured an in-memory cache. In-memory databases are not currently supported in Adaptive Server 15.5 Cluster Edition, or in any release earlier than Adaptive Server 15.5.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Caches**.
3. Click the **In-memory storage** tab.
4. Select a cache in the **In-memory storage** table.
   The tabs at the bottom of the screen are populated with information about the selected in-memory storage cache.

5. Click the tabs to display **In-memory Devices**, **In-memory Databases**, or **Cached Objects** details.

For more information on in-memory storage, see *Adaptive Server In-memory Database Users Guide*.

See also
*   *Monitoring Adaptive Server Data Caches* on page 204
*   *Monitoring the Adaptive Server Procedure Cache* on page 207
*   *Monitoring the Adaptive Server Statement Cache* on page 208
*   *Data Cache Statistics and Details* on page 206
*   *Procedure Cache Statistics and Details* on page 208
*   *Statement Cache Statistics and Details* on page 210
*   *In-memory Storage Statistics and Details* on page 211

In-memory Storage Statistics and Details
Interpreting the data on the In-memory Storage window.

The in-memory storage window displays information about Adaptive Server in-memory caches, the devices that are created from this cache, and the databases on these devices.

The In-memory Storage table shows details of in-memory storage including the size and unused size, in megabytes, and the number of partitions. Select a cache in the table to populate the tabs at the bottom of the screen with details about that cache.

Table 29. Tabs on the In-memory Storage window

In-memory Devices	Shows information about the devices that are created from in-memory storage. Details include name, size, space used, start page and number of pages of memory usage.
In-memory Database	Shows information such as name and size of databases that are created on in-memory cache.
Cached Objects	Lists the tables and table indexes in the selected data cache, and their cached size, in kilobytes. Click on the **Cached Size** column heading to sort the table by size.

See also
*   *Data Cache Statistics and Details* on page 206
*   *Procedure Cache Statistics and Details* on page 208
*   *Statement Cache Statistics and Details* on page 210
*   *Monitoring Adaptive Server Data Caches* on page 204

Manage Caches

Sybase Control Center allows you to create, delete, and generate data definition language for caches.

Creating a Cache

Sybase Control Center allows you to create new data caches.

1. In the left pane of the Administration Console, expand **Server > Space Management > Caches**.

2. Select **New**.

3. Fill out the appropriate information:
   - **Introduction** – select the server in which to create a cache.
   - (Optional) **Create this as a local cache** – In a cluster environment, you can select to create the cache as a local cache, then select the online instance on which to create the cache.
   - **Cache Name** – enter the name of the cache to create.
   - **Cache Size** – enter the size of the new cache, which must be at least 512KB, but can be no larger than the unconfigured amount remaining on the server.

     (Optional) To determine if the server can manage the cache size, enter the size and click **Calculate overhead**. The wizard calculates the overhead necessary for the specified cache size.
   - **Type of Cache** – choose what you want the cache to store:
     - Data and log pages
     - Only log pages
     - In-memory database – this option is available only on Adaptive Server version 15.5 and later

4. Click **Summary** to see the cache options you selected.

5. (Optional) Click **Preview** to view the Transact-SQL syntax used to create the cache.

See also

Cache Properties

Use the cache Properties option to modify default data cache sizes, buffer pool values, and cache bindings. If you created the cache in an in-memory database, use the Properties option to also view in-memory database and in-memory device information.

Click **Properties** on your cache to view its properties.

Option	Cache Properties
General	Current size – select the display format: • Pages • KB • MB • GB
Configuration	• Currently configured – You can change the size of the data cache. • Show in – Allows you to specify the format in which to show cache size, including pages, kilobytes, megabytes, and gigabytes. See *Managing Cache Configurations* on page 213.
Buffer Pool	Current buffer pool values for regular caches – You can add, change, or remove buffer pools. See *Managing Buffer Pools* on page 214.
Cache Bindings	Shows object bindings for databases, tables, and indexes for regular caches – You can add, change, or remove cache bindings. See *Managing Binding Options* on page 215.
In-Memory Database	Displays the in-memory database that is created on this cache. This option is not available for regular caches.
In-Memory Device	Displays the in-memory device list that occupies this cache. This option is not available for regular caches.

See also
*   *Managing Cache Configurations* on page 213
*   *Managing Binding Options* on page 215
*   *Managing Buffer Pools* on page 214

Managing Cache Configurations

Change the size of your data cache.

1. In the left pane of the Administration Console, expand **Server > Space Management > Caches**.

2. In the **Name** field, select the cache to modify.

3. Click the arrow and select **Properties** to see the cache properties.

4. Click **Configuration** to see the information for your data cache.

5. (Not available for in-memory storage caches) Choose whether the cache is stored as data and log pages, or only as log pages. You cannot change the type for default cache, which is configured for data and log pages.

6. (Not available for in-memory storage caches) In the **Currently configured** field, specify the size of the data cache.

   Note: **Current size** indicates how much unused space remains in your specified data cache, while **Available space** shows the amount of additional memory available in Adaptive Server for all caches.

   Calculate Overhead allows you to see how much overhead you need to manage your data cache.

See also
- *Managing Binding Options* on page 215
- *Managing Buffer Pools* on page 214
- *Cache Properties* on page 213

Managing Buffer Pools
You can add and change buffer pools for your data cache.

1. In the left pane of the Administration Console, expand **Server > Space Management > Caches**.

2. In the **Name** field, select the cache to modify.

3. Click the arrow and select **Properties** to see the cache properties.

4. Click **Buffer Pool** to see a list of existing buffer pool values.

5. Modify the buffer pool allocation for your data cache:

   - Click **Add** to add an additional memory pool to the existing data cache, and specify:
     - **I/O buffer size** – from the drop-down menu, select the size of your I/O buffer in kilobytes.
     - **Amount in pool** – set the desired size and select the size format. The default size format is MB.
     - **Wash size** – set the desired wash size—the point in the cache at which Adaptive Server writes dirty pages to disk for a memory pool—and size format. The default is KB.
     - **Local async prefetch limit** – set the percent of buffers in the pool that you can use to hold buffers that have been read into cache by asynchronoud prefetch, but have yet to be used.

- **Affected pool** – specify the amount of memory, in kilobytes, the new pool should take from the existing pool. The menu lists the existing buffer pool you added to the cache. Since there is only a 2KB page-sized pool in the cache, you can add a new buffer pool only by taking part of the size from the 2KB page-sized pool.
- Select a buffer pool and click **Change** to change the memory pool settings. You see the same Add/Change Memory Pool dialog, with fewer options to modify:
  - Wash size
  - Local async prefetch limit
  - Affected limit – since you do not affect other pools when you change an existing pool, set this to null.
- Select the buffer pool to delete, and click **Remove** to remove any additional buffer pools you created. You cannot remove the default buffer pool.

See also
- *Managing Cache Configurations* on page 213
- *Managing Binding Options* on page 215
- *Cache Properties* on page 213

Managing Binding Options
You can add and change object bindings for your data cache.

1. In the left pane of the Administration Console, expand **Server > Space Management > Caches**.
2. In the **Name** field, select the cache to modify.
3. Click the arrow and select **Properties** to see the cache properties.
4. Click **Cache Bindings**. You see a list of cached bindings for databases, tables, or indexes depending on what you select from the **Show object bindings for**.
5. Modify cache bindings for your data cache:

   - Add a binding – Select the database, table, or index and click **Bind** to bind a new object to the cache within your selected scope. If you do not select a scope, the default is database.
   - Delete a cache binding – Select the bound database, table, or index object and click **Unbind**.

   (Optional) Click **Properties** to see the detailed properties of the object you select.

See also
- *Managing Cache Configurations* on page 213
- *Managing Buffer Pools* on page 214
- *Cache Properties* on page 213

Generating DDL for Caches

You can generate and view data definition language (DDL) statements to create, modify, or remove data caches for objects such as databases, tables, and indexes.

1. In the left pane of the Administration Console, expand **Server > Space Management > Caches**.
2. In the **Name** field, select the cache to modify.
3. Click the arrow and select **Generate DDL**.
   You see the DDL for the selected cache.
4. (Optional) Click **Save** to export and save the DDL statement.

See also
* *Creating a Cache* on page 212
* *Deleting a Cache* on page 216

Deleting a Cache

Sybase Control Center helps you delete database objects, or the database itself.

Note: You cannot delete the default data cache.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your cache.
3. Click the Name field of the cache you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the cache.
6. Confirm the deletion.
7. Click **Finish**.

See also
* *Creating a Cache* on page 212
* *Generating DDL for Caches* on page 216

Clusters

Monitor Clusters

Monitor Adaptive Server cluster configurations. Monitored resources include memory usage, device I/O, engine CPU utilization, connection information, interprocess communication, workload management, load profiles, routes, and so on.

Cluster Instances

Monitor instances of an Adaptive Server cluster configuration including memory usage, device I/O, engine CPU utilization, and connection information.

See also

Monitoring Cluster Instances in Adaptive Server Cluster Configurations

The Cluster Instances window displays details about memory and device usage of all cluster instances in an Adaptive Server cluster configuration.

Use the Cluster Instances window to gather information about memory usage, processes, user connections, and device I/O per instance of the cluster.

1. In the Perspective Resources view, select a shared-disk cluster, click the drop-down arrow, and click **Monitor**. This opens the Adaptive Server monitor. Use one of these options to show the Cluster Instances screen:

    * Select **Cluster Instances** from the left pane.

   The Cluster Instances screen displays information pertaining to each cluster instance.

2. (Optional) Click a tab at the bottom of the screen to display more information about individual instances:

    * **Details** – displays two charts, **Memory Usage** and **Processes**. **Memory Usage** represents the memory (physical and logical) and cache usage (procedure, statement and data) of a single cluster instance. **Processes** gives a summary of running, blocking, and blocked processes on a single cluster instance.
    * **Advanced** – displays graphs **Device I/O per Sec**, **Active Connections**, and **Engine CPU Utilization**. **Device I/O per Sec** represents the reads and writes of the selected instance during a certain time period. **Active Connections** represents the number of user connections created on a selected instance during a certain time period. **Engine CPU Utilization** represents the aggregate CPU utilization by all engines of a selected instance.

See also

Cluster Instances Statistics and Details

Interpret the Cluster Instances window for instances of an Adaptive Server cluster configuration.

The Cluster Instances screen displays information pertaining to each cluster instance in graph and chart formats.

The **Memory Usage** chart represents the memory (physical and logical) and cache usage (procedure, statement and data) of a single cluster instance. The **Processes** chart gives a summary of running, blocking, and blocked processes on a single cluster instance.

The **Device I/O per Sec** graph represents the reads and writes of the selected instance during a certain time period. The **Active Connections** graph represents the number of user connections created on a selected instance during a certain time period. The **Engine CPU Utilization** graph represents the aggregate CPU utilization by all engines of a selected instance.

Right-click on a selected cache to use **Resize** to modify the size of the data cache, or **Add Buffer Pool** to change the configuration of your data cache buffer pool.

Table 30. Tabs on the Cluster Instances window

Details	Displays the Memory Usage and Processes charts.
Advanced	Displays the Device I/O per Sec, Active Connections, and Engine CPU Utilization graphs.
Workload Status	Displays the workload metric and the corresponding base metric value.

See also

<u>*Cluster Interconnect*</u>
Monitor interprocess communication in an Adaptive Server cluster configuration.

See also

Displaying the Cluster Overview
The Cluster Overview screen shows performance status for Adaptive Server cluster configurations.

Check the Cluster Overview window to find out whether the server is running, and details about memory usage, CPU utilization, recent alerts, and so on. Other windows in the Adaptive Server monitor display more detailed information about the status of individual server resources such as engines, databases, caches, and processes.

In Adaptive Server cluster configurations, this window allows you to check whether a particular cluster is running, how many instances of the cluster are down, and so on.

See also

- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 221
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 222

Cluster Interconnect Statistics and Details

Interpret the Interconnect window for Adaptive Server cluster configurations.

The Cluster Interconnect window displays information pertaining to each cluster instance, including instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches.

Table 31. Tabs on the Cluster Interconnect window

CIPC	Displays the current status for each CIPC link. Status information for each instance includes the instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches.
CIPC Links	Displays the current status for each CIPC link. Status information for each instance includes the ID, local and remote interfaces, passive and active states, and ages of passive and active states.
CIPC Mesh	Displays cluster instance name, channel name, far-end cluster instance, received and transmitted message counts, and dropped, re-sent and retried messages. There is a Message Send Queue summary with send queue and sent queue information. The send queue includes current messages waiting to be sent to the instance, while the sent queue includes sent messages whose notification has not been processed.
CIPC Endpoints	Displays cluster instance name, endpoint name, received and transmitted message counts, and received and transmitted byte counts. There is also a Message Received Summary with received queue and done queue information. The received queue includes messages queued for this logical endpoint. The done queue includes messages for this logical endpoint that have been processed, and await further action.

See also
- *Cluster Instances Statistics and Details* on page 218

Monitoring Interprocess Communication in Adaptive Server Cluster Configurations
The Cluster Interconnect window provides information about Interprocess Communication in an Adaptive Server cluster configuration.

The Cluster Interconnect window provides detailed interprocess communication information for each instance of the cluster.

1. In the Perspective Resources view, select a shared-disk cluster, click the drop-down arrow and click **Monitor**. This opens the Adaptive Server monitor.
2. To show the Cluster Interconnect screen, select **Cluster Interconnect** from the left pane. The Cluster Interconnect window displays information pertaining to each cluster instance, including instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches. Additional tabs show more information about CIPC (cluster interprocess communication) links, messages, channels, and endpoints.
3. (Optional) Click a tab at the top of the screen to display more CIPC information on individual instances. The tab **CIPC** is selected by default.

   - **CIPC** – displays the current status for each CIPC link. Status information for each instance includes the instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches.
   - **CIPC Links** – displays the current status for each CIPC link. Status information for each instance includes the ID, local and remote interfaces, passive and active states, and ages of passive and active states.
   - **CIPC Mesh** – displays cluster instance name, channel name, far-end cluster instance, received and transmitted message counts, and dropped, re-sent and retried messages. There is a Message Send Queue summary with send queue and sent queue information. The send queue includes current messages waiting to be sent to the instance, while the sent queue includes sent messages whose notification has not been processed.
   - **CIPC EndPoints** – displays cluster instance name, endpoint name, received and transmitted message counts, and received and transmitted byte counts. There is also a Message Received Summary with received queue and done queue information. The received queue includes messages queued for this logical endpoint. The done queue

includes messages for this logical endpoint that have been processed, and await further action.

See also

Workload Management

Monitor workload for an Adaptive Server cluster configuration including load profiles, load scores, routes, connections, and states.

See also

Monitoring Workloads in Adaptive Server Cluster Configurations

The Workload Management window provides, for each instance of the cluster, detailed information on logical clusters, load profiles, routes, and system information.

1. In the Perspective Resources view, select a shared-disk cluster, click the drop-down arrow and click **Monitor**. This opens the Adaptive Server monitor. Select **Workload Management** from the left pane.

2. (Optional) Click a tab at the top of the screen to display more workload information on individual instances. The tab **Logical Clusters** is selected by default.

   * **Logical Clusters** – displays the logical cluster name, state, active connections, base instances, active base instances, failover instances, active failover instances, and load profile name. The possible values for state are: online, offline, failed, inactive, and time_wait. The type may be application, login, or alias. The General, Base Instances,

Failover Instances, and Routes tabs at the bottom of the screen provide access to further information about a selected logical cluster.

(Optional) To see the system view, start-up mode, failover mode, down-routing mode, login distribution mode, and logical cluster role for the selected logical cluster, select **General**.

(Optional) To see the ID, name, and state of the selected logical cluster select **Base Instances**.

(Optional) To see the ID, name, state, and failover group of the selected logical cluster, select **Failover Instances**.

(Optional) To see the name and type of route associated with the selected logical cluster, select **Routes**.

- **Workloads** – displays a list of workloads with information on instance, load profile, load score, user connections, CPU busy, run queue length, I/O load, engine deficit, and user score.

  (Optional) To see charts depicting load score, and percentage use of the CPU, select **Details** for a selected workload.

  (Optional) To see charts depicting queue length, and I/O load, select **Advanced** for a selected workload.

- **Load Profiles** – displays a list of load profiles with information on name, type, minimum load score, login redirection, and dynamic migration.

  (Optional) To see a chart of metrics, select **Metric Weight** for a selected load profile. The metrics include user connections, CPU busy, run queue length, and so on, and the corresponding weights associated with the metric.

- **Routes** – displays a list of routes with the name of the route, the logical cluster it is defined on, and the type of route such as alias, application, or login.

See also

Workload Management Statistics and Details
Interpret the Workload Management window for Adaptive Server cluster configurations.

The Workload Management window displays, for each instance of the cluster, detailed information on logical clusters, load profiles, routes, and system information.

Table 32. Tabs on the Workload Management window

Logical Clusters	Displays the logical cluster name, state, active connections, base instances, active base instances, failover instances, active failover instances, and load profile name. The possible values for state are: online, offline, failed, inactive, and time_wait. The type may be application, login, or alias.
	Routes depicts the name and type of route associated with the selected logical cluster.
	Failover Instances depicts the ID, name, state, and failover group of the selected logical cluster.
	Base Instances depicts the ID, name, and state of the selected logical cluster.
	General depicts the system view, start-up mode, failover mode, down-routing mode, login distribution mode, and logical cluster role for the selected logical cluster.
Workload Scores	Displays a list of workloads with information on instance, load profile, load score, user connections score, CPU busy score, run queue length score, I/O load score, engine deficit score, and user score.
	Details displays charts depicting load score, and percentage use of the CPU, for a selected workload.
	Advanced displays charts depicting run queue length score, and I/O load score, for a selected workload.
	Note: Elements of the Adaptive Server monitor screens may be updated at different intervals depending on the mechanism that is used to collect the data. For example, the historical charts are updated based on the frequency of the **all_client_kpis** collection while the data in the tables on the screen are updated with the frequency that the user has set for the screen refresh interval on the Settings window.

Base Metric Values	Displays a list of instances and their user connections, CPU busy percentage, run queue length, I/O load, engine deficit percentage and user percentage.
Load Profiles	Displays a list of load profiles with information on name, type, minimum load score, login redirection, and dynamic migration. **Metric Weight** for a selected load profile displays user connections, CPU busy, run queue length, and so on, and the corresponding weights associated with the metric.
Routes	Displays a list of routes with the name of the route, the logical cluster it is defined on, and the type of route such as alias, application, or login.

See also

Manage Clusters

Create clusters, manage load profiles, add instances to a cluster, display properties of clusters, and manage routes.

In the cluster environment, each machine is referred to as a node, and each Adaptive Server as an instance. Connected instances form a cluster, working together to manage a single set of databases on the shared disks. In each case, the instances present as a single system, with all data accessible from any instance. The Cluster Edition assigns server process IDs (spids) that are unique to the cluster, so the spid identifies a single process across all instances in the cluster, with all data accessible from any instance.

You can configure an Adaptive Server cluster to run multiple servers as a shared-disk or private cluster. You can also create a shared cluster in Veritas Cluster Server (VCS) mode.

You cannot mix different operating systems or hardware architectures. However, each node can contain a different number of processors, cores, and RAM. This difference in

configuration can be accounted for within the cluster configuration file (and subsequently within the Workload Manager) once the Cluster Edition is installed and your cluster is built.

See the *Installation Guide Adaptive Server Enterprise Cluster Edition* for your platform for information about installing the software, platform support, and system requirements.

Configuration and descriptions of the features available in the Cluster Edition, see the *Clusters Users Guide* for detailed information about the features available in the Cluster Edition.

Starting the SCC Agent for a Cluster

For Adaptive Server versions 15.7 ESD1 and later. The SCC Agent enables the distributed management of the cluster.

1. From the $SYBASE directory of each instance, source the SYBASE.sh file.
2. Start the SCC Agent:
   $SYBASE/SCC-3_2/bin/scc.sh -port portNumber &

   Where *portNumber* is any available port.

   The default port number is 9999; however this may not be the correct port number for your configuration.

See also
- *Adding an Instance to a Cluster* on page 243
- *Registering a Cluster* on page 229
- *Creating a Cluster* on page 227

Starting the Unified Agent for a Cluster

For Adaptive Server versions 15.7 and earlier. The Unified Agent enables the distributed management of the cluster.

1. From the $SYBASE directory of each instance, source the SYBASE.sh file.
2. Start the Unified Agent:
   $SYBASE/UAF-2_5/bin/uafstart-up.sh -port portNumber &

   Where *portNumber* is any available port.

   The Unified Agent creates $SYBASE/UAF-2_5/nodes/<node_name>/log/agent.log, and sends agent output to the agent.log file. The default port number is 9999; however this may not be the correct port number for your configuration.

Preparing to Create a Cluster

Prerequisites for creating a cluster.

1. Install Adaptive Server Enterprise Cluster Edition.
2. Allocate enough storage space for the required database devices. Only raw partitions rather than filesystem-based devices are supported. Ensure that the raw partitions are

accessible from each node using the same access path. These database devices are created if they do not exist:

- master device
- sybsystemprocs device
- sybsystemdb device
- quorum device (no actual database visible from within Adaptive Server is created)

3. Allocate one device for the local temporary database on each node. For example, if you are creating a three-node cluster, then you need three individual raw partitions for temporary databases—one for each node.

4. Create the appropriate number of devices, such as data and log, for your user databases. Each instance of the cluster recognizes the same thing as the other, so you need to create the devices only once. For example, if you have a four-node cluster and you are connected to instance ASE, when you run a **disk init** command for datadev1 on ASE1, the instances ASE2, ASE3, and ASE4 recognize, and can write to datadev1 once **disk init** completes.

5. Start the agent from your Adaptive Server installation directory.

Creating a Cluster

Use Sybase Control Center to create and configure a shared-disk cluster.

Prerequisites

Complete the preparation steps for creating a cluster.

Task

1. In the Perspective Resources view, select a server, then select **Resource > Administration Console**.

2. In the left pane of the Administration Console, select **ASE Servers**, click the drop-down arrow, then select **Create Cluster**.

   Optionally you can click on **ASE Servers**, then from the menu bar, click **Folder > Create Cluster**.

3. On the Introduction screen, choose one of:

   - Provide a name for the cluster, or,
   - Import an existing cluster file.

   Note: Once you have entered the information to create a cluster into the wizard, you can export the information to a file on the final Summary screen.

4. On the Options screen:
   a) Select your platform. All nodes must use the same operating system and 64-bit hardware architecture.
   b) Choose **Shared** or **Private** installation.

c) Specify the number of instances allowed. The maximum number allowed is four.

5. On the Nodes screen, click **Add** to specify the hosts that will participate in the cluster. Provide the host name, the SCC agent port number, and the SCC agent user name and password.

6. On the Quorum Device screen, enter the full path to the device. The quorum device is a communications medium and synchronization point that provides a location to perform cluster membership management. It also provides a persistent place to store configuration data used by instances and the Unified Agent.

   If a cluster agent plug-in exists, you can either accept the existing plug-in or create a new plug-in.

7. The Cluster Environment screen is specific to a shared installation. List any start-up parameters and trace flags. For private installation, this information is specified when you define instances for the cluster.

   If a Veritas Cluster Server (VCS) base device is detected on the system, you are prompted to integrate the ASE Cluster Edition Dataserver with the Veritas Cluster Server. See *Using the Cluster Edition with the Veritas Cluster Server* in the *Clusters Users Guide*.

   If an LDAP environment is detected, see *Using LDAP as a directory service* in the *Adaptive Server Enterprise Cluster Edition Clusters Users Guide*.

   If the devices you specify in the following steps do not include SCSI-3 PGR capability or support I/O fencing, you can select to continue, not continue, or to ignore further messages regarding SCSI-3 PGR capability or I/O fencing support. This dialog does not appear if a VCS device is detected.

8. On the Master Device screen, enter the full path to the device, the device size, database size, and the page size of the instance.

9. On the System Procedures Device screen, enter the full path to the device, the device size, and the database size.

10. On the System Database Device screen, enter the full path to the device, the device size, and the database size.

11. (Optional) On the PCI Device screen, click to enable the PCI device, then enter the full path to the device, the device size, and the database size.

    See *Java in Adaptive Server Enterprise* for more information about PCI devices.

12. On the Network Protocol screen, verify or update the network protocol:

    a) The UDP network protocol.

    b) Choose an unused starting port number for the cluster's network interconnects.

    c) (Optional) Click the box if you have a secondary network available.

    d) (Optional) Click **Validate Port** to make sure the ports are not currently in use.

13. On the Instance screen, click **Add** to add instances to the cluster.
    You see the **Define Server Instance** screen with the tabs **General**, **Local System Temporary Database**, and **Private** (only if a Private installation is chosen).

14. On the General tab:

a) Select the node where the instance will be located.

b) Enter the instance name.

c) Select a query port number, then select **Validate Port** to make sure this port is available.

d) Select a primary and secondary network address. The Secondary option is enabled if you specified on the Network Protocol screen that a secondary address is available.

e) Provide a full path to the error log for this instance.

f) List any start-up parameters for the instance.

g) List any trace flags to use at start-up.

15. The Local System Temporary Database tab shows default values for local temporary databases. A shared-disk installation requires a local system temporary database. Accept the defaults or change the database device values.

16. (Optional) Click **Summary** to review the information you have provided.

17. (Optional) Click **Automatically start the cluster after it is successfully created**.

18. (Optional) Click **Export** to save the cluster information to a file.

19. Click **Next** to create the cluster.

> **Note:** You cannot start this cluster from the command line using a `run_server` file. To start the cluster, the agent must be running on each node and the cluster must be registered with Sybase Control Center.

See also

Registering a Cluster
Registering a cluster allows you to manage the cluster in Sybase Control Center as a resource.

Prerequisites

The agent must be running on each node.

The cluster does not need to be running to be registered as a resource.

Task

If your cluster is registered with Sybase Control Center, it appears in the Resource Explorer and can be added to one or more perspectives. If your cluster does not appear in the Resource Explorer, you must register the cluster.

1. From the Sybase Control Center main window, select **Resource > Register**.

---

2. On the Resource Type screen, enter the name of the cluster, then select the resource type.

3. (Optional) Provide a description of the cluster.

4. On the Connection Information screen, enter the name of a host that is running one of the instances of the cluster.

5. On the ASE Port number field, enter the port number of the host from step 4.

6. (Optional) Enter a default character set.

7. (Optional) Enter a default language.

8. (Optional) On the Authentication Information screen, enter the user name and password for the administrator.

9. (Optional) On the Options screen, click to add the cluster to the list of current perspective resources.

10. (Optional) Click to open the Resource Explorer screen. This option is available only if the window is not already open.

11. On the Resource Explorer screen, highlight the cluster name, then select **Resources > Add Resources to Perspective**.

12. From the application menu bar, select **View > Open > Resource Explorer**.
From the Perspective Resource view, you can open the Administration Console to register and authenticate a running agent, then start your cluster.

See also
- *Adding an Instance to a Cluster* on page 243
- *Starting the SCC Agent for a Cluster* on page 226
- *Creating a Cluster* on page 227
- *Registering the Agent and Starting a Cluster* on page 230

Registering the Agent and Starting a Cluster
Once a cluster is registered in Sybase Control Center, you can authenticate the agent, start the cluster, and authenticate the cluster.

Prerequisites

To start the cluster, the SCC agent must be running on each node and the cluster must be registered and available in the **Perspective Resources** screen.

Task

After you register and authenticate the agent, you see an option to start the cluster. Optionally, if the cluster is authenticated, you can choose to start the cluster from the Administration Console. Starting the cluster from the Administration Console provides a log trace that shows the start-up status along with any potential warning or error messages.

1. From the Perspective Resources view, select the cluster, click the drop-down arrow that appears next to the name, then select **Administration Console**.

2. In the left pane of the Administration Console, select **ASE Servers** to populate the right pane with a list of servers.

3. Select the cluster, click the drop-down arrow that appears next to the cluster name, then select **Register Agent**.
   You see the Server Properties screen.

4. In the Agent screen, enter the port number to the agent running the node and click **Register**. The default port number is 9999; however, this may not be the correct port number for your configuration.

5. Enter the user name and password for the user authorized as the agent and click **Authenticate.**
   Once the agent is authenticated, the screen updates with the cluster status and prompts you to start the cluster.

6. Click **Start Server**.
   The cluster status changes to Running.

7. From the Perspective Resources view, select the cluster, click the drop-down arrow that appears next to the name, then select **Authenticate**.

8. Select the credential type to authenticate the cluster.

9. (Optional) If the credential type is ASE Username, enter the user name and password authorized to authenticate the Adaptive Server and click **OK**.

See also
- *Registering a Cluster* on page 229

Displaying Cluster Configuration Values
Display or modify cluster configuration values.

1. In the Administration Console, select the cluster, click the drop-down arrow, and select **Configure**.

   Note: You must have **sa_role** to select **Configure** from the Administration Console.

   You see a table with fields parameter names, value, default value, maximum value, minimum value, and restart required. Editable columns are indicated by a "pencil" icon. To reset a value to the previously configured value, click the "Reset" icon that appears after you have edited a field.

2. Select the cluster configuration parameter to configure.
   a) Click **Configure Value** for the selected row.
   b) Enter the new value for the configuration parameter.

If the value is invalid, you see an error message. As some parameters require a restart, the changed value is in the **Pending Value** column until the server is restarted.

c) Click **Save All** to update the server with the new values or **Reset All** to restore the original values for the resource.

Starting a Cluster

Starting a cluster from the Administration Console after the agent is authenticated allows you to see a log trace with any warning and error messages.

Prerequisites

The agent must be running on each node and the cluster or clusters must be registered and available in the Perspective Resources screen.

Task

If the agent is authenticated and you choose to start the cluster from the Administration Console, you see a log trace in the dialog that shows the start-up status along with any potential warning or error messages. You can optionally start a cluster immediately after you register and authenticate the agent on the Server Properties screen; however, using this option does not provide a log trace.

1. In the Perspective Resources view, select multiple resources then select **Resource > Administration Console** from the menu bar.

2. In the left pane of the Administration Console, select **ASE Servers** to populate the right pane with a list of servers.

3. Select a cluster or multiple clusters, click the drop-down arrow that appears next to the name, then select **Start Cluster**. To start multiple servers at the same time, the servers must be the same type. You can start multiple Symmetric Multiprocessors (SMP) or Shared Disk Cluster (SDC) servers.
   The Start Cluster screen appears.

4. Click **Yes**.
   Start-up messages are displayed.

Done appears in the Start Cluster window when the start-up request has finished.

See also
- *Starting an Instance* on page 243

Stopping a Cluster

Shut down a cluster immediately or choose to wait for transactions to complete.

1. In the left pane of the Administration Console, select **ASE Servers** to populate the right pane with a list of servers.

Optionally, you can click the menu next to the server name, and select **Agent > Stop Cluster**. Using this option does not provide a log trace.

2. Select a cluster or multiple clusters, click the drop-down arrow that appears next to the name, then select **Stop Cluster**. To stop multiple servers at the same time, the servers must be the same type. You can stop multiple Symmetric Multiprocessors (SMP) or Shared Disk Cluster (SDC) servers.
   The Stop Cluster screen appears.

3. Choose:

   - **Yes** to wait for all statements and transactions to complete before shutting down each instance in the order specified in the cluster configuration file.
   - **Shut down immediately** to shut down the cluster immediately, without waiting for transactions to complete.

Done appears in the Stop Cluster window when the shut down request has finished.

Dropping a Cluster
Dropping a cluster allows you to remove the cluster and associated devices, logs, and files.

Prerequisites
Shut down the cluster. The agent must be authenticated to drop the cluster.

Task

When you drop a cluster it is removed completely and no longer available for restart.

1. In the left pane of the Administration Console, select **ASE Servers** to populate the right pane with a list of servers.
2. Select the cluster, click the drop-down arrow that appears next to the name, then select **Drop**.
   You see the Confirm Delete screen.
3. Select the devices, logs, and files associated with the selected cluster that you want to remove and click **OK**.

   Note: Due to certain file-system locking, the UAF plugin directory may not be deleted after you drop a cluster. Verify that the `$SYBASE/UAF-2_5/nodes/*/plugins/<cluster_name>` directory has been deleted. If the directory still exists, shut down the UAF agent, manually delete the directory, then restart the UAF agent.

Viewing a Cluster Log
View information about the success or failure of each start-up attempt of the cluster and operations of the cluster.

Prerequisites
The agent must be authenticated in order to view a cluster log.

Task

1. In the left pane of the Administration Console, select **ASE Servers** to populate the right pane with a list of servers.

2. Select the cluster, click the drop-down arrow that appears next to the name, then select **View Log**.
   You see the Server Log screen.

3. (Optional) Enter search criteria in the "Show messages matching" field to search for specific errors or messages.

4. (Optional) Click **Settings** to limit the number of messages to retrieve and select a regular expression to include in the retrieval process.

5. (Optional) Click **Options** to add or remove information columns, to copy a specific row, or copy the entire table.

Cluster Properties
Use the Properties wizard to view cluster and agent information, to configure the agent, and to stop or restart the cluster.

To initiate the Properties wizard, select the cluster from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	Server Properties
General	View cluster information such as: • Name and type • Version, build date, and build options • Edition – click **Cluster Edition** for licensing information. • Character set, sort order, and language • Host name, port number, platform and operating system (if the agent has been authenticated) • Page size • Status; running or stopped
Agent	View agent properties such as port number, user name, and version information. You can also authenticate or clear the authentication of your agent. View server properties such as version, server name, status, and start-up file. You can also start or stop your server using the agent.

Wizard Option	Server Properties
Cluster	View cluster configuration information such as: • Maximum instances • Installation mode – shared or private • Membership mode – specifies whether the cluster membership service functions autonomously in native mode or in cooperation with Veritas Cluster Server (VCS) mode • Quorum device – location of quorum device • Master device – location of master device • List instances, the instance address, and query port **Note:** If the agent is not authenticated, you cannot retrieve configuration information and a warning message is displayed.

Manage Logical Clusters

A logical cluster is a grouping of Adaptive Server instances operating on one or more nodes in the cluster. You can manage workloads running on clusters with separate routing, load balancing, and failover rules.

Adding a Logical Cluster

The logical cluster is an abstract representation of one or more instances in a physical shared-disk cluster used to manage workload, failover, and client application access to the cluster.

Each logical cluster has a set of instances on which it runs, and can have a set of instances to which it fails over. Routing rules direct incoming connections to specific logical clusters based on an application, user login, or server alias supplied by the client. Other rules can restrict the logical cluster to a bound connection or allow access by any authenticated connection.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Logical Cluster**.
2. Click the drop-down arrow and select **New**.
   You see the Add Logical Cluster wizard.
3. Select the shared cluster to which to add the logical cluster.
4. Assign a name to the logical cluster that is representative of the job it will perform, such as SalesLC.
5. Click **Add** to select the instances you want to add to the logical cluster. These instances are considered "base" instances.
   You can remove an instance from this list by clicking **Remove**.

6. In the Failovers screen, select **Add** to choose the failover server instances on which the logical cluster will run if one or more of the base instances fail. Any instance in the physical cluster can be a failover resource.

7. Select the group number for the failover instances. The failover group lets you specify preference and order in which failover instances are used. Lower-numbered groups are used first.

8. Click **Add** to add routed applications, logins, or aliases.

   Routing rules allow you to specify specific applications, logins, and aliases to which the logical cluster will connect. See the *Cluster Users Guide* for more information about routing. Add additional routes by selecting:

   - **Application Route** – specify the name of the application in the Add Application Route screen.
   - **Alias Route** – specify the alias name in the Add Alias Route screen.
   - **Login Route** – select the login name from the list in the Add Login Route screen (hold the Ctrl key to select multiple names).

   You can drop a listed route by selecting the route name and clicking **Remove**.

9. On the Load Profile screen, click **Change** to select a load profile to assign to the logical cluster.

   The Cluster Edition offers two system load profiles: `sybase_profile_oltp` for OLTP environments and `sybase_profile_dss` for DSS environments. You cannot modify or delete system load profiles.

10. On the Set Options page:

    a) Select **Cluster** or **Instance** for the system view.

       The system view determines whether the logical cluster users view the Cluster Edition as an entire cluster or as individual instances. This affects some queries and stored procedures.

    b) (Optional) Select **Automatically start logical cluster** to start the logical cluster when the shared cluster is started.

    c) Choose **instance** or **group** to specify whether the instances are brought online as a group or individually.

    d) (Optional) Select **Fail to any** to specify whether any instance, or a specific instance can be a failover resource.

    e) Choose the down-routing mode to specify how client connections are routed if the logical cluster designated by the routing rule is not available.

       - **system** – sends unroutable connections to the system logical cluster.
       - **open** – sends unroutable connections to the open logical cluster.
       - **disconnect** – disconnects unroutable connections.

    f) Choose a login distribution mode to specify how connections are distributed in logical clusters with multiple instances.

       Note: The login distribution mode does not affect single-instance logical clusters.

- **affinity** – specifies that a connection accepted by an instance is retained as long as the target logical cluster is running on that instance. If the load profile specifies a load threshold, and the load on the instance is too high, the workload manager redirects the connection to the least loaded instance in a logical cluster. If the target logical cluster is not running on the instance, the workload manager redirects the connection to the least loaded instance in a logical cluster.
- **round-robin** – specifies that incoming connections are distributed in a round-robin fashion among the instances hosting a logical cluster. For example, if "SalesLC" is running on "ase1" and "ase2", the workload manager sends the first connection to "ase1," the next to "ase2", and so on. Load scores are not included in the algorithm.

g) (Optional) Select **Open role** to indicate that all connections that are not routed to a logical cluster via an explicit routing rule are routed to the current open logical cluster. When you create a new cluster, the system logical cluster is automatically designated the open logical cluster. You can reassign the open role to another logical cluster. However, only one open logical cluster can exist per physical cluster.

11. (Optional) Click **Summary** to review your information.

12. Select **Finish** to build the logical cluster.

See also
- *Changing the State of a Logical Cluster* on page 238
- *Logical Cluster States* on page 238
- *Logical Cluster Properties* on page 240

Designating a Failover Logical Cluster
Designate the logical cluster as the failover.

Prerequisites
The logical cluster must be online to designate a failover cluster.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Logical Cluster**.

2. Select the logical cluster, click the drop-down arrow that appears next to the name, then select **Failover**.
The Confirm Failover screen appears.

3. Confirm the failover.

Designating a Failback Logical Cluster
Reverse a failover of a logical cluster by reinstating the original base resources.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Logical Cluster**.

2. Select the logical cluster, click the drop-down arrow that appears next to the name, then select **Failback**.
   The Confirm Failback screen appears.

3. Confirm the failback.

Logical Cluster States

A logical cluster and each instance associated with the cluster can have different states.

A logical cluster has a global state that determines whether the cluster is offline, online, or inactive. Logical clusters also have an instance state that describes the state of a particular instance associated with the logical cluster. For example, an online logical cluster may be online on its base instances and offline on its failover instances. This state may be independent of the actual Adaptive Server state, as a logical cluster may be offline on an instance that is actually up and running.

State	Global State	Instance State
online	A logical cluster is online and running on one or more instances.	The online logical cluster is accepting and managing connections on the current instance.
offline	A logical cluster is not running on any instance.	A logical cluster is not running on the current instance, which cannot accept connections or consume resources.
inactive	Similarly to the offline state, a logical cluster is not running on any instance. Inactive logical clusters are not started automatically and do not participate in failover. The cluster achieves the inactive state only through the deactivate command. Once inactive, the cluster comes online only through the online command.	Not applicable to instances.

See also
- *Logical Cluster Properties* on page 240
- *Adding a Logical Cluster* on page 235
- *Changing the State of a Logical Cluster* on page 238

Changing the State of a Logical Cluster

Change the state of a logical cluster or instances associated with the cluster to offline, online, or inactivate.

1. To change the state of a logical cluster or an instance, in the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Logical Clusters**.

   - To change the state of a cluster, or the global state, select the cluster, click the drop-down arrow that appears next to the name, then select the state: Offline, Deactivate, or Online.
   - To change the state of an instance, click the drop-down arrow that appears next to the name, then select **Properties > Base Instances**. Click the drop-down arrow next to the instance name and select the state: Online or Offline.

2. (Optional) When changing the state to **Offline**, select whether to take the cluster or instance offline immediately or whether to take it offline gracefully, choosing the amount of time for a reminder.

See also
- *Adding a Logical Cluster* on page 235
- *Logical Cluster States* on page 238
- *Logical Cluster Properties* on page 240

Deleting a Logical Cluster
Remove an offline logical cluster.

Prerequisites
A logical cluster must be offline before it can be deleted.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Logical Cluster**.
2. Select the cluster, click the drop-down arrow that appears next to the name, then select **Delete**.
3. Confirm the deletion.

Generating a DDL Script for a Logical Cluster
Generate DDL to create a logical cluster.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management** and select **Logical Clusters**.
2. Select the cluster, click the drop-down arrow that appears next to the name, then select **Generate DDL**.

   You can save the DDL in an external file on your local file system.

Logical Cluster Properties

Use the Properties wizard to view logical cluster information.

To initiate the Properties wizard, select the logical cluster from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	Logical Cluster Properties

General	View cluster information such as:
	
	• System view – determines how the logical cluster users view Cluster Edition, whether as an entire cluster, or as individual instances.
	• Automatically start logical cluster – determines whether you want the logical cluster started when the cluster starts.
	• Failover mode – select either instance or group to specify how to bring instances online as a group or individually.
	• Fail to any – specifies whether any instance, or only a specific instance, can be a failover resource.
	• Down routing mode – specifies how client connections are routed if the logical cluster designated by the routing rule is unavailable.
	   • system – sends unroutable connections to the system logical cluster.
	   • open – sends unroutable connections to the open logical cluster.
	   • disconnect – disconnects unroutable connections.
	• Login distribution mode – lets you specify how connections are distributed in logical clusters with multiple instances. The login distribution mode does not affect single-instance logical clusters.
	   • affinity – specifies that the instance accepting a connection retains it as long as the target logical cluster is running on that instance. If the load profile specifies a load threshold, and the load on the instance is too high, the workload manager redirects the connection to the least loaded instance in a logical cluster. If the target logical cluster is not running on the instance, the workload manager redirects the connection to the least loaded instance in a logical cluster.
	   • round-robin – specifies that incoming connections are distributed in a round-robin fashion among the instances hosting a logical cluster. For example, if "SalesLC" is running on "ase1" and "ase2", the workload manager sends the first connection to "ase1," the next to "ase2", and so on. Load scores are not included in the algorithm.
	• Logical cluster roles – indicates whether this logical cluster assumes an open role, meaning that all connections that are not routed to a logical cluster via an explicit routing rule are routed to the current open logical cluster. When you create a new cluster, the system logical cluster is automatically designated the open logical cluster. You can reassign the open role to another logical cluster. However, only one open logical cluster can exist per physical cluster.

Base Instances	List base instance names, their identification numbers, and their current state.

Wizard Op-tion	Logical Cluster Properties
Failover In-stances	List failover instance names, their identification numbers, their current state, and the failover to which they belong.
Routes	List routing applications, logins, and aliases that are associated with the logical cluster.
Load Profile	Change the load profile assigned to the logical cluster or view the current load profile information such as: • Metric weights assigned to the selected load profile • Minimum load score • Login redirection • Dynamic connection migration

See also

Manage Cluster Instances
Add instances to a cluster or change the status of current cluster instances.

Displaying Instance Configuration Values
Display or modify instance configuration values.

1. In the Administration Console, select the instance, click the drop-down arrow, and select **Configure**.

   Note: You must have **sa_role** to select **Configure** from the Administration Console.

   You see a table with fields parameter names, value, default value, maximum value, minimum value, and restart required. Editable columns are indicated by a "pencil" icon. To reset a value to the previously configured value, click the "Reset" icon that appears after you have edited a field.

2. Select the cluster configuration parameter to configure.
   a) Click **Configure Value** for the selected row.
   b) Enter the new value for the configuration parameter.
      If the value is invalid, you see an error message. As some parameters require a restart, the changed value is in the **Pending Value** column until the server is restarted.

c) Click **Save All** to update the server with the new values or **Reset All** to restore the original values for the resource.

See also
*   *Adding an Instance to a Cluster* on page 243
*   *Instance Properties* on page 246

Starting an Instance
Start an Adaptive Server instance connected to a cluster.

Prerequisites
The agent must be authenticated in order to add an instance.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Instances**.
2. Click the drop-down arrow that appears next to the instance name, then select **Start Server**.
    The Start Server screen appears.
3. Click **Yes** to confirm your request.
    Done appears in the Start Server window when the start server request has finished.

See also
*   *Starting a Cluster* on page 232

Adding an Instance to a Cluster
Add additional instances to an existing cluster.

Prerequisites
The maximum number of instances for the selected cluster must be set large enough to allow additional instances. Also, the agent must be running on the node on which you are adding the instance.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Instances**.
2. Click the drop-down arrow and select **New**.
    You see the Add Cluster Server Instance wizard.

3. On the Introduction screen, select the cluster to which to add an instance.

4. (Optional) On the Cluster information screen, you can import an existing instance file.

5. On the Nodes screen, click **Add** to specify the hosts that participate in the cluster. Provide the host name, the agent port number, and the agent user name and password.

6. The Quorum Device screen provides information about the designated quorum device. Once the device is accepted, the cluster agent plug-in is deployed to the agent host. If the duplicate agent plug-in already exists, you to create a new plug-in or accept the existing plug-in.

7. The Cluster Environment screen is specific to a shared installation. For a private installation, this information is specified when you are defining instances for the cluster.

   If a Veritas Cluster Server (VCS) base device is detected on the system, you are prompted to integrate the ASE Cluster Edition Dataserver with the Veritas Cluster Server. See *Using the Cluster Edition with the Veritas Cluster Server* in the *Clusters Users Guide*.

   If an LDAP environment is detected, see *Using LDAP as a directory service* in the *Clusters Users Guide*.

   If the devices you specify in the following steps do not include SCSI-3 PGR capability or support I/O fencing, you are prompted to continue, not continue, or to ignore further messages regarding SCSI-3 PGR capability or I/O fencing support. This dialog does not appear if a VCS device is detected.

8. Click **Add** to add instances to the cluster.
   You see the Define Server Instance screen with the tabs General, Local System Temporary Database, and Private (if a private installation is chosen).

9. On the General tab:
   a) Select the node where the instance will be located.
   b) Enter the instance name.
   c) Select a query port number, then select **Validate Port** to make sure this port is available.
   d) Select a primary and secondary network address. The Secondary option is enabled if you specified on the Network Protocol screen that a secondary address is available.
   e) Provide a full path to the error log for this instance.
   f) List any start-up parameters for the instance.
   g) List any trace flags to use at start-up.

10. On the Local System Temporary Database tab default values are provided for the local temporary databases. A local system temporary database must be create for a shared disk installation. Accept the defaults or change the database device values.

11. (Optional) On the Private tab:
    a) Sybase home – the Adaptive Server release directory (value of `$SYBASE`).
    b) Adaptive Server Enterprise home – the full path to the `ASE-15_0` directory (value of `$SYBASE/$SYBASE_ASE`).
    c) Interfaces directory – path to the directory containing the interfaces file.

    d) Sybase environment script – full path to the script that you use to set the environment variables.

    e) Cluster configuration file – full path to the cluster configuration file.

12. (Optional) Click **Summary** to review the information you have provided.

13. (Optional) Click **Export** to save the instance information to a file.

14. Click **Next** to add the instance to the selected cluster.

> **Note:** After adding instances, start each one separately.

See also
- *Registering a Cluster* on page 229
- *Starting the SCC Agent for a Cluster* on page 226
- *Creating a Cluster* on page 227
- *Displaying Instance Configuration Values* on page 242
- *Instance Properties* on page 246

Stopping an Instance
Shut down an instance immediately or wait for transactions to complete.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Instances**.

2. Click the drop-down arrow that appears next to the instance name, then select **Start Server**.
The Start Server screen appears.

3. Choose:

- **Yes** to wait for all statements and transactions to complete before shutting down.
- **Shut down immediately** to shut down the cluster immediately, without waiting for transactions to complete.

Done appears in the Stop Server window when the shut down request has finished.

Deleting an Instance
Remove an offline cluster instance.

Prerequisites
The agent must be authenticated and the cluster instance must be offline before it can be deleted.

Task

---

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Instances**.
2. Click the drop-down arrow that appears next to the instance name, then select **Delete**.
3. Confirm the deletion.

Instance Properties
Use the Properties wizard to view instance information.

To initiate the Properties wizard, select the instance from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	Server Properties
General	View instance information such as: • Name and type • Version, build date, and build options • Edition – click **Cluster Edition** for licensing information. • Character set, sort order, and language • Host name, port number, platform and operating system (shown if agent is authenticated) • Page size • Status; running or stopped • Location of Sybase home directory

See also

Viewing an Instance Log
View information about the success or failure of each start-up attempt of the cluster and operations of the cluster.

Prerequisites
The agent must be authenticated in order to view an instance log.

Task

1. In the left pane of the Administration Console, select **ASE Servers > Cluster Management > Instances** to populate the right pane with a list of instances.
2. Select the instance, click the drop-down arrow that appears next to the name, then select **View Log**.

You see the Server Log screen.

3. (Optional) Enter search criteria in the "Show messages matching" field to search for specific errors or messages.

4. (Optional) Click **Settings** to limit the number of messages to retrieve and select a regular expression to include in the retrieval process.

5. (Optional) Click **Options** to add or remove information columns, to copy a specific row, or copy the entire table.

Manage Load Profiles

Load profiles let you define the operating criteria for a logical cluster.

Load profile criteria are typically called "load score metrics," with the value associated for each criteria rolled into a "score" for each instance in the logical cluster that uses the load profile. You can periodically compare load scores for different instances within a logical cluster to detect when the workload is undesirably skewed to one or more instances, or to determine if an instance is underutilized.

Instances included in multiple logical clusters can be impacted by multiple load profiles, so take care when associating instances with multiple logical clusters and when defining and applying load profiles.

The Cluster Edition includes two system load profiles: `sybase_profile_oltp` for OLTP environments and `sybase_profile_dss` for DSS environments. You cannot modify or delete system load profiles. However, you can duplicate them and modify the duplicates to create your own load profiles.

The load profile status reports the:

- Name – the name for the load profile configuration.
- Type – the load profile type: system or user.
- Metric weights – relative weight assigned to each metric in the load profile. Metrics include:
  - User connections – the weight of users connected to the particular load profile.
  - CPU busy – the weight of CPUs that are currently busy.
  - Run queue length –the weight of the run queue.
  - IO load – the weight of the I/O load.
  - Engine deficit – the weight of the engine deficit.
  - Users – the weight for a metric the user chooses to measure.
- Thresholds – configured difference (as a percentage) in the load between two instances in a logical cluster at which point the following can occur:
  - Login redirection – used for connection-time load balancing and routing connections to a logical cluster. If necessary, an instance directs a client to stop the current login attempt and try connecting to instances it supplies as a list of available network addresses.

- Dynamic migration – (also known as the hysteresis value) the dynamic migration configuration.
- Minimum load score – load score necessary to trigger login redirection and dynamic migration.

Adding a Load Profile

Use a load profile to consolidate all configurable aspects of the workload scoring system into a single named entity.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Load Profile**.
2. Click the drop-down arrow and select **New**.
   You see the Add Load Profile wizard.
3. Select a cluster for the load profile.
4. Enter a name for your profile.
5. Specify the metric weights for the load profile.

   When a load profile is associated with a logical cluster, the workload manager calculates a load score for each instance in the logical cluster. This is calculated using the weight you entered for each metric, the raw value of each metric for the instance, and the workload measurement algorithm.

   The metrics measured by the server are:
   - User connections – the capacity of an instance to accept a new connection, based on resource availability.
   - CPU busy – the capacity of an instance to accept additional work.
   - Run-queue length – the number of runnable tasks on a system. The run-queue length measures the processing backlog, and is a good indicator of relative response time.
   - I/O load – outstanding asynchronous I/Os.
   - Engine deficit – the difference in the number of online engines among instances in the cluster. Engine deficit is measurable only when instances in the cluster have unequal numbers of engines. Engine deficit adds a metric for maximum relative capacity to the load score.
   - User-defined – an optional customer-supplied metric specific to the user's environment.
6. Enter values for:
   - Minimum Load Score – the load score is not a percentage, but is the minimum score the workload manager requires before it redirects work to other instances. The minimum load score is meaningful when compared to the load scores of other instances in the logical cluster using the load profile.
   - Login Redirection (%) – the load threshold for determining how best to distributed incoming connections.
   - Dynamic Connection Migration (%) – the load threshold that determines whether to distribute existing connections. The load threshold is the percentage difference

between the load on the current instance and the load on the least-loaded instance currently participating in a logical cluster. That value must be met before the Cluster Edition redirects a login or migrates an existing connection.

The percentages for Login Redirection and Dynamic Connection Migration are independent, and do not need to add up 100.

7. (Optional) Click **Summary** to verify your selected options.

See also
- *Load Profile Properties* on page 249

Deleting a Load Profile
Delete a user created load profile.

You cannot modify or delete the system load profiles, `sybase_profile_oltp` and `sybase_profile_dss`.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management** and select **Load Profile**.
2. Click the drop-down arrow that appears next to the load profile name, then select **Delete**.
3. Confirm the deletion.

Generating a DDL Script for a Load Profile
Generate DDL to create a load profile.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management** and select **Load Profile**.
2. Select the load profile, click the drop-down arrow that appears next to the name, then select **Generate DDL**.

   You can save the DDL in an external file on your local file system.

Load Profile Properties
Use the Properties wizard to view load profile information.

To initiate the Properties wizard, select the load profile from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	Load Profile Properties
General	View the name and type of load profile.
Metric Weights	Change the metric weights currently assigned to the selected load profile.

Wizard Option	Load Profile Properties
Thresholds	Change the thresholds currently assigned to the selected load profile.

See also
*   *Adding a Load Profile* on page 248

Manage Routes
Routes allow you to direct client connections to specific logical clusters.

Adding a Route
Assign routes to a logical cluster for applications, logins, and aliases to manage the cluster workload.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management**, then select **Routes**.
2. Click the drop-down arrow and select **New**.
   You see the Add Route wizard.
3. Select a cluster from which the incoming client connections will be routed.
4. Specify the type of route:

   *   Alias – associate a server name with a logical cluster. Applications can choose a logical cluster from server aliases placed in the interfaces file. These aliases use unique server names. The alias entry can point anywhere in the physical cluster. The workload manager sends it to the correct instances via login redirection
   *   Application – establish a relationship between an application name and a logical cluster.
   *   Login – establish a relationship between an Adaptive Server login and a logical cluster.
5. Enter a name for your route.
6. Specify the name of the logical cluster to which this route is associated.
7. (Optional) Click **Summary** to review your information.

See also
*   *Route Properties* on page 251

Deleting a Route
Close a client connection to a logical cluster.

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management** and select **Routes**.

2. Click the drop-down arrow that appears next to the route name, then select **Delete**.

3. Confirm the deletion.

See also
- *Route Properties* on page 251

Route Properties
Use the Properties wizard to view route information.

To initiate the Properties wizard, select the route from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	Route Properties
General	View the name, route type, and the logical cluster to which the incoming client connections are routed.

See also
- *Adding a Route* on page 250
- *Deleting a Route* on page 250

Manage Auxiliary Servers
Configure backup servers and XP backup servers for clusters.

Configuring a Backup Server
Configure a backup server for an entire cluster or multiple backup servers to run on a per-instance basis.

Prerequisites
The agent must be authenticated in order to configure a backup server.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management > Auxiliary Servers**.

2. Select **Backup Servers**.

3. Click the drop-down arrow, then select **Configure Backup Server**.
You see the Configure Backup Server wizard.

4. Select the cluster for which to configure a backup server.

5. Specify the type of configuration:

- Configure a backup server for the entire cluster.

  The Specify Properties screen shows a default backup server name, error log path, and port number. Optionally adjust any of these values for your configuration environment.

- Configure multiple backup servers, that run on a per-instance basis.

  When configuring a backup server on a per-instance basis:

a) Choose either a dedicated server policy or a round-robin policy for the global server configuration.

b) (Optional) Select **Enable global high availability mode** for backup servers configured on a per-instance basis.

c) The Specify Properties screen shows a default error log path and port number for each instance. Optionally adjust either of these values for your configuration environment.

6. (Optional) Click **Summary** to review your information.

See also

- *Configuring an XP Server* on page 253
- *Backup Server Properties* on page 252
- *XP Server Properties* on page 254

Deleting a Backup Server

Delete a cluster backup server.

Prerequisites

The agent must be authenticated in order to delete a backup server.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management > Auxiliary Servers**.

2. Select **Backup Servers**.

3. Click the drop-down arrow that appears next to the backup server name, then select **Delete**.

4. Confirm the deletion.

Backup Server Properties

Use the Properties wizard to view cluster backup server information.

To initiate the Properties wizard, select the backup server from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	Backup Server Properties
General	• Name and type. • For a backup server configured for the entire cluster – host and associated port numbers. • For multiple server configurations – instances and associated port numbers.
Global Configuration **Note:** Available only for multiple backup server configurations.	• Type of server policy, either dedicated or round-robin. • Global high-availability mode.

See also
- *XP Server Properties* on page 254
- *Configuring a Backup Server* on page 251
- *Configuring an XP Server* on page 253

Configuring an XP Server
Configure a backup XP server for a cluster.

Prerequisites
The agent must be authenticated in order to configure an XP server.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management > Auxiliary Servers**.
2. Select **XP Servers**.
3. Click the drop-down arrow, then select **Configure XP Server**.
   You see the Configure XP Server wizard.
4. Select the cluster for which to configure an XP server.
5. The Specify Port Numbers screen provides a default port number. Optionally adjust the port number for your configuration environment.
6. (Optional) Click **Summary** to review your information.

See also
- *Configuring a Backup Server* on page 251
- *Backup Server Properties* on page 252
- *XP Server Properties* on page 254

Deleting an XP Server
Delete a cluster backup XP server.

Prerequisites
The agent must be authenticated in order to delete an XP server.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Cluster Management > Auxiliary Servers**.
2. Select **XP Servers**.
3. Click the drop-down arrow that appears next to the XP server name, then select **Delete**.
4. Confirm the deletion.

XP Server Properties
Use the Properties wizard to view cluster backup XP server information.

To initiate the Properties wizard, select the XP server from the Administration Console, click the drop-down arrow that appears next to the name, then select **Properties**.

Wizard Option	XP Server Properties
General	• Name and type • Instances and associated port numbers

See also
- *Backup Server Properties* on page 252
- *Configuring a Backup Server* on page 251
- *Configuring an XP Server* on page 253

Databases

Monitor Adaptive Server databases.

Monitor Databases
From the Perspective Resources view, use the Monitor option on your Adaptive Server to display database statistics and monitor database activity.

Determining the Backup Status of a Database

Find out when a database's last backup started, whether the last backup failed, whether a backup is currently in progress, and more.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Databases**.

   You can also display the Databases window by clicking a **Databases** link on another window in the Adaptive Server monitor.

3. Locate your database in the Databases table.

   The table shows:
   - The date and time at which the last backup started
   - Whether a backup is currently in progress
   - Whether the last backup failed
   - Whether the transaction log is full
   - Whether there are suspended processes associated with this database

For more information on backups, see the chapters on developing a recovery plan and on backing up and restoring user databases in the Adaptive Server *System Administration Guide*, Volume 2.

See also
- *Displaying Resources Used by a Database* on page 255
- *Displaying Information About Segments Used by a Database* on page 258
- *Database Statistics and Details* on page 257

Displaying Resources Used by a Database

View disk usage, running processes, unused indexes, and devices and segments associated with an Adaptive Server database.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Databases**.

   You can also display the Databases screen by clicking a **Databases** link from a different window in the Adaptive Server monitor.

3. Select a database in the Databases table.

   The tabs at the bottom of the screen are populated with information on the database you selected.

> **Note:** When you select a database, Sybase Control Center calculates space usage before displaying any data, and this may take 30 seconds or more for a large database.

4. Click a tab to see information about the resource you are interested in:

   - **Details** – shows disk usage, including the size of reserved and unreserved data and log segments.
   - **Running Processes** – shows the server process ID (spid), login, host, command, and transaction for each process that is currently using the database. Each spid number is a link; click it to see more information about that process.
   - **Devices Used** – shows the size and usage allocation (data or log) for each device that provides storage for this database. Each device name is a link; click it to see more information about that device.
   - **Segments Used** – shows the size and unused portion of each segment, in megabytes. Each segment name is a link; click it to see more information about that segment.
   - **Unused Indexes** – shows the name and table of each unused index.
   - **Frequently Used Tables** – shows usage statistics for frequently used tables.

See also
- *Determining the Backup Status of a Database* on page 255
- *Displaying Information About Segments Used by a Database* on page 258
- *Database Statistics and Details* on page 257

Modifying Database Sizes

Sybase Control Center allows you to increase the size of your databases from the Database window.

1. In the Adaptive Server monitor, select **Databases**.
2. From the Databases table, select the database to configure.
3. Right-click a row and select **Extend**.
4. Select the **Extend** menu item from the context menu.
   You see the **Extend Database Size** wizard .
5. (Optional) Select a **Device Name** on which to extend the database.
6. (Optional) Specify the amount of space you want to allocate to the log and data segments.
7. Click **OK**.
   The dialog box closes if the operation succeeds, else Sybase Control Center displays an error and the dialog box stays open.

Database Statistics and Details

The Databases window shows a variety of detailed statistics, including the status, of active Adaptive Server databases.

Note: In Adaptive Server cluster configurations, the Global Databases table provides information about global databases, and the Local Databases tables provides information about local, temporary databases, grouped by cluster instance.

The Databases table lists the databases in the current Adaptive Server by name. If a database is unavailable, for example, because it is quiesced or is offline, the Name column includes the reason.

For Adaptive Server 15.5 and later, the Databases table lists the type, durability, and DML logging status for each database. The Databases table also includes for each database, the ID, and current status information, backup status, whether the transaction log is full, and whether there are suspended processes. Processes may be suspended when the transaction log fills up.

The type of database is indicated for temporary, in-memory, proxy and archive databases, and left blank for all other databases. The Durability column indicates if a database is recoverable.

The tabs at the bottom of the screen display information about the selected database. For a large database, it might take 30 seconds or more for the information to appear.

Table 33. Tabs in the Databases View

Details	Displays information about space usage, with pie charts for data segments and log segments. If this database does not have a log segment, the pie chart on the right shows combined data and log segment usage.
Running Processes	Displays information about processes that are currently using this database, including the process ID, login, host, command, and transaction name. Click a process ID in the SPID column to switch to the **Processes** view's information about that process.
Devices Used	Displays information about devices on which this database stores its data, including the device name, the amount of space used on that device, and the usage allocation (data or log). Click a device in the Name column to switch to the **Devices** view's information about that device.
Segments Used	Displays information about segments used by this database, including the segment name, the size of the segment in megabytes, and the amount of free space in the segment. Click a segment in the Name column to switch to the **Segments** view's information about that segment.

Unused Indexes	Lists indexes in this database that have not been used since the Adaptive Server was last restarted.
Frequently Used Tables	Displays information about tables in this database that have been used since the Adaptive Server was last restarted, including the table name, index ID, logical and physical reads, lock requests and waits, and contention statistics.

The Local Temporary Databases screen provides information about local, temporary databases grouped by cluster instances. For each cluster instance, Sybase Control Center lists the ID and current status information, including backup status; whether the transaction log is full, and whether there are suspended processes.

The tabs at the bottom of the Local Temporary Databases screen display information about the selected database. Click the Temporary Database Activity tab to see information about log requests and wait status, device read and write values, and charts depicting I/O activity in the selected temporary database.

See also
- *Determining the Backup Status of a Database* on page 255
- *Displaying Resources Used by a Database* on page 255
- *Displaying Information About Segments Used by a Database* on page 258

Displaying Information About Segments Used by a Database
For a database, get details about space usage, devices that make up each segment, and tables and indexes that use it.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Databases**.
3. Select a database in the Databases table at the top of the window.

   The tabs at the bottom of the window are populated with information on the database you selected.
4. Click the **Segments Used** tab.
5. Each segment name is a link; click it to see more information about that segment.
   The Segments window appears. In the Segments window, the tabs at the bottom of the screen are populated with information about the selected segment.
6. Click the tabs to see information about space usage on the segment, devices that make up the segment, and tables and indexes that are allocated on the segment.
7. If the database uses more than one segment, return to the Databases screen to identify and click through to the remaining segments.

For more information on segments, see the Adaptive Server *System Administration Guide*, Volume 2.

See also
- *Determining the Backup Status of a Database* on page 255
- *Displaying Resources Used by a Database* on page 255
- *Database Statistics and Details* on page 257
- *Determining the Space Used by a Table on a Segment* on page 361
- *Extending a Segment in Adaptive Server* on page 361
- *Segment Statistics and Details* on page 362

Manage Databases

From the Perspective Resources view, use the Administration Console option on your Adaptive Server to create databases, modify their properties, and perform other administrative tasks.

Creating a User Database

Create a user database using the Sybase Control Center Administration Console.

Prerequisites

Consider these database attributes:

- Size :
  - **sp_estspace** helps you estimate table and index space requirements based on the definition of a specific table. See the *Adaptive Server Reference Manual* .
  - Space for planned views, stored procedures, defaults, rules, and triggers
  - Size of the transaction log
  - Space for anticipated expansion
- Database device location, and whether there is enough space on that device.
- Transaction log location – Sybase recommends that you store the transaction log on a different device than the data.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **User Databases**.
3. Click the drop-down arrow and select **New**.
   You see the Add User Database wizard.
4. On the Introduction screen, select the server in which to create a database.

5. On the Database screen, enter the name of the database you want to create.

6. (Optional) On the Devices screen, enter the size of the new database. If you do not enter a size, the default size allocated is 3MB. You can specify separately the amount of space to allocate to the log and data segments.

7. (Optional) On the Options screen, select:

   - **With override** to store the log and data on the same logical device - this is not recommended by Sybase.
   - **For load** if you want the database to be used for loading a database dump.

8. (Optional) On the Durability Level screen, choose one of these levels to increase the performance of the server by reducing the recoverability in case of a system crash:

   - NO_RECOVERY – there is no guarantee that, at runtime, committed transactions are written to the disk.
   - AT_SHUTDOWN – all committed transactions are written to disk during a normal server shutdown.
   - FULL – a complete recovery of committed transactions is possible after a crash.

   Note: These options only apply to nonclustered Adaptive Server version 15.5 and later.

9. (Optional) On the Data Compression screen, select:

   - Data compression for the entire database – choose either page-level or row-level compression. If you choose neither option, then data is not compressed.
   - LOB compression – choose from levels 0 – 9, 100, or 101.
   - In-row LOB length – choose the length of the LOB column to be saved in-row. To disallow in-row LOB storage in the database, set the length to 0.

   Note: These options only apply to nonclustered Adaptive Server version 15.7 and later.

10. (Optional) On the Guest User screen, select **Create guest user** to create a guest user who can access the database with limited privileges.

11. (Optional) Click **Summary** to see the database options you have selected.

See also

Creating a Temporary Database

Create a temporary database using the Sybase Control Center Administration Console.

Prerequisites

Consider these database attributes:

- Size :
    - **sp_estspace** helps you estimate table and index space requirements based on the definition of a specific table. See the *Adaptive Server Reference Manual* .
    - Space for planned views, stored procedures, defaults, rules, and triggers
    - Size of the transaction log
    - Space for anticipated expansion
- Database device location, and whether there is enough space on that device.
- Transaction log location – Sybase recommends that you store the transaction log on a different device than the data.

Task

Sybase Control Center provides a wizard to create a temporary database.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **Temporary Databases**.
3. Click the drop-down arrow and select **New**.
   You see the Add Temporary Database Wizard.
4. On the Introduction screen, select the server in which to create a database.
5. (Optional) For a cluster environment, click **Create this as a local temporary database**, then select the online instance on which to create the database.

   For the Cluster Edition, you can create local user temporary databases on private devices, but you must create local system temporary databases on shared devices. Local file system devices (block devices) can be used for managing the storage needs of temporary data in the cluster. These devices are added as private devices and can only be used by local user temporary databases.
6. On the Database screen, enter the name of the database you want to create.
7. (Optional) For a cluster environment, on the Temporary Database Group screen, select the group in which to add the local temporary database.

   You can only add a local temporary database to a temporary database group. Global temporary databases cannot be added to a temporary database group.

8. (Optional) On the Devices screen, enter the size of the new database. If you do not enter a size, the default size allocated is 3MB. You can specify separately the amount of space to allocate to the log and data segments.

9. (Optional) On the Options screen, select:

   - **With override** to store the log and data on the same logical device - this is not recommended by Sybase.

10. (Optional) On the Data Compression screen, select:

    - Data compression for the entire database – choose either page-level or row-level compression. If you choose neither option, then data is not compressed.
    - LOB compression – choose from levels 0 – 9, 100, or 101.
    - In-row LOB length – choose the length of the LOB column to be saved in-row. To disallow in-row LOB storage in the database, set the length to 0.

11. (Optional) On the Temporary Database Group screen, select the database group that the temporary database belongs to.

12. (Optional) Click **Summary** to see the database options you have selected.

See also
- *Creating a User Database* on page 259
- *Creating a Proxy Database* on page 262
- *Creating an Archive Database* on page 263
- *Creating an In-Memory Database* on page 264
- *Creating an In-Memory Temporary Database* on page 266
- *Creating a Temporary Database Group* on page 267
- *Mounting an Adaptive Server Database* on page 275
- *Database Properties* on page 276

Creating a Proxy Database
Create a proxy database using the Sybase Control Center Administration Console.

Prerequisites

Consider these database attributes:

- Size :
  - **sp_estspace** helps you estimate table and index space requirements based on the definition of a specific table. See the *Adaptive Server Reference Manual* .
  - Space for planned views, stored procedures, defaults, rules, and triggers
  - Size of the transaction log
  - Space for anticipated expansion
- Database device location, and whether there is enough space on that device.

- Transaction log location – Sybase recommends that you store the transaction log on a different device than the data.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **Proxy Databases**.
3. Click the drop-down arrow and select **New**.
   You see the Add Proxy Database wizard.
4. On the Introduction screen, select the server in which to create a database.
5. On the Database screen, enter the name of the database you want to create.
6. (Optional) On the Devices screen, enter the size of the new database. If you do not enter a size, the default size allocated is 3MB. You can specify separately the amount of space to allocate to the log and data segments.
7. (Optional) On the Default Location screen, enter the name of the remote location where you want to store your proxy database. Select **For Proxy Update** to get metadata automatically from the remote location while creating proxy tables.
8. (Optional) Click **Summary** to see the database options you have selected.

See also

Creating an Archive Database
Create an archive database using the Sybase Control Center Administration Console.

Prerequisites

Consider these database attributes:

- Size :
  - **sp_estspace** helps you estimate table and index space requirements based on the definition of a specific table. See the *Adaptive Server Reference Manual* .
  - Space for planned views, stored procedures, defaults, rules, and triggers

---

- • Size of the transaction log
- • Space for anticipated expansion
- Database device location, and whether there is enough space on that device.
- Transaction log location – Sybase recommends that you store the transaction log on a different device than the data.

Task

Sybase Control Center provides a wizard to create an archive database.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **Archive Databases**.
3. Click the drop-down arrow and select **New**.
   You see the Add Archive Database wizard.
4. On the Introduction screen, select the server in which to create a database.
5. On the Database screen, enter the name of the database you want to create.

   To enter the name of an archive database, you must first select a scratch database. You can mark a database as a scratch database by selecting the Scratch Database option from the database property sheet.

   For information on scratch databases, see the *System Administration Guide: Volume 2*.
6. (Optional) On the Devices screen, enter the size of the new database. If you do not enter a size, the default size allocated is 3MB. You can specify separately the amount of space to allocate to the log and data segments.
7. (Optional) Click **Summary** to see the database options you have selected.

See also

Creating an In-Memory Database

Create an in-memory database using the Sybase Control Center Administration Console.

Prerequisites

Consider these database attributes:

- Size :
    - **sp_estspace** helps you estimate table and index space requirements based on the definition of a specific table. See the *Adaptive Server Reference Manual* .
    - Space for planned views, stored procedures, defaults, rules, and triggers
    - Size of the transaction log
    - Space for anticipated expansion
- Database device location, and whether there is enough space on that device.
- Transaction log location – Sybase recommends that you store the transaction log on a different device than the data.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **In-Memory Databases**.
3. Click the drop-down arrow and select **New**.
   You see the Add In-Memory Database wizard.
4. On the Introduction screen, select the server in which to create a database.
5. On the Database screen, enter the name of the database you want to create.
6. On the Devices screen, enter the size of the new database. If you do not enter a size, the default size allocated is 3MB. You can specify separately the amount of space to allocate to the log and data segments.
7. (Optional) On the Options screen, select:

    - **With override** to store the log and data on the same logical device - this is not recommended by Sybase.
    - **For load** if you want the database to be used for loading a database dump.
    - Specify a template database that is copied over to create the in-memory database.
8. (Optional) On the Guest User screen, select **Create guest user** to create a guest user who can access the database with limited privileges.
9. (Optional) Click **Summary** to see the database options you have selected.

See also

Creating an In-Memory Temporary Database

Create an in-memory temporary database using the Sybase Control Center Administration Console.

Prerequisites

Consider these database attributes:

- Size :
  - **sp_estspace** helps you estimate table and index space requirements based on the definition of a specific table. See the *Adaptive Server Reference Manual* .
  - Space for planned views, stored procedures, defaults, rules, and triggers
  - Size of the transaction log
  - Space for anticipated expansion
- Database device location, and whether there is enough space on that device.
- Transaction log location – Sybase recommends that you store the transaction log on a different device than the data.

Task

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **In-Memory Temporary Databases**.
3. Click the drop-down arrow and select **New**.
    You see the Add In-Memory Temporary Database wizard.
4. On the Introduction screen, select the server in which to create a database.
5. On the Database screen, enter the name of the database you want to create.
6. (Optional) On the Devices screen, enter the size of the new database. If you do not enter a size, the default size allocated is 3MB. You can specify separately the amount of space to allocate to the log and data segments.
7. (Optional) On the Options screen, select:

    - **With override** to store the log and data on the same logical device - this is not recommended by Sybase.
8. (Optional) On the Temporary Database Group screen, select the database group that the temporary database belongs to.
9. (Optional) Click **Summary** to see the database options you have selected.

See also
- *Creating a User Database* on page 259
- *Creating a Temporary Database* on page 261
- *Creating a Proxy Database* on page 262

- *Creating an Archive Database* on page 263
- *Creating an In-Memory Database* on page 264
- *Creating a Temporary Database Group* on page 267
- *Mounting an Adaptive Server Database* on page 275
- *Database Properties* on page 276

Creating a Temporary Database Group

Create a temporary database group using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **Temporary Database Groups**.
3. Click the drop-down arrow and select **New**.
   You see the Add Temporary Database Group wizard.
4. On the Introduction screen, select the server in which to create a database group.
5. On the Group Name screen, enter the name of the temporary database group.
6. On the Databases screen, specify the temporary database to be added to the temporary database group.
7. On the Bindings screen, click **Bind Application** to specify applications to be bound to the temporary database group.

   - **Bind Application** to specify applications to be bound to the temporary database group.
   - **Bind Login** to specify logins to be bound to the temporary database group.

     Note: If you change the binding of a login to a different group, the old binding is no longer valid.

8. (Optional) Click **Summary** to see the database options you have selected.

See also
- *Creating a User Database* on page 259
- *Creating a Temporary Database* on page 261
- *Creating a Proxy Database* on page 262
- *Creating an Archive Database* on page 263
- *Creating an In-Memory Database* on page 264
- *Creating an In-Memory Temporary Database* on page 266
- *Mounting an Adaptive Server Database* on page 275
- *Database Properties* on page 276

Backing Up Databases

Sybase Control Center helps you back up a database and its transaction log.

Prerequisites

- Ensure that you can connect to the Backup Server from each Adaptive Server you administer.
- Decide on the backup media you will use, and create dump devices that identify the physical backup media to Adaptive Server.
- Ensure that the login of the person starting the Backup Server has write permissions for the physical backup dump device, and that the dump device is available.

Task

Although Adaptive Server has automatic recovery procedures to protect you during power outages and computer failures, your best protection against media failure is regular and frequent backups of system and user databases. See the *System Administration Guide* for details on backup and recovery.

Note: Sybase Control Center checks to see if a Backup Server is available. If it is not, the Backup and Restore wizards will not start.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
    - **User Databases**
    - **System Databases**
    - **Temporary Databases**
    - **Proxy Databases**
    - **In-Memory Databases**
    - **In-Memory Temporary Databases**
3. Click the Name field of the database you want to back up.
4. Click the arrow and select **Backup**.
5. In the Backup Database wizard, choose from these options:

Table 34. Inputs for the Backup Database wizard

Input	Description
Type of backup	Back up either the entire database or the transaction log. If you back up only the transaction log, you can choose to: • Create either a new transaction log or a new transaction entry in the log • Delete the inactive part of the log
Physical File	Specify the dump device path for the backup.
Options	Optionally, specify: • Compression level – row-level, or page-level. • Remote Backup Server name, if different from **SYB_BACKUP**. • A password to protect the backup from unauthorized access. If you specify a password for backup, you must use the same password while restoring the database.
Dump Performance	Specify the amount of data to be dumped by the Backup Server: Default, Maximum, Minimum, or Advanced. Clicking **Advanced** allows you to specify threshold values (in percentages) for reserved and allocated pages. See **sp_dumpoptimize** in the *Adaptive Server Reference Manual*.

6. (Optional) Click **Summary** to verify your selected options.

7. Click **Finish** to start the backup. Sybase Control Center displays backup messages from the Adaptive Server.

See also

• *Restoring Databases* on page 269

Restoring Databases
Sybase Control Center helps you restore a database backup and its transaction log.

Prerequisites
Decide whether you will load the backup into a new database with the **for load** option, or into a preexisting database.

Task

You cannot load a database backup that was created on a different operating system, or with an earlier version of Adaptive Server.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.

2. Select one of:

- **User Databases**
- **System Databases**
- **Archive Databases**
- **In-Memory Databases**

3. Right-click the Name field of the database you want to back up.

4. Click the arrow and select **Restore**.

5. In the Restore Database wizard, choose from these options:

Table 35. Inputs for the Restore Database Wizard

Input	Description
Type of restore	Restore the entire database or the transaction log.
Physical file	Specify the dump device path for the restored database.
Options	Optionally, specify: • Remote Backup Server name, if different from **SYB_BACKUP**. • A password to access the backup. If you specified a password for backup, the same password must be used while restoring the database.

6. (Optional) Click **Summary** to verify your selected options.

7. Click **Finish** to start the restore process. Sybase Control Center displays restore messages from the Adaptive Server.

See also
- *Backing Up Databases* on page 268

Viewing Database Statistics

Sybase Control Center obtains database statistics by executing the Adaptive Server **optdiag** utility.

Sybase Control Center obtains database statistics by executing the Adaptive Server **optdiag** utility.

Note: To execute this command, you must have an agent configured for your Adaptive Server.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.

2. Select one of:
   - **User Databases**
   - **System Databases**
   - **Temporary Databases**
   - **Proxy Databases**

- **Archive Databases**
- **In-Memory Databases**
- **In-Memory Temporary Databases**

3. Click the Name field of the database.

4. Click on the arrow and select **Database Statistics**.
   You see table, page details such as data and empty page counts, space utilization, and other statistics for the selected database.

See also
- *Changing Database Ownership* on page 278
- *Modifying Database Storage Allocations* on page 279
- *Modifying the Transaction Log Cache and the Log I/O Buffer Size* on page 280
- *Changing Database Options* on page 280
- *Database Properties* on page 276

Checkpointing Databases

Use the **checkpoint** command to force Adaptive Server to write modified data pages from memory to disk.

When you issue a checkpoint, Adaptive Server freezes all current data-modifying transactions while writing to the disk. See the *Adaptive Server Reference Manual: Commands*.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
   - **User Databases**
   - **System Databases**
   - **Temporary Databases**
   - **Proxy Databases**
   - **Archive Databases**
   - **In-Memory Databases**
   - **In-Memory Temporary Databases**
3. Click the Name field of the database.
4. Click on the arrow and select **Checkpoint**.
5. Confirm that you want to run **checkpoint** on the current database.

See also
- *Checking Database Consistency* on page 272

Checking Database Consistency

Use the database consistency check to check the logical and physical consistency of a database.

Regular database consistency checks detect, and often correct, index and page allocation errors resulting in corrupted tables.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
   - **User Databases**
   - **System Databases**
   - **Temporary Databases**
   - **Proxy Databases**
   - **Archive Databases**
   - **In-Memory Databases**
   - **In-Memory Temporary Databases**
3. Click the Name field of the database.
4. Click the arrow and select **Check Consistency**.
5. In the Database Consistency Checker wizard, choose from these options:

Table 36. Inputs for the Database Consistency Checker Wizard

Input	Description
Check overall consistency	Run **dbcc checkdb**, which checks each table and index in the selected database.
	To skip checking nonclustered indexes on users tables, select **Ignore non-clustered indexes**; leave it unselected to check all indexes on all tables in the database.
	The generated report for each undamaged table shows the number of data pages and data rows.
Check allocation	Run **dbcc checkalloc**, which checks page allocation.
	To fix allocation errors, select **Fix allocation errors**. Adaptive Server automatically places the database in single-user mode while executing **dbcc checkalloc** and then returns the database to multiuser mode when processing is complete.
	The generated report shows the amount of space allocated and used by each database table, including the system tables. For each table or index, the report shows the number of pages and extents (8-page blocks of allocated space) used.

Input	Description
Check system cata-logs	Execute **dbcc checkcatalog** and check for consistency within and be-tween the system tables found in a database. The generated report lists the segments used by the database.

6. Click **Finish** to start the consistency check.

See also

- *Checkpointing Databases* on page 271

Placing a Database in Quiesce-Hold

Use **Quiesce Hold** to block updates to a database during a copy operation.

Quiesce hold allows you to block updates to one or more databases while you perform a disk unmirroring or external copy of each database device. Because no writes are performed during this time, the external, secondary copy of the database is identical to the primary image. While the database is in the quiescent state, read-only queries to operations on the database are allowed. You can load the external copy of the database onto a secondary server, ensuring that you have a transactionally consistent copy of your primary image.

Only database owners or system administrators can quiesce a database.

Note: If there are distributed or multidatabase transactions in the database in prepared state, Sybase Control Center waits for 5 seconds for those transactions to complete. If they do not complete in 5 seconds, the quiesce database hold operation fails.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.

2. Select one of:
- **User Databases**
- **System Databases**
- **Temporary Databases**
- **Proxy Databases**

3. Click the Name field of the database you want to quiesce.

4. Click the arrow and select **Quiesce Hold**.

5. In the Quiesce Database Hold wizard, choose from these options:

Table 37. Inputs for the Quiesce Database Hold Wizard

Input	Description
Tag Name	A tag name for the quiesce hold operation

Input	Description
External Dump Option	Copy the database while updates to specified databases are suspended with the **Quiesce Hold** command. You must also specify: • Manifest File – Specify the path for the manifest file. • Evaluate Dependencies – If you have not selected all the databases to be quiesced, allow the wizard to generate a list of databases that must be quiesced, along with your selected database, to ensure that the **Quiesce Hold** succeeds. The list of unselected databases that must be quiesced are indicated in the dependency matrix.

6. (Optional) Click **Summary** to verify your selected options.
7. Click **Finish** to start the quiesce-hold process.

See also
• *Placing a Database in Quiesce-Release* on page 274

Placing a Database in Quiesce-Release
Use **Quiesce Release** to resume database updates that were suspended by a **Quiesce Hold** command.

Issue **quiesce release** only when the external copy operation has completed.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
   • **User Databases**
   • **System Databases**
   • **Temporary Databases**
   • **Proxy Databases**
3. Click the Name field of the database in which you want to resume updates.
4. Click the arrow and select **Quiesce Release**.
5. Enter the tag information to release the database hold.

   Note: If you have **mon_role** permissions, you can select a tag from the displayed list of tags. Otherwise, enter the tag name in the provided text input box.
6. Click **Finish** to start the quiesce-release process.

See also
• *Placing a Database in Quiesce-Hold* on page 273

Mounting an Adaptive Server Database

Mount a user database on an Adaptive Server.

The **mount** command attaches the database to the destination or secondary Adaptive Server. **mount** decodes the information in the manifest file and makes the set of databases available online. The Adaptive Server also adds database devices, if necessary, and activates them, creates the catalog entries for the new databases, recovers them, and puts them online.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **User Databases**.
3. Select **Mount**.
   You see the Mount Database Wizard.
4. Select the server to which to attach the database.
5. Specify the path of the manifest file, and select **With Verify** to verify the devices specified on the manifest.
6. Verify that the device paths listed in the Device Specification screen are correct. Click any row to change the device path of the corresponding device.
7. (Optional) Click **Summary** to see the database options you have selected.

See also
* *Unmounting an Adaptive Server database* on page 275

Unmounting an Adaptive Server database

Unmount a database from an Adaptive Server.

When you unmount a database, you remove the database and its devices from an Adaptive Server. The unmount command shuts down the database. All tasks using the database are terminated. The database and its pages are not altered and remain on the operating system devices.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Click **User Databases**.
3. Click the Name field of the database to unmount, and select **Unmount.**
   You see the Unmount Database Wizard.
4. Specify the location where Adaptive Server will create the manifest file.
5. Select **Yes** on the Evaluate Dependencies screen to view any unselected databases that must be selected for the **unmount** command to succeed.

   Note: The **unmount** command fails unless you select all the databases on a device.

6. Select the databases listed in the Dependency Matrix screen for **unmount** to succeed. If no databases are listed in the Unselected Databases column, there are no dependencies.

7. Override referential integrity checks by selecting **With override**.

> **Note:** When the referencing database is dropped by the **unmount** command with an override, you cannot drop the referential constraints.

8. Enter a delay for distributed or multi-database transactions in prepared state to complete before the **unmount** command is activated. If the transactions do not complete in the specified time period, Adaptive Server does not execute the **unmount** command.

9. (Optional) Click **Summary** to see the database options you have selected.

See also
*   *Mounting an Adaptive Server Database* on page 275

Database Properties
Use the database Properties wizard to modify database options, cache options and storage allocation, extend log buffers, and change the owner.

Click **Properties** on your database to initiate the Properties wizard.

Wizard Option	Database Properties

General	• **Change Owner** – See *Changing Ownership of a Database* on page 278.
	• Instance name (cluster environment only)
	• Data cache – From the drop-down menu, select the cache to which you want to bind the database.
	• Durability level – Select one of:
	  • NO_RECOVERY – there is no guarantee that, at runtime, committed transactions are written to the disk.
	  • AT_SHUTDOWN – all committed transactions are written to disk during a normal server shutdown.
	  • FULL – a complete recovery of committed transactions is possible after a crash.
	• **Note:** These options only apply to nonclustered Adaptive Server version 15.5 and later.
	
	Default location – Specify the default storage location to be used for remote tables if no storage location is provided via the stored procedure sp_**addobjectdef**. See the section on sysdatabases in the System Tables chapter of the *Adaptive Server Reference Manual*.
	• **DML logging** – Click to enable DML logging.
	• **Database guest user** – Select if guest users are configured on the database.
	• **Resynchronize proxy tables** – Select this option to force re-synchronization of proxy tables in the proxy databases. See the **alter database** command in the *Adaptive Server Reference Manual* .

Devices	• Database devices – You can add or remove devices associated with a selected database. See *Modifying Database Storage Allocations* on page 279. • Transaction log – You can move the transaction log to a different location. See *Modifying Database Storage Allocations* on page 279.
Transaction Log	• Transaction log buffer size – You can modify the I/O buffer size of the transaction log. See *Modifying Transaction Log Buffer Size* on page 280.
Options	• Server configuration options – See *Changing Database Options* on page 280.
Usage	• Table and index space – Sybase Control Center displays a graph of the space used by the tables and indexes of your database. Use these values to determine if you have enough unreserved space to accommodate new database objects.

See also

Changing Database Ownership

Use the Database Properties wizard to change the owner of a database.

You can change ownership of a database to a user who is not a current user of the database and who does not have a current alias in the database.

Note: Only system administrators can change database ownership.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
    *   **User Databases**
    *   **System Databases**
    *   **Temporary Databases**
    *   **Proxy Databases**
    *   **Archive Databases**
    *   **In-Memory Databases**
    *   **In-Memory Temporary Databases**
3. Click the Name field of the database.
4. Click the arrow and select **Properties**.
    Sybase Control Center displays the Properties wizard.
5. On the default General screen, click **Change Owner**.
6. From the list, choose the login name for the new owner of the database. Additionally, you can choose to transfer all the aliases and their permissions to the new owner.
7. Click **OK**.

See also
- *Modifying Database Storage Allocations* on page 279
- *Modifying the Transaction Log Cache and the Log I/O Buffer Size* on page 280
- *Changing Database Options* on page 280
- *Viewing Database Statistics* on page 270
- *Database Properties* on page 276

Modifying Database Storage Allocations
Use the Database Properties wizard to add or modify space allocations for the database.

Note: Only system administrators can change database storage allocations.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
   - **User Databases**
   - **System Databases**
   - **Temporary Databases**
   - **Proxy Databases**
   - **Archive Databases**
   - **In-Memory Databases**
   - **In-Memory Temporary Databases**
3. Click the Name field of the database.
4. Click the arrow and select **Properties**.
   Sybase Control Center displays the Properties wizard.
5. Click **Devices**.
   You see the list of devices to which the database is allocated.
6. (Optional) Modify storage allocation for your database:
   - Click **Add** to add additional space from a different device for your database. Specify whether the space is to be allocated for data or for the transaction log.
   - Click **Remove** to remove the space allocated to your database from a device. You can only remove devices that are added using the **Add** option.
   - Click **Move log** to move the transaction log of a database, with log and data on the same device, to a separate device. See **sp_logdevice** in the *Adaptive Server Reference Manual*.
7. (Optional) Click **Create log or data fragment with override** to force Adaptive Server to allocate the data and log devices as specified, even if data and log are specified on the same device.
8. (Optional) Click **Preview** to see the SQL statements for your command.
9. Click **Apply**.

See also
- *Changing Database Ownership* on page 278
- *Modifying the Transaction Log Cache and the Log I/O Buffer Size* on page 280
- *Changing Database Options* on page 280
- *Viewing Database Statistics* on page 270
- *Database Properties* on page 276

Modifying the Transaction Log Cache and the Log I/O Buffer Size

Use the Database Properties wizard to modify the transaction log cache and the log I/O buffer size.

Change the size of the transaction log cache and log I/O buffer by binding the log to a cache of different size. The log buffer size determines the number of I/O transactions that can be stored in the transaction log I/O cache.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.
2. Select one of:
   - **User Databases**
   - **System Databases**
   - **Temporary Databases**
   - **Proxy Databases**
3. Click the Name field of the database.
4. Click the arrow and select **Properties**.
   Sybase Control Center displays the Properties wizard.
5. Click **Transaction Log**.
   You see the list of caches; the highlighted cache is the one currently configured for your database I/O buffer.
6. Select a different cache and click **Apply**.

See also
- *Changing Database Ownership* on page 278
- *Modifying Database Storage Allocations* on page 279
- *Changing Database Options* on page 280
- *Viewing Database Statistics* on page 270
- *Database Properties* on page 276

Changing Database Options

Use the Database Properties wizard to change database options using **sp_dboption**.

Note: Only database owners or system administrators can change options settings for individual databases.

Note: A newly created database has the same default settings as the **model** database.

1. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Databases**.

2. Select one of:
   - **User Databases**
   - **System Databases**
   - **Temporary Databases**
   - **Proxy Databases**

   Note: You cannot update any database options for the master database, or for archive databases.

3. Click the Name field of the database.

4. Click the arrow and select **Properties**.
   Sybase Control Center displays the Properties wizard.

5. Click **Options** to see the list of options that you can set for this database.

   Database options that you can set include:
   - **abort tran on full log** – determines how Adaptive Server treats active transactions when the database's log becomes critically low on space:
     - To cancel all user queries that need to write to the transaction log until space in the log has been freed, select this option.
     - To suspend transactions and awaken them when space has been freed, unset this option.
   - **allow nulls by default** – affects the ability of columns in newly created database tables to accept NULL values:
     - If you select this option, columns in newly created tables allow null values unless the column definitions explicitly state "not null."
     - If you do not select this option, nulls are not allowed unless the column definitions explicitly permit them.
   - **allow wide dol row** – allow wide, variable-length data-only locked (DOL) rows in user databases.

     Note: allow wide dol row is supported by Adaptive Server version 15.7 and later.
   - **async log service** – provide greater scalability in Adaptive Server and higher throughput in logging subsystems for high-end symmetric multiprocessor systems.

     Note: async log service is supported by Adaptive Server version 15.5 and later.
   - **auto identity** – automatically adds a 10-digit IDENTITY column in a new table when a user creates the table without specifying a primary key, a unique index, or an IDENTITY column.
   - **dbo use only** – restricts database access to the database owner.
   - **ddl in tran** – allows users to include DDL syntax within their transactions.

Generally, avoid using Data Definition Language commands inside transactions. For more information about this option, see the *Adaptive Server Reference Manual*.

- **delayed commit** – when enabled, all local transactions use delayed commits so that control returns to the client without waiting for the I/O on log pages to complete, and I/O is not issued on the last log buffer for delayed commit transactions. **delayed commit** is supported by Adaptive Server version 15.5 and later.

  Note: Delayed commit is not used if you enable both **delayed commit** and **async log service** for a database.

- **enforce dump tran sequence** – when set to "true", prevents operations that disallow a subsequent dump transaction.

  Note: enforce dump tran sequence is supported by Adaptive Server version 15.7 and later.

- **identity in nonunique indexes** – automatically includes an IDENTITY column in a table's index keys, so that all indexes created on the table are unique.

- **no chkpt on recovery** – sets the database so that a checkpoint record is added to the database after it is recovered due to restarting Adaptive Server.

  This checkpoint, which ensures that the recovery mechanism is not re-run unnecessarily, changes the sequence number on the database. If the sequence number on the secondary database has been changed, a subsequent dump of the transaction log from the primary database cannot be loaded into it.

  Select this option if you keep an up-to-date copy of a database. This prevents the secondary database from getting a checkpoint from the recovery process so that subsequent transaction log dumps from the primary database can be loaded into it.

- **no free space acctg** – determines whether the database enables free-space accounting and execution of threshold actions for non log segments.

  Suppressing free-space accounting speeds recovery time because the free-space counts are not recomputed for those segments. However, it disables updating the rows-per-page value stored for each table, so system procedures that estimate space usage may report inaccurate values.

  Note: System security officers can change the **no free space acctg** option.

- **read only** – prevents modification of any data in the database.

- **scratch database** – the database that stores the **sysaltusages** table. See the *System Administration Guide: Volume 2*.

- **select into/bulk copy/pllsort** – allows users to perform nonlogged operations. Nonlogged operations include **select into** for permanent tables, the bulk-copy utility **bcp**, and the **writetext** utility.

  You need not select this option to allow **select into** for temporary tables or to run **bcp** on a table with indexes, because inserts are logged.

Attempting to dump the transaction log in a database after unlogged changes have been made to the database with **select** only or bulk-copy produces an error message instructing you to use **dump database** instead.

- **single user** – allows only one user at a time to use the database.
- **trunc log on chkpt** – truncates the transaction log (removes committed transactions) every time the database is checkpointed.

  If you select this option, you cannot dump the transaction log. You may want to select this option during development work, when backups of the transaction log are typically not needed.

  > **Note:** If you select **trunc log on chkpt** for development purposes, clear it periodically and dump the transaction log. If you never dump the transaction log, it continues to grow, and eventually you run out of space in the database.

- **unique auto_identity index** – if a database's **auto_identity** is turned on, newly created tables automatically get a column named SYB_IDENTITY_COL. This helps maintain data integrity, since unique IDs are commonly used.

See also
- *Changing Database Ownership* on page 278
- *Modifying Database Storage Allocations* on page 279
- *Modifying the Transaction Log Cache and the Log I/O Buffer Size* on page 280
- *Viewing Database Statistics* on page 270
- *Database Properties* on page 276

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**
   - **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

You can save the DDL in an external file on your local file system.

Deleting a Database

Sybase Control Center helps you delete database objects, or the database itself.

Note: Deleting a database deletes all the objects of a database.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your database.
3. Click the Name field of the database you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the database.
6. Confirm the deletion.
7. Click **Finish**.

Defaults

Create and modify defaults using the Sybase Control Center Administration Console.

Manage Defaults

Specify a default value that can be referenced by an object if no value is explicitly supplied.

Creating a Default

Create a default using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Defaults**.
2. Click the drop-down arrow and select **New**.
3. On the Introduction screen, select the server, database, and owner of the new default.
4. Enter the name of the default.
5. Enter the expression to define the value of the default. Expressions must be constants, mathematical expressions, or built-in functions.
6. (Optional) Click **Summary** to verify your selected options.

Defaults Properties

Use the Properties wizard to access and modify information on defaults.

Click **Properties** on your default to initiate the Properties wizard.

Wizard Option	Properties of Defaults
General	• View the name, type, database, owner, creation date, and expression of the default. To change the default name, enter the new name and select **Apply.**
SQL	• View the SQL statements for creating the default.
Referenced By	• View the name, type, owner, and properties of objects that this default references. Click an object, then click **Properties** to view the object properties.

Deleting a Default

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your default.
3. Click the Name field of the default you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the default.
6. Confirm the deletion.
7. Click **Finish**.

Devices

Monitor the devices used by Adaptive Server.

Monitor Devices

Determining Device I/O Response and I/O per Second

Find out how long a device is taking to respond to I/O requests and what its I/O rate is.

High response time can indicate problems in the functioning of the physical device or the storage layer, problems with the configuration of the storage layer, or that the device is busy.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Devices** .

   You can also display the Devices screen by clicking a **Devices** link on another window in the Adaptive Server monitor.
3. (For shared-disk clusters only) Select **Global** to display global devices, and **Local** to display devices for instances of Adaptive Server clusters.

4. In the Devices table, select the device to monitor.

5. The **IO Response Time** column and the **Device IO/Sec** graph on the Details tab display details on the I/O activity on the selected device..

See also
- *Modifying Device Sizes* on page 286
- *Device Statistics and Details* on page 286

Modifying Device Sizes

Sybase Control Center allows you to increase the size of your devices from the Devices window.

1. In the Adaptive Server Monitor, select **Devices.**
   Sybase Control Center displays the Devices table.

2. Select the device to configure.

3. Right-click and select **Resize**.
   Sybase Control Center displays the device resize dialog with the name of the selected device, allocated size, input field for increased size, unit of size, and an option that allows you to specify whether to initialize the device.

4. Input the amount by which to increase the device size.
   The dialog box now displays the new device size that is calculated based on the input. If there is an error, it is indicated in the dialog box.

5. Click **OK**.

   For more information on devices, see the Adaptive Server *System Administration Guide*, Volume 2.

See also
- *Determining Device I/O Response and I/O per Second* on page 285
- *Device Statistics and Details* on page 286

Device Statistics and Details

Interpret the Devices screen for Adaptive Server.

Note: For Adaptive Server Cluster configurations, the **Devices** table is called **Global Devices** when you select the **Global** tab, and **Private Database Device** when you select the **Local** tab. Information in the Private Database Device table is grouped by cluster instance.

The Devices table displays information about all devices that store databases for this Adaptive Server. A device can be a whole disk drive, or any part of a disk or file system. The charts are populated by data from collection_ase_all_client_kpis, covering the current trend period.

The **Devices** table includes device semaphore statistics. The device semaphore controls access to device I/O; a high ratio of **Device Semaphore Waits** to **Device Semaphore Requests**

indicates contention. If **IO Wait Time** is high enough to cause concern, you may want to redistribute the data on the physical devices.

The tabs at the bottom of the screen show information about the device selected in the **Devices** table.

Table 38. Tabs on the Devices Screen

Details	Displays two charts: • A pie chart showing space usage on the selected device. Includes used and unused space, in megabytes, and as percentages of all the available space on the device. The title above the chart indicates the total available space. • **Device IO/Sec** – a line graph showing the rate of I/O per second on the selected device over the current trend period. The graph shows the sum of reads, writes, and asynchronous prefetch (APF) reads. Because the graph shows a rate, and the read, write, and APF read figures in the table are changes since the last refresh, the values do not correspond. **Note:** Sybase Control Center does not display the Device IO/Sec graph for in-memory devices.
Advanced	Displays two charts: • **Device IO Response Time** – a line graph showing the response time, in milliseconds, for I/O operations performed on the selected device. • **Device APF Reads/Sec** – a line graph showing the rate of asynchronous prefetch read operations, per second, on the selected device. APF reads indicate that table scans are taking place. **Note:** Sybase Control Center does not display the Advanced tab for in-memory devices.
IO Distribution	Appears only for **Global Devices** on Adaptive Server shared-disk clusters. Details for master device I/O activities for each instance include I/O wait time, response time, APF reads and request and wait times for device semaphores.

See also
• *Engine Statistics and Details* on page 294
• *Process Statistics and Details* on page 324
• *Determining Device I/O Response and I/O per Second* on page 285
• *Modifying Device Sizes* on page 286
• *Setting Up Statistics Collection* on page 120

Manage Devices

Sybase Control Center allows you to create, delete, and generate data definition language for database, device, and dump devices.

Displaying a Device Object

You can use a wizard to look at device objects.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Devices**.
2. Click:

   - **Database Devices** to view a list of database devices
   - **In-Memory Devices** to view a list of in-memory devices

   Both lists display:
   - **Name** – name of the database device.
   - **Server** – name of the server in which the database resides.
   - **Size** – displays the megabytes used by the device.
   - **Unused size** – displays the amount of unallocated space for the device, in megabytes.
   - **Physical name** – name of the physical device.

Creating a Database Device

Use a wizard to create new database devices.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Devices**.
2. Click the drop-down arrow on **Database Devices** and select **New**.
3. Select the server in which to create the database device.
4. (Optional) For a cluster environment, click on **Create this as a private database device**, then select the instance on which to create the private device.

   Private devices are used only by local user temporary databases. For the Cluster Edition, you can use block devices for managing the storage needs of temporary data in the cluster. These devices are added as private devices and can only be used by local user temporary databases. See, *"Using Temporary Databases" in the Clusters User Guide* got more information.
5. On the Device Name and Path screen, enter:

   - **Device name** – the logical device name used in the **create database** and **alter database** commands.

- **Device path** – the physical device name, usually in the form of a full path for the new file, or in UNIX, a raw device partition. If you do not specify a device path, the wizard fills this field with the device name with a `.dat` file extension.

6. On the Advanced Options screen, specify:
   - The device size in megabytes
   - The device number – a unique number that identifies this device on the server, **Lookup** to see what devices use what device numbers on the server. The default in this field is the wizard's recommended number.
   - Starting address – The virtual starting address, or the offset, for Adaptive Server to begin using the database device. Defines the starting address for this device, and is a virtual offset in 2KB blocks. The default is 0. See *Other optional parameters for **disk init*** in the *System Administration Guide* for information about **vstart.**
   - Skip initialization of device – select to speed up the resizing of the device.

7. Select a write option:
   - **Data sync** – Ensures that writes to the database device occur on the physical storage medium. This allows Adaptive Server to recover data from the device when a system failure occurs.
   - **Direct IO** – Transfers the data directly to disk, bypassing the operating system's buffer cache.
   - **Cached IO** – Turns off the data sync option, and any writes to the database device are buffered into the file system. During system failures, Adaptive Server does not recover any data that has not been updated to the physical medium.

8. (Optional) On the Mirroring screen, you can click **Mirror the database device** and specify the path for the duplicate device.

   Note: If the server is not configured to enable disk mirroring, the options for the Mirroring wizard page are grayed out.

See also
- *Creating a Dump Device* on page 290
- *Creating an In-Memory Device* on page 289
- *Deleting a Database, In-Memory, or Dump Device* on page 293
- *Generating a DDL Script* on page 293
- *Database Device Properties* on page 291

Creating an In-Memory Device
Use a wizard to create an in-memory device—or a cache device—in a cache created for an in-memory database. This device resides on an in-memory storage cache, and allows you to create in-memory databases.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Devices**.

2. Click the drop-down arrow on **In-Memory Devices** and select **New**.

3. (Adaptive Server version 15.5 and later only) Select the server on which to create the in-memory device.

   Note: The In-Memory Device wizard is only available for Adaptive Server 15.5 and later releases. If a valid server user is not selected, the wizard will be disabled and error messages will be displayed at the bottom of the server selection page.

4. On the Device Name screen, enter the logical device name, which Adaptive Server uses in its **create database** and **alter database** commands.

5. The In-Memory Storage screen displays a list of caches on which to create your in-memory device. On this screen, you can:

   • Add – displays the Specify Cache Device and Size wizard page, and lets you choose a cache to create the device in. The default size of an in-memory device is 6MB. If the in-memory storage is smaller than 6MB, the device size automatically matches the in-memory storage size.

   • Edit – allows the in-memory device to require more space from the in-memory storage. You cannot, however, increase the size of the storage itself.

   • Remove – removes the selected cache.

See also
• *Creating a Database Device* on page 288
• *Creating a Dump Device* on page 290
• *Generating a DDL Script* on page 293
• *Deleting a Database, In-Memory, or Dump Device* on page 293
• *In-Memory Device Properties* on page 292

Creating a Dump Device
Use a wizard to create a dump device on a server.

A dump device is a tape, partition, or file used for database or transaction dumps.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Dump Devices**.

2. Click the drop-down arrow on **Dump Devices** and select **New**.

3. On the Introduction screen, select the server on which to create the dump device.

4. On the Device Name and Path screen, enter:

   Dump device path – The physical device path.

   Dump device name – The name of the dump device.

5. On the Advanced Options screen, specify the type of device to create:

- Disk dump device
- Tape dump device – If you choose this type, specify the device size, in megabytes.

See also
- *Creating a Database Device* on page 288
- *Creating an In-Memory Device* on page 289
- *Generating a DDL Script* on page 293
- *Deleting a Database, In-Memory, or Dump Device* on page 293
- *Dump Device Properties* on page 292

Database Device Properties
Use the database device properties option to view the device's general information, as well as analyze its mirror status, databases that occupy the space on the device, and segments contained in the device. You can also modify the write option and increase the space for this device.

Click **Properties** on your database device to initiate the Properties wizard.

Wizard Option	Database Device Properties
General	• Name and type of device • I/O fencing support (cluster environment only) • Instance name (cluster environment only) • Physical name and path • Space allocated – you can change the value, in megabytes. • Default device – specify whether to set this device as the default device. • Write option – choose data sync, direct I/O, or cached I/O (data sync off) for this device.
Mirror	Displays whether disk mirroring is enabled. If it is, you can select: • Mirror device – turn on mirror device, choose whether the mirror is written after the primary is written or in parallel with parallel writes, and specify the mirror path in a file name relative to the server. • Disable mirror – select either the primary or secondary device to disable, and whether this disabling is temporary or permanent.
Databases	Displays a list of databases that occupy space on the device. Click **Properties** to see the database properties without going to the Database view.
Segments	Displays a list of segments contained in the device. Click **Properties** to view segment properties without going into the segment view.

See also
*   *Creating a Database Device* on page 288
*   *Deleting a Database, In-Memory, or Dump Device* on page 293
*   *Generating a DDL Script* on page 293

<u>*In-Memory Device Properties*</u>
Use the In-Memory Device Properties wizard to view general information, and analyze the list of databases and segments contained in the in-memory device.

Click **Properties** on your in-memory device to initiate the Properties wizard.

Wizard Option	In-Memory Device Properties
General	Displays: • In-memory storage object name • Type information • Status • Current size – check the size of the in-memory device's storage by pages, kilobytes, megabytes, or gigabytes.
In-Memory Database	Displays the in-memory database created on the in-memory device. Click **Properties** to see the database properties.
Segment	Displays a list of segments contained in the in-memory device. Click **Properties** to see the segment properties without going into the Segment view..

See also
*   *Creating an In-Memory Device* on page 289
*   *Generating a DDL Script* on page 293
*   *Deleting a Database, In-Memory, or Dump Device* on page 293

<u>*Dump Device Properties*</u>
Use the Dump Devices Properties wizard to view information about the dump device.

Click **Properties** on your dump device to initiate the Properties wizard. The wizard displays **General** on the left column, and you can view:

*   **Name** – is the name of of the dump device.
*   **Type** – is the type of dump device, such as tape device.
*   **Physical name** – is the full path of dump device.
*   **Capacity (MB)** – is the storage capacity, in megabytes.

See also
- *Creating a Dump Device* on page 290
- *Generating a DDL Script* on page 293
- *Deleting a Database, In-Memory, or Dump Device* on page 293

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**
   - **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a Database, In-Memory, or Dump Device
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your device.
3. Click the Name field of the device you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the device.
6. Confirm the deletion.
7. Click **Finish**.

Engines

Monitor Engines
Monitor Adaptive Server engines.

Find out how heavily Adaptive Server engines are being used.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Engines**.
3. In the Engines table, select the engine you want to monitor.
   The I/O Processing and Garbage Collection tables and the Engine CPU Utilization graph at the bottom of the screen are populated with data for the selected engine.

See also
*   *Engine Statistics and Details* on page 294

Engine Statistics and Details
Interpret the Engines screen for Adaptive Server.

The Engines screen displays information about all processing engines for this Adaptive Server. The charts are populated by data from the collection_ase_all_client_kpis, covering the current trend period.

The Engines table identifies engines by number, and gives CPU utilization percentages, status, start date and time, number of connections, and the operating system process identifier (**OS PID**) for each one.

Note: For Adaptive Server shared-disk clusters, the Engines table has information that is grouped by cluster instances.

The area at the bottom of the screen shows information about the selected engine selected.

IO Processing table	Provides counts of disk I/O checks, checks without waits, polls, and completed operations over the current trend period.
Garbage Collection table	For the current trend period, provides the garbage collector's maximum queue size and counts of pending items, high water mark items, and overflows.
Engine CPU Utilization graph	A line graph showing CPU utilization for this engine as a percentage. If Adaptive Server is performing poorly, use the information from this graph to determine how busy the engines are.

See also
*   *Displaying Engine CPU Utilization* on page 294

Manage Engine Groups

Manage Adaptive Server engine groups.

Creating Engine Groups

Create groups of Adaptive Server engines or processes that run in parallel.

Engine groups are useful only in multiprocessor systems.

1. In the left pane of the Administration Console, expand **ASE Servers > Performance**, then select **Engine Groups**.
2. Click the drop-down arrow and select **New**.
   You see the Add Engine Group wizard.
3. Select the Adaptive Server on which to create the engine group.
4. For a cluster environment, select an online cluster instance on which to create the engine group.
5. Specify the name of the engine group.
6. Select the engines in this engine group.
7. (Optional) Click **Summary** to verify your selected options.

See also

Engine Groups Properties

In the Engine Groups screen, select **Properties** on an engine group name to initiate the Properties wizard.

Wizard Option	Engine Group Properties
General	• Name and type • For a cluster environment, the online instance on which the engine group was created is displayed. • Properties – select an execution class and click **Properties** to view the properties of the execution class.

Wizard Option	Engine Group Properties
Engines	Select an engine from the list of engines, and click: • Add – to add the engine to your engine group. • Remove – to remove the engine from an engine group. **Note:** You cannot remove the last engine from the group.

See also
*   *Thread Pool Properties* on page 406
*   *Execution Classes Properties* on page 310
*   *Thread Statistics and Details* on page 404
*   *Creating a Thread Pool* on page 405
*   *Creating Execution Classes* on page 310
*   *Creating Engine Groups* on page 295

Deleting an Engine Group
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your engine group.
3. Click the Name field of the engine group you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the engine group.
6. Confirm the deletion.
7. Click **Finish**.

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
    *   **Compiled Objects**
    *   **Schema Objects**
    *   **Security**
    *   **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**

Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.

3. Click in the Name field of the specific object for which you want the DDL script.

4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Extended Stored Procedures

Create, delete, modify, and administer extended stored procedures using the Sybase Control Center Administration Console.

Manage Stored Procedures

Manage extended stored procedures using the Sybase Control Center Administration Console.

Creating an Extended Stored Procedure

Create an extended stored procedure using the Sybase Control Center Administration Console.

1. In the Administration Console view, select **ASE Servers > Compiled Objects > Extended Stored Procedures**.

2. Click the drop-down arrow and select **New**.
   You see the Create Extended Stored Procedure wizard.

3. On the Introduction screen, select the server, database and owner of the new extended stored procedure.

4. Enter the name of the extended stored procedure.

5. Enter the name of the dynamic link library or shared library that is executed when an application invokes the extended stored procedure.

6. (Optional) Click **Preview** to see the SQL statements for your command.

7. (Optional) Click **Summary** to verify your selected options.

See also

Extended Stored Procedures Properties

Use the extended stored procedure Properties wizard to access and modify information on extended stored procedures.

Click your extended stored procedure and select **Properties** to initiate the Properties wizard.

Wizard Option	Extended Stored Procedure Properties
General	View the name, type, database, owner, creation date, and dynamic link library (DLL) path of the stored procedure. The DLL need not exist when you create the extended stored procedure, but it must exist when you execute the extended stored procedure.
Permissions	Grant and revoke permissions on an extended stored procedure to users, groups, or roles. Choose the **Grant** option to allow the grantee to further grant permissions to other users. Select an object in the table of permissions, and click **Properties** to view the object properties. See *Granting Permissions on Database Objects* on page 313 and *Revoking Permissions on Database Objects* on page 314.
Referenced By	View the name, type, owner, and properties of objects that are referenced by this extended stored procedure. Click an object, then click **Properties** to view the object properties.

See also

- *Manage Stored Procedures* on page 311
- *Creating an Extended Stored Procedure* on page 297

Granting Permissions on Database Objects
Grant permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.
2. Click the **Name** field of the object on which to grant permissions.
3. Click the drop-down arrow and select **Properties**.
4. Select **Permissions**.
5. Click **Grant** to grant access permissions for the selected object.
You see the Grant Permission wizard.
6. Select the type of grantee:

   - **Users**
   - **Groups**
   - **Roles**
7. Select one or more grantees.
8. Select the types of permissions allowed for the selected grantees.

9. (Optional) Click **Summary** to verify your selected options.

Revoking Permissions on Database Objects

Revoke permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.
2. Click the **Name** field of the object on which to revoke permissions.
3. Click the drop-down arrow and select **Properties**.
4. Select **Permissions**.
5. Select the grantee, then click **Revoke** to revoke access permissions to the object.
    In the Revoke Permissions wizard, each type of permission and the current granted permissions are shown in cells.
6. Choose one of:

    •   **Revoke all permission.**

    •   Individual cells to revoke the currently granted permissions. The cell changes to show an "x", indicating that the permission type is no longer granted.
7. (Optional) Click **Preview** to see the SQL statements for your command.

Generating a DDL Script

Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
    •   **Compiled Objects**
    •   **Schema Objects**
    •   **Security**
    •   **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
    Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

    You can save the DDL in an external file on your local file system.

Deleting an Extended Stored Procedure

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your extended stored procedure.
3. Click the Name field of the extended stored procedure you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the extended stored procedure.
6. Confirm the deletion.
7. Click **Finish**.

Functions

Manage scalar and SQLJ functions using the Sybase Control Center Administration Console.

Manage Scalar Functions

Create, delete, modify, and administer scalar functions using the Sybase Control Center Administration Console.

Creating a Scalar Function

Create a scalar function using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects > Functions**, then select **Scalar Functions**.
2. Click the drop-down arrow and select **New**.
    You see the Create Scalar Function wizard.
3. On the Introduction screen, select the server, database, and owner of the new function.
4. Enter the name of the function.
5. Select the datatype of the value returned by the function.
6. On the Compilation Option screen, you can elect to have your function recompiled every time it is executed. This is useful if you expect parameter values to change frequently. If you do not select this option, the function is compiled only the first time it is executed.
7. On the SQL Editor screen, provide the SQL statements for the scalar function. Ensure that all objects referenced by the function exist in the database.
8. (Optional) Click **Preview** to see the SQL statements for your command.
9. (Optional) Click **Summary** to verify your selected options.

See also

- *Scalar Function Properties* on page 301

Scalar Function Properties

Use the scalar function properties wizard to access and modify information on stored procedures.

Click your scalar function and select **Properties** to initiate the Properties wizard.

Wizard Option	Scalar Function Properties
General	• View the name, type, database, owner, and creation date and group number of the function.
SQL	• View the SQL statements for creating the function.
Parameters	• View the name, type, mode, and order of all the function parameters. The mode value indicates whether it is an input or an output parameter. The order is a numeric value that indicates the place of the parameter in the list of parameters. • To change the parameters, change the definition of the function by dropping and re-creating the function.

See also
- *Creating a Scalar Function* on page 300
- *Creating a SQLJ Function* on page 302

Generating a DDL Script

Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**
   - **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.

3. Click in the Name field of the specific object for which you want the DDL script.

4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a Scalar Function

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your scalar function.

3. Click the Name field of the scalar function you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the scalar function.

6. Confirm the deletion.

7. Click **Finish**.

Manage SQLJ Functions

Create, delete, modify , and administer SQLJ functions using the Sybase Control Center Administration Console.

See also

• *Manage SQLJ Procedures* on page 315

Creating a SQLJ Function

Create a SQLJ function using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects > Functions**, then select **SQLJ Functions**.

2. Select **New**.
   You see the Create SQLJ Function wizard.

3. On the Introduction screen, select the server, database, and owner of the new function.

4. Enter the name of the function.

5. Specify the external name, which identifies the Java method, class, and an optional package name.

6. Select the datatype of the value returned by the function.

7. On the SQL Properties screen, select:
   • Null input – select to either return null if input is null, or to execute the function with null input.
   • Modifies SQL data – indicate that the Java method invokes SQL operations and modifies SQL data in the database.

- Exportable – specify if this function may be run on a remote server using the Adaptive Server OmniConnect™ feature. Both the procedure and the method it is built on must exist on the remote server.
- Deterministic option – include the keywords deterministic or not deterministic for compatibility with the SQLJ standard. However, Adaptive Server does not make use of this option.

8. On the SQL Editor screen, provide the SQLJ statements for the function. Ensure that all objects referenced by the function exist in the database.

9. (Optional) Click **Preview** to see the SQL statements for your command.

10. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Scalar Function* on page 300
- *Scalar Function Properties* on page 301
- *Creating a SQLJ Procedure* on page 315
- *SQLJ Function Properties* on page 303
- *SQLJ Procedure Properties* on page 316

SQLJ Function Properties

Use the SQLJ function properties wizard to access and modify information on SQLJ functions.

Click your SQLJ function and select **Properties** to initiate the Properties wizard.

Wizard Option	SQLJ Function Properties
General	• View the name, type, database, owner, and creation date and group number of the function.
SQL	• View the SQL statements for creating the function.
Parameters	• View the name, type, mode, and order of all the function parameters. The mode value indicates whether it is an input or an output parameter. The order is a numeric value that indicates the place of the parameter in the list of parameters. • To change the parameters, change the definition of the function by dropping and re-creating the function.

See also
- *SQLJ Procedure Properties* on page 316
- *Creating a SQLJ Function* on page 302

- *Creating a SQLJ Procedure* on page 315

Generating a DDL Script

Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**
   - **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a SQLJ Function

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your SQLJ function.
3. Click the Name field of the SQLJ function you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the SQLJ function.
6. Confirm the deletion.
7. Click **Finish**.

Networks

Manage remote servers.

Managing Remote Servers

Add, delete, or configure remote servers using Sybase Control Central.

Configuring Adaptive Server for Remote Procedure Calls

Configure Adaptive Server installations to allow request for execution of stored procedures on a remote server from a local server. The results of this request is called a remote procedure call (RPC).

Your choice of RPC handling method affects Adaptive Server configuration and login mapping for remote servers. Configuration options for RPC handling are site handler and Component Integration Services (CIS).

The default method for handling interaction between local and remote servers is through a site handler. A site handler creates a physical connection between the local server and remote server. Then it creates a logical connection for each RPC to the remote server. Adaptive Server creates a site handler for each remote server it connects to. Site handler is used only for connections between two Adaptive Server installations.

You can enabling CIS for Adaptive Server to request execution of stored procedures and access data on a remote server as if it were on the local server. CIS RPC handling is always used for connections involving proxy tables.

The principal difference between the two methods of handling RPCs is how the remote server views the RPC:

- If you use site handler, the remote Adaptive Server detects that the logical connection is made by another remote server and performs remote server verification through *sysremotelogins*.
- If you use CIS RPC handling, the remote server sees the RPC as an ordinary client connection. There is no verification using *sysremotelogins*. Therefore, connections must have a valid Adaptive Server login account established prior to the connection request. You cannot use trusted mode. Use of CIS RPC handling allows you to include RPCs in a transaction. Work done by an RPC can be committed or rolled back along with the other work performed in the transaction.

See also

- *Testing a Remote Server Connection* on page 307
- *Adding a Remote Server* on page 305

Adding a Remote Server

To gain access to a remote server, it must be defined on the local Adaptive Server.

To add a remote server you must register and authenticate the agent for Adaptive Server.

1. In the left pane of the Administration Console, expand **Networks**, then select **Remote Servers**.

2. Click the drop-down arrow and select **New**.
   You see the Add Remote Server wizard.

3. On the Introduction screen, select the local server.

4. On the Remote Server Name screen, specify the local name for the remote server.

5. Enter a network name for the remote server.

6. Specify the server class of the remote server.
   If Component Integration Services (CIS) is enabled, specify the server class of the remote server. If CIS is not enabled, accept the default server class: **ASEnterprise**.

7. (Optional) Click **Summary** to see the SQL statement and verify your selected options.

See also
* *Configuring Adaptive Server for Remote Procedure Calls* on page 305

Deleting a Remote Server

Delete a remote server from the Adaptive Server system tables.

1. In the left pane of the Administration Console, expand **Networks**, then select **Remote Servers**.

2. Select a server from the right pane.

3. Click the drop-down arrow and select **Delete**.

4. Click **OK** to confirm.

Remote Server Properties

Use the Remote Server Properties wizard to test a connection, change the server class, map local logins and remote logins, and change configuration options.

Wizard Option	Remote Server Properties
General	Displays remote server general information including the remote server name, and type, the network name, and server class.
	Provides an option to test the connection to a remote server.
	See *Testing a Remote Server Connection* on page 307.
Options	Provides options for configuring a remote server. Setting options is only available for a user created remote server.
	See *Setting Options for a Remote Server* on page 307.

Wizard Option	Remote Server Properties
Login Mapping	Provides options to: • Manage default mapping for remote logins when called from a remote server. • Manage remote logins specifically mapped to local logins when called from a remote server. See *Manage Remote Server Logins Mappings* on page 308.
CIS Mapping	Manage the mapping of local logins or roles to remote logins when access to the remote server is through Component Integration Services. See *Manage CIS Roles and Logins Mappings* on page 309.

Testing a Remote Server Connection
Verify that a connection can be established between the local server and the remote server.

1. In the Administration Console, select **Networks > Remote Servers**.
   You see the Remote Server Properties window.
2. In the General window, select the server class for the remote server and click **Test Connection**.

See also
• *Configuring Adaptive Server for Remote Procedure Calls* on page 305

Setting Options for a Remote Server
View or change or remote server options.

Set remote server configuration options:

server cost	Component Integration Services only – specifies the cost of a single exchange under the user's control, on a per-server basis.
cis hafailover	Component Integration Services only – if enabled, instructs Open Client to use automatic failover when connections fail. In this case, CIS connection failures automatically failover to the server specified in directory services (such as the interface file and ldap server) as the failover server.
enable login redirection	Is used by the Adaptive Server workload manager to send incoming connections to specific instances based on the logical cluster configuration and the cluster's current workload.
External engine auto start	Specifies that EJB Server starts up each time Adaptive Server starts up. The default is true; starting Adaptive Server also starts up EJB Server.

mutual authentication	Verifies the identity of the client and the server. The local server initiating the remote connection can request mutual authentication for all remote connection requests to target an Adaptive Server. This allows the client to verify the identity of the remote server.
negotiated logins	Component Integration Services only – this option is necessary if CIS connections to XP server or Backup Server are required. When enabled, Omni connects to the specified server and establishes a callback handler that can respond appropriately to login challenges from XP Server and Backup Server.
net password encryption	Specifies whether to initiate connections with a remote server with the client side password encryption handshake or with the normal (unencrypted password) handshake sequence. The default is false, no network encryption.
readonly	Component Integration Services only – specifies that access to the server named is read only.
relocated joins	Relocation joins permits joins between local and remote tables to locate to remote server.
security mechanism	External software that provides security services for a connection.
server logins	Component Integration Services only – to fully support remote logins, Client-Library provides connection properties that enable CIS to request a server connection. This connection is recognized at the receiving server as a server connection (as opposed to an ordinary client connection), allowing the remote server to validate the connection as if the connection were made by a site handler.
timeouts	When unset (false), disables the normal timeout code used by the local server, so the site connection handler does not automatically drop the physical connection after one minute with no logical connection.
use message confidentiality	Data is encrypted over the network to protect against unauthorized disclosure.
use message integrity	Verifies that communications have not been modified during transport.

Managing Remote Server Login Mappings
Add, remove, and configure remote server logins mappings.

1. In the Administration Console, select **Networks > Remote Servers**.
2. Select **Properties**.
3. In the Login Mapping window, choose how the logins from a remote server will be mapped to a local server.

- None – a particular remote login is mapped to a particular local login name. For example, user joe on the remote server might be mapped to joesmith.
- Map to local logins with the same names – all logins from one remote server can use their remote names.
- Map all to a single local login – all logins from one remote server can be mapped to one local name. For example, all users sending remote procedure calls from the MAIN server are be mapped to remusers.

> **Note:** Mapping more than one remote login to a single local login reduces individual accountability on the server. Audited actions can be traced only to the local server login, not to the individual logins on the remote server.

4. (Optional) To map a particular remote login to specific local login name, click **Add**.
   a) In the Add Specific Mapping window, specify the remote login name, then select the local login name.
   b) (Optional) Click on Trusted Password to indicate that the remote logins are trusted. Using the Trusted Password option reduces the security of your server, as the passwords are not verified.
5. (Optional) To remove a mapping of particular remote login specific local login name, select the login, then click **Remove**.

See also

Managing CIS Roles and Logins Mappings
Add, remove, and configure remote server CIS roles and logins mappings.

Logins and roles for CIS RPC handling are mapped on the local server level. By default, your local login is used as the remote login.

1. In the Administration Console, select **Networks > Remote Servers**.
2. Select **Properties**.
3. To add a mapping for CIS, click **Add** in the CIS Mapping window
   a) In the Add Login Mapping window, select the local login or role, then specify the remote login name.
   b) Type in the remote password and confirm the password.
4. (Optional) To remove a CIS mapping, select the remote login name, then click **Remove**.

See also

Execution Classes
Use the Administration Console to manage execution classes.

Manage Execution Classes

Create, delete, and modify execution classes using the Sybase Control Center Administration Console.

Creating Execution Classes

Create execution classes that can be bound to logins or tasks.

1. In the left pane of the Administration Console, expand **ASE Servers > Performance**, then select **Execution Classes**.
2. Click the drop-down arrow and select **New**.
   You see the Add Execution Classes wizard.
3. Select an Adaptive Server that is configured to run in threaded mode.
4. Specify the name of the execution class you want to create.
5. Select the priority of the execution class to High, Medium, or Low. This priority determines the priority of the tasks that are bound to the execution class. It also determines the priority of the tasks run by the logins associated with the execution class.
6. Specify the affinity of the execution class. This is the thread pool or engine group associated with the execution class. If the Adaptive Server version is running in threaded mode, the tasks bound to the execution class can run only on the thread from the chosen thread pool.
7. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Thread Pool* on page 405
- *Creating Engine Groups* on page 295
- *Thread Pool Properties* on page 406
- *Execution Classes Properties* on page 310
- *Engine Groups Properties* on page 295
- *Thread Statistics and Details* on page 404

Execution Classes Properties

In the Execution Classes screen, select **Properties** on an execution class name to initiate the Properties wizard.

Wizard Option	Execution Class Properties
General	• Priority – modify the priority of your execution class. • Properties –modify the thread pool properties. See *Thread Pool Properties* on page 406.

Wizard Option	Execution Class Properties
Bindings	Change the process, login, or stored procedure that is bound to this execution class. See *Modifying Bindings to Execution Classes* on page 311.

See also
- *Thread Pool Properties* on page 406
- *Engine Groups Properties* on page 295
- *Thread Statistics and Details* on page 404
- *Creating a Thread Pool* on page 405
- *Creating Execution Classes* on page 310
- *Creating Engine Groups* on page 295

Modifying Bindings to Execution Classes
Change the scope and bindings of execution classes.

1. In the left pane of the Administration Console, expand **ASE Servers > Performance**, then select **Execution Classes**.
2. Click the **Name** field of the execution class to modify.
3. Click the drop-down arrow and select **Properties**.
4. Click **Bindings**.
5. Click **Bind** to bind objects to an execution class.
   a) Select the scope of the execution class.
   b) Select the login to bind to the execution class.
6. (Optional) Select a login bound to the execution class and click **Unbind** to release the binding.
7. (Optional) Select a login bound to the execution class and click **Properties**.

Procedures
Manage stored and SQLJ procedures using the Sybase Control Center Administration Console

Manage Stored Procedures
Create, delete, or modify stored procedures using the Sybase Control Center Administration Console.

Stored procedures are named collections of SQL statements and flow control statements. A stored procedure that performs a **select**, **execute**, or data modification command must have the same owner as the object acted upon.

A system administrator, a database owner, or a user or group with **create procedure** permission can create a stored procedure.

See also
- *Extended Stored Procedures Properties* on page 297
- *Creating an Extended Stored Procedure* on page 297

Creating a Stored Procedure
Create a stored procedure using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects > Procedures**, then select **Stored Procedures**.
2. Click the drop-down arrow and select **New**.
    You see the Create Stored Procedure wizard.
3. On the Introduction screen, select the server, database, and owner of the new procedure.
4. Enter the name of the procedure.
5. On the Compile Option screen, you can opt to have your procedure recompiled every time it is executed. This is useful if you expect parameter values to change frequently. If you do not select this option, the procedure is compiled only the first time it is executed.
6. On the Stored Procedure Group screen, you can specify a group number to which to add the stored procedure. Grouping together all stored procedures that belong to a certain application lets you drop all procedures with a single command.
7. On the SQL Editor screen, provide the SQL statements for the procedure. Ensure that all objects referenced by the procedure exist in the database.
8. (Optional) Click **Preview** to see the SQL statements for your command.
9. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a SQLJ Procedure* on page 315
- *Extended Stored Procedures* on page 297
- *Stored Procedure Properties* on page 312

Stored Procedure Properties
Use the stored procedure Properties wizard to access and modify information on stored procedures.

Click **Properties** on your stored procedure to initiate the Properties wizard.

Wizard Option	Stored Procedure Properties
General	• View the name, type, database, owner, creation date, and group number of the procedure.
SQL	• View the SQL statements for creating the procedure.
Parameters	• View the name, type, mode, and order of all the procedure parameters. The mode value indicates whether it is an input or an output parameter. The order is a numeric value that indicates the place of the parameter in the list of parameters. • To change the parameters, change the definition of the stored procedure by dropping and re-creating the procedure.
Permissions	• Grant and revoke permissions on a procedure to users, groups, or roles. Choose the **Grant** option to allow the grantee to further grant permissions to other users. Select an object in the table of permissions, and click **Properties** to view the object properties. • See *Granting Permissions on Database Objects* on page 313 and *Revoking Permissions on Database Objects* on page 314.
Referenced By	• View the name, type, owner, and properties of objects that this procedure references. Click an object, then click **Properties** to view the object properties.
References	• View the name, type, owner, and properties of objects that this procedure references. Click an object, then click **Properties** to view the object properties.

See also
- *Extended Stored Procedures* on page 297
- *Creating a Stored Procedure* on page 312
- *Creating a SQLJ Procedure* on page 315

Granting Permissions on Database Objects
Grant permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.

---

2. Click the **Name** field of the object on which to grant permissions.

3. Click the drop-down arrow and select **Properties**.

4. Select **Permissions**.

5. Click **Grant** to grant access permissions for the selected object.
   You see the Grant Permission wizard.

6. Select the type of grantee:

   • **Users**
   • **Groups**
   • **Roles**

7. Select one or more grantees.

8. Select the types of permissions allowed for the selected grantees.

9. (Optional) Click **Summary** to verify your selected options.

Revoking Permissions on Database Objects

Revoke permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.

2. Click the **Name** field of the object on which to revoke permissions.

3. Click the drop-down arrow and select **Properties**.

4. Select **Permissions**.

5. Select the grantee, then click **Revoke** to revoke access permissions to the object.
   In the Revoke Permissions wizard, each type of permission and the current granted permissions are shown in cells.

6. Choose one of:

   • **Revoke all permission.**

   • Individual cells to revoke the currently granted permissions. The cell changes to show an "x", indicating that the permission type is no longer granted.

7. (Optional) Click **Preview** to see the SQL statements for your command.

Generating a DDL Script

Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:

- **Compiled Objects**
- **Schema Objects**
- **Security**
- **Space Management**

2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.

3. Click in the Name field of the specific object for which you want the DDL script.

4. Click the drop-down arrow, then select the option to generate DDL.

You can save the DDL in an external file on your local file system.

Deleting a Procedure
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your procedure.

3. Click the Name field of the procedure you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the procedure.

6. Confirm the deletion.

7. Click **Finish**.

Manage SQLJ Procedures
Create, delete, modify , and administer SQLJ procedures using the Sybase Control Center Administration Console.

SQLJ procedures are named collections of SQLJ statements and flow control statements. A stored procedure that performs a **select**, **execute**, or data modification command must have the same owner as the object acted upon.

A system administrator, a database owner, or a user or group with **create procedure** permission can create a stored procedure.

See also
- *Manage SQLJ Functions* on page 302

Creating a SQLJ Procedure
Create a SQLJ procedure using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects > Procedures**, then select **Stored Procedures**.

2. Click the drop-down arrow and select **New**.
   You see the Create SQLJ Procedure wizard.

3. On the Introduction screen, select the server, database, and owner of the new procedure.

4. Enter the name of the procedure.

5. Specify the external name, which identifies the Java method, class, and an optional package name.

6. On the SQL Properties screen, select:
   - Modifies SQL data – indicate that the Java method invokes SQL operations and modifies SQL data in the database.
   - Dynamic result set – set the number of rows returned. The default number of returned rows is 1.
   - Deterministic option – include the keywords deterministic or not deterministic for compatibility with the SQLJ standard. However, Adaptive Server does not make use of this option.

7. On the SQL Editor screen, provide the SQL statements for the procedure. Ensure that all objects referenced by the procedure exist in the database.

8. (Optional) Click **Preview** to see the SQL statements for your command.

9. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Stored Procedure* on page 312
- *Extended Stored Procedures* on page 297
- *Stored Procedure Properties* on page 312
- *Creating a SQLJ Function* on page 302
- *SQLJ Function Properties* on page 303
- *SQLJ Procedure Properties* on page 316

SQLJ Procedure Properties

Use the SQLJ procedure Properties wizard to access and modify information on SQLJ procedures.

Click your SQLJ procedure and select **Properties** to initiate the Properties wizard.

Wizard Option	SQLJ Procedure Properties
General	• View the name, type, database, owner, creation date, and group number of the procedure.
SQL	• View the SQL statements for creating the procedure.

Wizard Option	SQLJ Procedure Properties
Parameters	• View the name, type, mode, and order of all the procedure parameters. The mode value indicates whether it is an input or an output parameter. The order is a numeric value that indicates the place of the parameter in the list of parameters. • To change the parameters, change the definition of the stored procedure by dropping and re-creating the procedure.
Permissions	• Grant and revoke permissions on a procedure to users, groups, or roles. Choose the **Grant** option to allow the grantee to further grant permissions to other users. Select an object in the table of permissions, and click **Properties** to view the object properties. • See *Granting Permissions on Database Objects* on page 313 and *Revoking Permissions on Database Objects* on page 314.
Referenced By	• View the name, type, owner, and properties of objects that are referenced by this procedure. Click an object, then click **Properties** to view the object properties.
References	• View the name, type, owner, and properties of objects that this procedure references. Click an object, then click **Properties** to view the object properties.

See also

Granting Permissions on Database Objects
Grant permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.
2. Click the **Name** field of the object on which to grant permissions.
3. Click the drop-down arrow and select **Properties**.
4. Select **Permissions**.
5. Click **Grant** to grant access permissions for the selected object.

You see the Grant Permission wizard.

6. Select the type of grantee:

- **Users**
- **Groups**
- **Roles**

7. Select one or more grantees.

8. Select the types of permissions allowed for the selected grantees.

9. (Optional) Click **Summary** to verify your selected options.

Revoking Permissions on Database Objects
Revoke permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.

2. Click the **Name** field of the object on which to revoke permissions.

3. Click the drop-down arrow and select **Properties**.

4. Select **Permissions**.

5. Select the grantee, then click **Revoke** to revoke access permissions to the object.
   In the Revoke Permissions wizard, each type of permission and the current granted permissions are shown in cells.

6. Choose one of:

- **Revoke all permission.**

- Individual cells to revoke the currently granted permissions. The cell changes to show an "x", indicating that the permission type is no longer granted.

7. (Optional) Click **Preview** to see the SQL statements for your command.

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
- **Compiled Objects**
- **Schema Objects**
- **Security**
- **Space Management**

2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.

3. Click in the Name field of the specific object for which you want the DDL script.

4. Click the drop-down arrow, then select the option to generate DDL.

You can save the DDL in an external file on your local file system.

Deleting a SQLJ Procedure
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your SQLJ procedure.

3. Click the Name field of the SQLJ procedure you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the SQLJ procedure.

6. Confirm the deletion.

7. Click **Finish**.

Processes

Monitor Adaptive Server processes.

Monitor Processes
Monitor processes using the Sybase Control Center Administration Console.

Identifying Resource-Intensive Processes
Find the user processes that are consuming the most system resources on the selected Adaptive Server.

You can choose a system resource (CPU, disk I/O, incoming network traffic, or outgoing network traffic) and display information about the user processes on the selected server that are using the chosen resource most intensively. For each system resource, you can rank the processes by cumulative or most recent activity values. Each bar in the graph represents the value of the selected metric for a process.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Processes**.
The **All Processes** tab is selected. A bar graph shows the five user processes that are using the most cumulative CPU cycles.

3. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.

4. Use the menu to the right of the bar graph to change the system resource. You can choose:

   - **CPU Cumulative** – cumulative CPU activity since the Adaptive Server started or the counter wrapped.
   - **CPU Activity** – CPU activity, per second, since the last screen refresh.
   - **Disk I/O Cumulative** – cumulative disk I/O since the Adaptive Server started or the counter wrapped.
   - **Disk I/O Activity** – I/O activity per second since the last screen refresh.
   - **Incoming Network Traffic Cumulative** – cumulative incoming network traffic since the Adaptive Server started or the counter wrapped.
   - **Incoming Network Traffic Activity** – incoming network traffic per second since the last screen refresh.
   - **Outgoing Network Traffic Cumulative** – cumulative outgoing network traffic since the Adaptive Server started or the counter wrapped.
   - **Outgoing Network Traffic Activity** – outgoing network traffic per second since the last screen refresh.

5. (Optional) Move your mouse over a bar in the graph to display the server process ID (SPID) and the value of the selected system resource metric for the process.

6. (Optional) Select **Only display user processes below** to filter out system processes, and only display user process information.

7. (Optional) Click a bar in the graph to highlight information about that process in the table below.

8. (Optional) Click a bar in the graph or a row in the table to display information about that process in the Details, SQL, and Wait Events tabs at the bottom of the screen.

See also
- *Displaying Wait Events for a Process* on page 324
- *Displaying the SQL Query Associated with a Process* on page 323
- *Identifying the Lead Blocker in a Chain* on page 322
- *Identifying Blocked Processes and Blocking Processes* on page 321
- *Terminating Blocking Processes* on page 322
- *Process Statistics and Details* on page 324

Identifying Blocked Processes and Blocking Processes
Find user processes that are blocked and the processes that are blocking them.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Processes**.
3. Click the **All Processes** tab.
4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
5. Check the table below the bar graph for rows highlighted in red, and with a lock icon, which indicate blocked processes.

   The Blocked by SPID column identifies the blocking process. Blocking processes are also shown in the table, highlighted in yellow.
6. Click a red table row to display information about the blocked process in the the Details, SQL, and Wait Events tabs at the bottom of the screen.
7. Click the row for the blocking process (yellow) to display its information in the tabs.

   Note: Identifying the lock held by a blocking process is not always straightforward. For example, the blocking process does not necessarily hold a page lock; it might hold a table lock. For this reason, Sybase Control Center shows the lock request process that is blocking another process, not the blocking lock.
8. Click the **Blocked Processes** tab to display additional information about blocked processes, including details about the lock, the row number, the page number, and the lock configuration options.

   For non-clustered servers, Sybase Control Center displays the number of free locks, active locks, and maximum used locks. For clustered servers, Sybase Control Center displays these values on a per instance basis, and you must select **Display Lock Configuration** to display these values.

Next
For information on handling blocked processes, see the locking reports chapter of the Adaptive Server *Performance and Tuning Series: Locking and Concurrency Control*.

See also

Terminating Blocking Processes
Terminate a blocking process from the Processes window.

Sybase Control Center allows you to terminate a blocking process.

1. In the Adaptive Server monitor, select **Processes**.
2. Select the blocking process, or a set of blocking processes, that you wish to terminate.
3. Right-click the selected row and select **Terminate.**
4. Select **Yes** to relay the terminate request to the Adaptive Server, or **No** to close the dialog box without performing the terminate operation.

See also
* *Displaying Wait Events for a Process* on page 324
* *Displaying the SQL Query Associated with a Process* on page 323
* *Identifying the Lead Blocker in a Chain* on page 322
* *Identifying Resource-Intensive Processes* on page 319
* *Identifying Blocked Processes and Blocking Processes* on page 321
* *Process Statistics and Details* on page 324

Identifying the Lead Blocker in a Chain
Find a process that is blocking several other processes.

When Process A blocks Process B, which blocks Process C—and so on—the blocking processes form a chain.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Processes**.
3. Click **Blocked Processes**.

    You see a table with information about blocked and blocking processes, including the lock requests on the basis of which processes are blocked.
4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
5. The table on the Blocked Processes tab has an entry for each lead blocker; click the arrow to expand the entry and show all the blocked processes in the chain.

    Note: Identifying the lock held by a blocking process is not always straightforward. For example, the blocking process does not necessarily hold a page lock; it might hold a table lock. For this reason, Sybase Control Center shows the lock request process that is blocking another process, not the blocking lock.

6. Select a process to populate the tabs at the bottom of the screen with information about that process.

7. Click the tabs to see:

Option	Description
Details	Details about the selected process, including the initiating program, transaction information, and network statistics
SQL	The SQL statement and query plan for the selected process
Wait Events	Information about wait events for the selected process, including number of waits, wait times, and wait descriptions

8. (Optional) Click the **All Processes** tab at the top of the window.
   In the table below the bar graph, rows highlighted in yellow are blocking processes. Rows highlighted in red are blocked processes.

Next

For information on handling blocked processes, see the *Adaptive Server Performance and Tuning Series: Locking and Concurrency Control*. You can find Adaptive Server documentation on the Sybase Product Documents Web site at *http://sybooks.sybase.com*.

See also

- *Displaying Wait Events for a Process* on page 324
- *Displaying the SQL Query Associated with a Process* on page 323
- *Identifying Resource-Intensive Processes* on page 319
- *Identifying Blocked Processes and Blocking Processes* on page 321
- *Terminating Blocking Processes* on page 322
- *Process Statistics and Details* on page 324

Displaying the SQL Query Associated with a Process

See the SQL statement and query plan for an Adaptive Server user process.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Processes**.

3. Click either the **All Processes** tab or the **Blocked Processes** tab.

4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.

5. To select a process, click a row in the table.

6. At the bottom of the window, click the **SQL** tab.

The tab shows the SQL query and query plan for this process.

See also
- *Displaying Wait Events for a Process* on page 324
- *Identifying the Lead Blocker in a Chain* on page 322
- *Identifying Resource-Intensive Processes* on page 319
- *Identifying Blocked Processes and Blocking Processes* on page 321
- *Terminating Blocking Processes* on page 322
- *Process Statistics and Details* on page 324

Displaying Wait Events for a Process
Get information about wait events that are affecting an Adaptive Server process.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Processes**.
3. Click **All Processes**.

   Each row of the table below the bar graph displays information about a process.
4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
5. To select a process, click a row in the table.
6. Click **Wait Events**.

For more information on wait events, see the wait events chapter in the Adaptive Server *Performance and Tuning Series: Monitoring Tables* guide.

See also
- *Displaying the SQL Query Associated with a Process* on page 323
- *Identifying the Lead Blocker in a Chain* on page 322
- *Identifying Resource-Intensive Processes* on page 319
- *Identifying Blocked Processes and Blocking Processes* on page 321
- *Terminating Blocking Processes* on page 322
- *Process Statistics and Details* on page 324

Process Statistics and Details
Interpret Adaptive Server process information.

Lock icons in the SPID column of the Processes table identify processes that are blocked (a grayed-out lock) or blocking (a gold lock) other processes. Other columns of the Processes table include the family ID which is the parent SPID value, processes blocked by SPID, CPU activity, CPU cumulative activity, disk I/O activity and disk I/O cumulative activity.

On the Blocked Processes tab, Sybase Control Center shows the lock request process that is blocking another process, not the blocking lock itself. A yellow warning icon appears on the Blocked Processes tab label when there are any blocked processes.

Table 39. Color indicators in the Processes table

Color	Indicates that the process is...
Blue	Selected
Green	Executing a query
Yellow	Blocking another process
Red	Blocked by a lock held by another process

Table 40. Tabs

Details	Displays information about the selected process, including program name, current SQL command, client machine IP address and name, transaction name and start time, physical and logical reads, and CPU utilization.
SQL	If the selected process is active, displays the SQL statement and query plan for the query that the process is executing.
Wait Events	Displays a list of all events the selected process has waited for. Wait events are internal Adaptive Server states that represent conditions that cause a process to stop. Common wait events include waiting: • On the Adaptive Server scheduler runnable queue for a CPU to become available • For disk I/O to complete • For a lock on a table to be released

See also

Replication Agents

Set replication parameters and monitor repAgent threads in Adaptive Server.

In Adaptive Server, there is one RepAgent thread for each database from which data is replicated. The RepAgent thread reads the transaction log of a primary database. It sends the

transaction as Log Transfer Language (LTL) commands for replicated tables and replicated stored procedures to the primary Replication Server, which converts the LTL commands into SQL and applies the SQL to the replicate database.

For more information about replication, see the *ASE Replicator Users Guide*.

Monitor Replication Agents

Monitor RepAgent threads for replicate databases on the selected Adaptive Server.

Monitoring RepAgent Threads

Display the status and transaction log details of the RepAgent threads running in Adaptive Server.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, schedule collection jobs, and verify that you have permission to perform this task.

Task

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Replication Agent**.

   Note: If Sybase Control Center does not detect a Replication Management Agent Plug-in, or detects one that is incompatible with your Adaptive Server version, the **Replication Agent** option is greyed out.

3. Select a RepAgent thread in the table.
   You see the **Log Size** and **Activity** graphs.

For more information about replication, see the *ASE Replicator Users Guide*.

See also

- *Setting Replication Parameters* on page 326
- *Replication Agent Statistics and Details* on page 327

Setting Replication Parameters

Configure replication parameters to improve server performance. Sybase Control Center for Replication allows you to configure the parameters for Replication Server, Replication Agent, connection and logical connection, route, and Adaptive Server RepAgent thread.

Replication Agent Statistics and Details
Interpret the Replication Agent information for Adaptive Server.

The Replication Agent screen displays the name, status, and controlling Replication Server of the RepAgent thread for the current Adaptive Server. The Transaction Log Details tab displays the log data of the selected RepAgent thread.

Transaction Log Details tab

Log Size	Displays a graph of the log size, in megabytes.
Activity	Displays a log count activity graph of the of the scanned and processed log records.

For more information about replication, see the *ASE Replicator Users Guide*.

See also

Rules

Manage rules using the Sybase Control Center Administration Console.

Manage Rules

Create, delete, modify, and administer rules using the Administration Console.

Creating a Rule
Create a rule using the Sybase Control Center Administration Console.

Only a database owner, or a user or group with **create rule** permission can create a rule.

1. In the left pane of Administration Console, expand **ASE Servers > Compiled Objects**, then select **Rules**.
2. Click the drop-down arrow and select **New**.
    You see the Create Rule wizard.
3. On the Introduction screen, select the server, database, and owner of the new rule.
4. Enter the name of the rule.
5. Enter the expression that is used to evaluate the data. You can use any expression that is valid in a **where** clause.
6. (Optional) Click **Preview** to see the SQL statements for your command.
7. (Optional) Click **Summary** to verify your selected options.

See also

* *Rule Properties* on page 328

Rule Properties

Use the rules Properties wizard to access information on rules and objects that they reference.

Click **Properties** on your rule to initiate the Properties wizard.

Wizard Option	Rule Properties
General	View the name, type, database, owner, creation date, and rule expression.
SQL	View the SQL statements for creating the rule.
Referenced By	View the name, type, owner, and properties of objects that referenced by this rule. Click an object, then click **Properties** to see the object properties.

See also

* *Creating a Rule* on page 327

Generating a DDL Script

Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   * **Compiled Objects**
   * **Schema Objects**
   * **Security**
   * **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a Rule

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your rule.

3. Click the Name field of the rule you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the rule.

6. Confirm the deletion.

7. Click **Finish**.

Security

Use server-level security features such as logins, login profiles, and roles, and database-level security features such as encrypted columns, users, and groups.

Manage Encryption Keys

Database columns can be encrypted with keys that are created with user-defined or login passwords.

Encryption Keys

In each database, you can create a key that is used to encrypt columns. Creating a key on each database minimizes cross-database key integrity problems. Such key problems can happen in distributed systems, particularly when you are dumping and loading, or mounting and unmounting databases.

Note: You can create encryption keys only you have:

- System security officer or key custodian role
- Permissions to execute **create encryption key**

If you are a key owner, allow other users to access encryption keys by either:

- Creating an encryption key with a user-defined password and sharing it with each user who accesses key-encrypted data, or
- Giving each user a copy of the base encryption key, and allowing him or her to change the key-copy password.

Encryption Keys with User-Defined Passwords

Using encryption keys with user-defined passwords creates a highly secure system where even database owners and system administrators cannot access encrypted data. You can also require that the key encryption method itself use a user-defined password.

Adaptive Server provides recovery for lost base-key passwords.

When data is encrypted, system security officers, key-custodians, and users with permission to create encryption keys can also create base keys. System security officers can also grant base key creation permission to users with no other permissions.

Whoever creates the base key is the "key owner." To control access to encrypted data, only key owners and system security officers can change the base key password.

Encryption Keys with Login Passwords

To avoid excess passwords, you can authorize users to access encrypted data using their login password. Using login passwords to access key-encrypted data:

- Enables access to encrypted data without users explicitly supplying passwords.
- Involves fewer passwords to track.
- Reduces the need for the key custodian to replace lost passwords.

Key Copies

Key owners can allow data access to other users by making copies of the base key—called key copies. A key copy is an additional password for the base key that can be changed as soon as it is assigned to a user, or key-copy owner. Only the key copy owner can change the key-copy password.

You can make key copies for designated users if you are the base key owner or a system security officer. Key copies of the base key are not new keys themselves; they are additional passwords for the base key. Key copy assignees should change their user-defined password for the key copy as soon as the key copy is assigned to them.

The key copy is encrypted with the login password as soon as the assignee logs in and accesses the key copy.

Note: The base key can be encrypted by the system encryption password or a user-defined password. Key copies can be encrypted by a login password or by a user-defined password. The recovery key copy can only be encrypted by a user-defined password. Keys encrypted with the system encryption password cannot have key copies.

Key recovery requires you to create a special key copy designated for the recovery of the base key. This is called the recovery key. If you lose your password, use the recovery key to access the base key.

Creating a System Encryption Password

The system security officer (SSO) creates the default system encryption password. Adaptive Server encrypts keys using the Advanced Encryption Standard (AES) algorithm. The system encryption password is encrypted and stored in the database.

If system encryption passwords are too short or easy to guess, the security of encryption keys can be compromised.

Note: Keys encrypted using the system encryption password cannot have key copies.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.
2. Click **System Encryption Passwords**.

3. Click the drop-down arrow and select **New**.
   You see the Add System Encryption Password wizard.

4. Select the Adaptive Server and database containing the keys that the system encryption password encrypts.

5. Enter the new password, and confirm it.

6. (Optional) Click **Summary** to verify your selected options.

See also

* *Modifying and Deleting a System Encryption Password* on page 331

Modifying and Deleting a System Encryption Password

You can change or delete the system encryption password using Sybase Control Center.

Only the system security officer (SSO) can change the system encryption password.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.

2. Click **System Encryption Passwords**.

3. Find the row containing the database in which you want to change the encryption password, and right-click the **Name** field.

4. (Optional) To change the system encryption password:
   a) Select **Change Password**.
   b) Enter the old and new passwords, and confirm the new password.

5. (Optional) To delete the system encryption password, select **Delete**, and confirm it on the next screen.

See also

* *Creating a System Encryption Password* on page 330

Creating a Master Key

Create the master key for the database.

Adaptive Server versions 15.7 and later support master-key functionality. The master key:

* Is a database-level key, created by a user with sso_role or keycustodian_role.
* Is used as a key encryption for user-defined encryption keys.
* Replaces the system encryption password as the default key encryption key for user-defined keys.

Note: Sybase does not recommend that you create system encryption passwords after you have created master keys.

- Can be used with the dual master key as a composite key to provide dual control and split knowledge for all user-created keys. Alternatively, the master key can be used as a composite key with a column encryption key's explicit password.
- Can be altered to add key copies. Master key copies provide access to the dual-master key for unattended start-up, to support recovery of the master key, and to allow users other than the base-key owner to set the encryption password.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.

2. Click **Master Keys**.

3. Select **New**.
   You see the Add Master Key wizard.

4. Select the Adaptive Server and database where the encryption key is being defined.

5. Enter a password for the master key and confirm it.

6. (Optional) Click **Summary** to verify your selected options.

See also
- *Modifying, Regenerating, and Deleting a Master Key* on page 334
- *Dual Control and Split Knowledge* on page 335
- *Master Key Properties* on page 335

Creating a Column Encryption Key
Create a column encryption key using a specified encryption method.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.

2. Click **Column Encryption Keys**.

3. Select **New**.
   You see the Add Column Encryption Key wizard.

4. Select the Adaptive Server and database where the encryption key is being defined.

5. Select the key owner.

6. Enter an encryption key name.

7. Select these parameters for the Advanced Encryption Standard (AES) encryption algorithm:
   - Key length – choose 128, 192, or 256, depending on the level of security you need.
   - Default key – select this key as the default key to allow users to create encrypted columns without specifying the key.
   - Encryption method – select one of these:
     - User-defined password – provide a password and confirm it. Select **With dual control** to encrypt with the master key and a user-defined password.

- Master key – enable encryption using the master key. Select **With dual control** to encrypt with the master key and a user-defined password.

  Note: If you select dual control, the master key must already exist in the database and you must supply the master key password.

- System encryption password – enable encryption using the system encryption password.

8. Select the initialization vector to be either random (the default) or null. Use the initialization vector padding to increase the security of encrypted data by increasing the cryptographic variance of the cipher text.

9. Select the pad value to be either random or null (the default). If pad is set to random, Adaptive Server uses datatype padding when the length is less than one block.

10. (Optional) Click **Summary** to verify your selected options.

See also
- *Modifying and Deleting a Column Encryption Key* on page 333
- *Executing SQL Statements* on page 193
- *Modifying, Regenerating, and Deleting a Master Key* on page 334
- *Column Encryption Keys Properties* on page 336

Modifying and Deleting a Column Encryption Key

Change the encryption key, with the option of adding dual control.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.

2. Click **Column Encryption Keys**.

3. Find the row containing your column encryption key and right-click the **Name** field.

4. (Optional) To change the encryption key:

   a) Select **Change Password**.

   b) Select one of these methods to encrypt the key:
      - Using a user-defined password – enter the old and new passwords, and confirm the new password. If **With dual control** is selected, both the master key and user-defined password are used to encrypt the column encryption key. Both master and dual-master keys are used for encryption if both exist in the database.
      - Using the master key – the server encrypts your key using the master key.
      - Using the system encryption password – the server encrypts your key using the system encryption password. Before choosing this option, ensure that a system encryption password exists in your database.

   (Optional) Select **With dual control** to use master and dual-master keys to control the security of your column encryption keys.

5. (Optional) To set the column encryption key, select **Supply Password**, and confirm it on the next screen.

6. (Optional) To delete the column encryption key, select **Delete**, and confirm it on the next screen.

See also
- *Creating a Column Encryption Key* on page 332
- *Executing SQL Statements* on page 193
- *Modifying, Regenerating, and Deleting a Master Key* on page 334
- *Column Encryption Keys Properties* on page 336

Modifying, Regenerating, and Deleting a Master Key
Modify existing passwords or regenerate the master key.

Use **Change Password** when a password is compromised.

Use **Regenerate** to periodically change the key encryption keys as a good key management practice. Adaptive Server replaces the master or dual-master key with a new value and reencrypts all column encryption keys that are encrypted by the master or dual-master keys.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.

2. Click **Master Keys**.

3. Find the row that contains your master key and right-click the **Name** field.

4. (Optional) To change the master key:
   a) Select **Change Password**.
   b) Enter the old and new passwords, and confirm the new password.

   Note: If a key has key copies, you cannot modify the key to encrypt it with the system encryption password.

5. (Optional) To set the master key, select **Supply Password**, and confirm it on the next screen.

6. (Optional) To delete the master key, select **Delete**, and confirm it on the next screen.

7. (Optional) To regenerate the master key:
   a) Select **Regenerate**
   b) Enter the old and new passwords, and confirm the new password.

See also
- *Creating a Master Key* on page 331
- *Dual Control and Split Knowledge* on page 335
- *Master Key Properties* on page 335

Master Key Properties

Properties of master keys and key copies.

Click **Properties** on your master key to initiate the Master Key Properties wizard to modify these properties:

Wizard Option	Encryption Properties
General	• **Name** – change the name of the master key. • **Change Owner** – change the owner of the master key. • Key recovery – indicates if this key has a recovery copy. • **Master key startup file** – specify if the master key has an automatic start-up copy. The automatic startup copy is used to access the master or dual-master keys when a server, configured for **automatic master key access**, is started.
Key Copies	• Assignees and other information about keys – lists the types of passwords and assignees for the key, and information about whether the key is recoverable.

See also
* *Dual Control and Split Knowledge* on page 335
* *Creating a Master Key* on page 331
* *Modifying, Regenerating, and Deleting a Master Key* on page 334

Dual Control and Split Knowledge

Adaptive Server version 15.7 and later provide dual-control and split-knowledge encryption.

Adaptive Server allows you to use a combination of system keys at the database level called the master key and the dual master key. You must have **sso_role** or **keycustodian_role** to create the master key and dual master key. The master key and the dual master key must have different owners.

With Sybase Control Center, you can provide passwords for the master keys using the **Supply Password** option for encryption keys. You can also use the **Execute SQL** option of the Administration Console to provide the password using SQL. The passwords to both these keys are not stored in the database.

Master and dual master keys act as key encryption keys (KEKs), and are used to protect other keys, such as column encryption keys and service keys. Once created, master and dual master keys become the default protection method for column encryption keys. There can only be one master and one dual master key for a database.

The dual master key is needed only for dual control of column encryption keys. Once the master key is created, it replaces the system encryption password as the default key encryption key for user-created keys.

A composite key comprising the master key and dual master key provides dual control and split-knowledge security for all user-created keys. Alternately, a composite key may also be created using the master key and the column encryption key's password. When master and dual master keys are configured in a database, Adaptive Server uses the combination to encrypt passwords when you issue **create table**, **alter table** or **select into** commands specifying dual control.

See also
- *Master Key Properties* on page 335
- *Creating a Master Key* on page 331
- *Modifying, Regenerating, and Deleting a Master Key* on page 334
- *Executing SQL Statements* on page 193

Column Encryption Keys Properties
Properties of column encryption keys and key copies.

Click **Properties** on your column encryption key to initiate the Column Encryption Key Properties wizard to modify these properties:

Wizard Option	Encryption Properties
General	• **Name** – change the name of the encryption key.
	• **Change Owner** – change the owner of the encryption key.
	• **Default** – select this key as the default key to allow users to create encrypted columns without specifying the key.
	• Key recovery – indicates if this key has a recovery copy.
	• **Init vector** – use the initialization vector padding to increase the security of encrypted data by increasing the cryptographic variance of the cipher text.
	• **Pad** – if pad is set to random, Adaptive Server uses datatype padding when the length is less than one block.

Wizard Option	Encryption Properties

Key Copies	• Assignees and other information about keys – list the types of passwords and assignees for the key, and information about whether the key is recoverable.
	• Assignees and other information about key copies – click the **Type of password** field for your key copy, and select **Properties** to see the general properties of key copies, including database, owner, assignee, and type of password.
	• Key-copy management –
	  • Create a new key copy by clicking on **New**. See *Creating Key Copies* on page 337.
	  • Delete a key copy by clicking on the **Type of password** field for your key copy, and selecting **Delete**.

Object Permissions	• Grantees and other object information – list the grantees and grantee types for the key, and information whether select is granted. • Permissions – click **Grant** or **Revoke** to modify permissions to users, groups, or roles. See *Granting Encryption Permissions to Roles, Users and Groups* on page 338.
Dependencies	• Encrypted columns – list the columns encrypted by this key, and their databases and tables.

See also
- *Creating a Column Encryption Key* on page 332
- *Modifying and Deleting a Column Encryption Key* on page 333
- *Executing SQL Statements* on page 193
- *Modifying, Regenerating, and Deleting a Master Key* on page 334

Creating a Key Copy
Create key copies specifying an encryption method.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.
2. Click **Column Encryption Keys**.
3. Find the row containing your column encryption key and click the **Name** field.
4. Click the arrow and select **Properties**.
   Sybase Control Center displays the Properties wizard.
5. Click **Key Copies**.

6. Select **New**.
   You see the Add Key Copy wizard.

7. Enter the password for the base key.

8. Enter the assignee for the key copy. The assignee cannot be the key owner.

9. (Optional) Designate this key copy as the recovery key copy.

10. Select one of these encryption methods for the key copy:
    - User-defined password – provide a password and confirm it.
    - Login password – enable encryption using the login password.
    - System encryption password – enable encryption using the system encryption password.

11. (Optional) Click **Summary** to verify your selected options.

Granting Encryption Permissions to a Role, User, or Group
Grant permission to access the encryption key.

1. In the left pane of the Administration Console, expand **ASE Servers > Security > Encryption Keys**.

2. Click **Column Encryption Keys**.

3. Find the row containing your column encryption key and click the **Name** field.

4. Click the arrow to display a menu, and select **Object Permissions**.

5. Click **Grant** to allow other users, groups, or roles to access the encryption key.
   You see the Grant Permission wizard.

6. Select one of users, groups, or roles for access to the encryption key.

7. Select the grantee from the list of possible users, groups, or roles.

8. Select the key permissions to be granted.

9. (Optional) Click **Summary** to verify your selected options.

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**
   - **Space Management**

2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**

Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.

3. Click in the Name field of the specific object for which you want the DDL script.

4. Click the drop-down arrow, then select the option to generate DDL.

    You can save the DDL in an external file on your local file system.

Deleting an Encryption Key

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your encryption key.

3. Click the Name field of the encryption key you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the encryption key.

6. Confirm the deletion.

7. Click **Finish**.

Manage Login Profiles

Login Accounts can be managed with login profiles that define attributes for individual logins, a subset of logins, or all logins.

Login Profiles

A login profile is a collection of attributes specific to login accounts. You can manage attributes of login accounts by creating login profiles and associating the profile with a login account. You can manage attributes of a large number of login accounts by defining login profile as: the default for all login accounts, a subset of login accounts, or individual login accounts.

Note: Login profiles are supported in multiple-server or a single-server environments.

Options for Creating or Modifying Login Profiles

When you create or modify a login profile, you can :

- Assign a default database and default language
- Assign an authentication mechanism
- Track the last login
- Define a stale login inactivity period
- Execute a login script

Login Profile Precedence Rules

Login profiles attributes are associated with login accounts using this precedence:

1. Attribute values from a login profile bound to the login

2. Attribute values from a default login profile
3. Values that have been specified using **sp_passwordpolicy** under these circumstances:
    * A default login profile does not exist
    * A login profile has not been defined and bound to the account
    * The login profile is set to be ignored
4. The default value for the attribute

Creating a Login Profile
Create a login profile using the Sybase Control Center Administration Console.

Note: Only a system security officer can create, modify, or delete login profiles.

1. In the Perspective Resources view, select the servers on which the login profile is to be created and select **Administration Console**.

    Note: Login profiles are supported in multiple-server or single-server environments.

2. Select **ASE Servers > Security > Login Profiles**.
3. Select **New**.
   The Add Login Profile wizard appears.
4. On the Introduction window, select the individual servers on which to create the login profile, or select to create multiple login profiles on all available servers.

   If you create multiple login profiles on different servers, the names of the login profiles are the same on all servers. However, the login profiles on each server can have different default databases, default languages, or authentication mechanisms.

5. On the Login Profile Name window, enter the name of the login profile to create.
6. (Optional) Select:
    a) **With attributes derived from an existing login account** – to transfer existing login account values to a new login profile.
    b) **As default for all login accounts** – if you want the new login profile to be the default for all login accounts on the selected servers.
7. (Optional) On the Default Database window, click **Specify default database** to choose a database to be used as the default for the login profile. Select one of:

    * **Use common default database for the login profile on all servers**, then select the default database.
      The list of available of databases is based the databases that are common on all servers that have been selected. No servers to select from indicates that there are no common databases available.
    * **Use default database for the login profile on individual server**, then select the default database .

8. (Optional) On the Default Language window, select **Specify default language** to choose a language to be used as the default for the login profile. Select one of:

- **Use common default language for the login profile on all servers**, then select the default language.

  us_english is the default language, but you can install locale character sets. The additional installed languages and the default language constitute the list of available languages.
- **Use default language for the login profile on individual server**, then select the default language.

9. (Optional) On the Authentication window, select **Specify authentication** to choose an authentication mechanism for the login profile. Select one of:

   - **Use common authentication for the login profile on all servers** , then select the authentication mechanism.
   - **Use default authentication for the login profile on individual server**, then select the authentication mechanism.

   If you select **ANY** (the default) as the authentication mechanism, Adaptive Server checks for a defined external authentication mechanism and uses it, if it exist. Otherwise, the ASE mechanism is used. The default is ANY.

10. (Optional) On the More Options window, choose from:

Table 41. Options for Login Profiles

Option	Description
Track last login	Specify whether to enable last login updates. The default is to track the last login.
Stale login inactivity period	Specify the length of time a login account can remain inactive before it is locked due to inactivity.
Login script	Specify a script to be invoked on login.

11. (Optional) Click **Summary** to verify your selected options.

12. Click **Finish** to create the login profile.

Login Profile Properties
View or modify login profile property defaults using the Login Properties wizard.

Note: Login profiles can be supported in a multiple-server or a single-server environment.

Wizard Option	Login Properties

General	• Default database – if not specified, the master database.
	• Default language – if not specified, us_english is the default.
	• Authentication – specify the external authentication mechanism:
	  • ANY
	  • ASE
	  • KERBEROS
	  • LDAP
	  • PAM
	If you select **ANY** (the default) as the authentication mechanism, Adaptive Server checks for a defined external authentication mechanism and uses it, if it exist. Otherwise, the ASE mechanism is used. The default is ANY.
	• Track last login – specify whether to enable last login updates. The default is to track the last login.
	• Stale login inactivity period – specify the length of time a login account can remain inactive before it is locked due to inactivity.
	• Login script – specify a script to be invoked on login.

Logins	Displays the login accounts that are bound to the selected login profile. See *Displaying Logins Assigned to Login Profiles* on page 343.
Roles	You can add or remove roles that have been granted to the selected login profile. See *Managing Roles Granted to Login Profiles* on page 342.

Managing Roles Granted to a Login Profile

Use the Login Profile properties view to add roles to login profiles or remove roles from login profiles.

1. From the Administration Console, select **ASE Servers > Security > Login Profiles**.
2. Select a login profile on which to grant or remove roles, or view currently granted roles.
3. Select **Properties**.
4. On the Login Profile Properties window, select **Roles**.

   The Name field shows the roles that are granted to the selected login profile.

   • To grant roles, click **Add** and select one or more roles. Optionally, select **Active By Default** to indicate the role must be automatically activated on login.
   • To remove a role, select the role from the name list and click **Remove**.

Deleting a Login Profile

Use the Delete command to drop a login profiles.

1. In the Perspective Resources window, select the servers on which the login profile is defined, then select **Administration Console**.

2. Select **ASE Servers > Security > Login Profiles**.

3. Click the login profile to delete, then select **Delete**.

4. (Optional) From the Confirm Delete Login Profile window, select **Drop with override** to forcefully drop login profiles that are bound to login accounts. Login accounts that are bound to the deleted login profile will be reassigned to the default login profile.

5. (Optional) Select **Preview** to view the properties of the login profile.

Displaying Logins Assigned to a Login Profile
Display login profiles and the bindings of login accounts to login profiles.

1. In the Perspective Resources window, select the servers on which the login profiles have been defined and select **Administration Console**.

2. Select **ASE Servers > Security > Login Profiles**.

3. Choose a login profile and select **Properties**.

4. Select **Login**.

   You can find additional details about the login profile in the General and Roles properties options.

Transferring Login Attributes to a Login Profile
Use attributes of an existing login account to create a login profile.

1. In the Perspective Resources window, select the servers on which the login profile is to be created, then select **Administration Console**.
   The servers you choose here are used for some of the options in the steps below.

2. Select **ASE Servers > Security > Login Profiles**.

3. Select **New**.
   You see the Create Login Profile wizard.

4. On the Introduction window, select the individual servers on which to create the login profile, or create multiple login profiles on all available servers.

5. Enter the name of the new login profile.

   If you are creating a login profile on different servers, a login profile is created on each server. All the login profiles will have the same name, but can have different attributes.

6. Select **With attributes derived from an existing login account**.

7. (Optional) Click **As default for all login accounts** to set the new login profile or profiles as the default for all login accounts on the selected servers.

8. (Optional) On the Select Login Name window, select the login account from which to derive attributes. Select one of:

- **Use common login on all servers**, then select the login account.

  The list of available of login accounts is based on login accounts that are common on all of the servers that have been selected. An empty list indicates that here are no common login accounts.

  A login profile is created on each server based on attributes of the common login account. The name of the login profile is the same on each server.
- **Use login account on individual server**, then select the login accounts.

  A login profile is created on each server based on attributes from different login accounts. The name of the login profile is the same on each server.

Manage Logins

Each Adaptive Server user must have a login account identified by a unique login name and a password.

Note: Only a system security officer can create, modify, or delete login accounts.

Login Accounts

Users must have a login account with a unique name and password to access a server. When a login account is added to one or more servers, the account is given a unique system user ID, which identifies the users regardless of the server being used. Once a login account is created, a user account is created for users to access individual database. Login profiles can be associated with a login account to manage attributes such as the default database, default language, authentication mechanism, tracking the login, setting inactivity periods, and invoking login scripts.

Options for Managing Login Accounts

- Grant roles to logins
- Map users to logins
- Map client users to logins
- Assign login profiles to login accounts (in Adaptive Server version 15.7 and higher).
- Lock login accounts
- Expire login accounts
- Set the number of failed logins
- Configure passwords parameters at the server level
- Change the password for a specific login

Creating a Login

The system security officer (SSO) creates a login account for each user.

Note: Only a systems security officer can create, modify, or delete login accounts.

1. In the Perspective Resources window, select the servers on which the login account is to be added, then select **Administration Console**.

The servers you choose here are used for some of the options in the steps below.

2. Select **ASE Servers > Security > Logins**.

3. Select **New**.
   You see the Add Login wizard.

4. On the Introduction window, select the individual servers on which to create a login account, or select all servers.
   Whether you can associate a login profile to the new account depends on the software version the server is running. Login profiles, which support the management of login attributes such as default database, default language, and authentication mechanism are available only in software versions 15.7 and later.

5. On the Login Name window, enter a name for the login account you want to create, then enter a password.

6. (Optional) Enter a full name for the account.
   Specifying a full name for the account allows easier identification of the account owner.

7. On the Login Profile window, choose whether to use a login profile and designate the servers on which the login profile will be used. Choose one of:

   - **Ignore login profile in creating login**.
   - **Use common login profile for the login on the servers with version 15.7 and above**. This option is available only when an existing login profile is available.
   - **Use different login profile for the login on individual server**. This option is available only when an existing login profile is available.

8. (Optional) On the Default Database window, select a default database for the login account. Choose one of:

   - **Use common default database for the login on all servers**
   - **Use default database for the login on individual server**

   If the software version is 15.7 or later and a login profile is assigned to the login, you cannot specify a default database.

9. (Optional) On the Optional Parameters window, select the default language for the new login account.

   If the software version is 15.7 or later and a login profile is assigned to the login, you cannot specify a default language.

10. (Optional) Select the authentication mechanism for the new login account.

11. On the Database Access window, select the databases that the login account can access.

    The step adds a user account of the same name as the login to the selected database.

12. (Optional) Click **Summary** to verify your selected options.

Login Properties

Use the Administration Console option on your Adaptive Server to view or modify login properties.

The Login Properties wizard allows you to perform these actions:

Wizard Option	Login Properties
General	- Full name – allows for easier identification name for the login account. - Default database – if not specified, the master database. - Default language – if not specified, us_english is the default. - Authentication – specify the external authentication mechanism: - ANY - ASE - KERBEROS - LDAP - PAM If you select **ANY** (the default) as the authentication mechanism, Adaptive Server checks for a defined external authentication mechanism and uses it, if it exist. Otherwise, the ASE mechanism is used. The default is ANY. - Temp DB binding – binds logins to a temporary database of the default temporary database group.
Parameters	- Invalid password or NULL – specify new password for the login account. - Password has expired – the account owner must change the login password. - Account is locked – lock the login account. - Password last set – indicates when the password was changed. - Max failed logins – the number of login attempts allowed, after which the account is locked. - Min password length – minimum password length required for the login account. - Password expiration intervals (days) – the number of days until the password expires. - CPU time accumulated – tracks of the amount of CPU time. - I/O time accumulated – tracks the amount time spent processing input and output operations.
Databases Owned	Displays a list of databases that are owned by the specified login account. See *Displaying Login Account Properties* on page 348.

Wizard Option	Login Properties
Roles	Displays a list of roles granted to the account. You can add or remove roles that have been granted to the selected login. See *Managing Roles Granted to Login Accounts* on page 347.
Users	Displays a list of users or aliases that are bound to the account. You can add or remove users from the account. See *Managing Users Mapped to Logins* on page 347.

Managing Roles Granted to Logins

Use the Administration Console option to grant roles to logins or remove roles from logins.

1. From the Administration Console, select **ASE Servers > Security > Logins**.
2. Select the login accounts on which to grant or remove roles, or view currently granted roles.
3. Select **Properties**.
4. On the Logins Properties window, select **Roles**.

   - To grant roles, click **Add** and select one or more roles. Optionally, select Active By Default to indicate the role must be automatically activated on login.
   - To removed a role, select it from the name list, then click **Remove**.

Managing Users Mapped to Logins

Use the Administration Console option to manage the mapping of users to logins.

1. From the Administration Console, select **ASE Servers > Security > Logins**.
2. Select the login account on which to add or remove users, or view users mapped to the selected login account.
3. Select **Properties**.
4. On the Logins Properties window, select **Users**.

   Users currently mapped to the selected login account are listed in the **Name** field.

   - To add a user to the login account, click **Add** , then select one or more users.
   - To remove users from a login account, select a user and select **Remove**.
   - To see the attributes and properties assigned to a user, select **Properties**.

Configuring Login Password Properties

Use the Administration Console option to manage password properties for login accounts.

1. Select **ASE Servers > Security > Logins > Configure Login Passwords**.

2. On the Servers Selection window, select the individual servers on which to configure passwords, or select all servers.

3. Click **Configuration**.

4. Select options from the table to configure password complexity options.

5. (Optional) On the Expiration window, select **Expire login accounts** to specify that the owners of the login accounts must change the login password.

   • (Optional) To expire passwords for specific login accounts or login accounts matching specified characters, select **Expire passwords**.

   • (Optional) To expire passwords that have not been changed by a specified date, select **Expire stale passwords** and specify a cut-off date.

6. On the Lock Inactive Accounts window, check **Lock inactive login accounts** to locked accounts due to inactivity.
   To lock inactive accounts, **enable last login updates** on the Configuration screen must be checked.

7. Specify the number of days the account can remain inactive before the account is locked.

Changing a Login Password
Change passwords and parameters for login accounts using the Administration Console.

1. Select **Administration Console**.

2. Select **ASE Servers > Security > Logins**.

3. Select the login account for which to change the password, then select **Properties**.

4. Select **Parameters**.

5. (Optional) Set the **Min password length**, **Max failed logins**, and **Password expiration interval**.

6. Click **Change Password**.

7. Enter the current password for the login and the new password.

Displaying Login Account Properties
View or modify properties of login accounts using the Administration Console.

1. In the Perspective Resources view, select the servers in which the login accounts have been defined.

2. Select **Administration Console**.

3. Select **ASE Servers > Security > Logins**.

4. Select a login account, then select **Properties**.

5. Select **Properties**, then select:

   • **General** – to view or change settings for defaults.

- **Parameters** – to view or change password and login settings.
- **Databases Owned** – to view databases owned by the selected login.
- **Roles** – to view roles granted to the login.
- **Users** – to view users mapped to the login.
- **Clients** – to view clients mapped to the login.

Assigning Login Profiles to a Login

Manage attributes of login accounts by assigning a login profile to an individual login, a subset of logins, or all logins.

Note: Login Profiles are supported in software version 15.7 and higher

---

1. Select **Administration Console**.
2. From the Administration Console, select **ASE Servers > Security > Logins**.
3. Select the login account to which to assign a login profile, then select **Properties**.
4. From the Login Properties window, select **General**.
5. Unselect **Ignore login profile**.
6. Select the name of the login profile.
    The available login profiles are those that have been defined on the same server as the selected login account.

Deleting a Login

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your login.
3. Click the Name field of the login you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the login.
6. Confirm the deletion.
7. Click **Finish**.

Manage Groups

Permissions can be granted to groups to access database objects.

You must have permission from the system security officer to work with other groups.

The database owner grants and revokes group encryption permissions.

When Adaptive Server is configured to restrict decrypt permission, only the system security officer can grant decrypt permission on tables, columns, and views. When restricted decrypt permission is turned off, the system security officer or the database owner can grant decrypt permission.

Command permissions allow the group to execute **create** commands. Database owners can assign command permissions to groups in the databases they own.

> **Note:** Sybase Control Center reports only explicitly granted and revoked permissions as well as those that users obtain by belonging to a group. For example, iSybase Control Center does not report on permissions associated with a login role.

Creating a Group
Create a group in a database.

1. In the Administration Console view, select **ASE Servers > Security > Groups**.
2. Select **New**.
   You see the Add Group wizard.
3. Select the server and database in which to create a group.
4. Enter the name of the group to create.
5. (Optional) Click **Summary** to verify your selected options.

See also
* *Creating a User* on page 352
* *Transferring Ownership of a Database Object* on page 353
* *Groups Properties* on page 350
* *Users Properties* on page 354

Groups Properties
Properties of groups.

Click **Properties** on your group to initiate the Group Properties wizard to modify these database properties:

Wizard Option	Database Properties

General	• Users – to change the users in your group, use:
	    • **Add** – select a user and click **Apply** to add the user to your group.
	    • **Remove** – select a user and click **Apply**.

Wizard Option	Database Properties
Command Permissions	Permissions to create database objects – select the kind of permissions to grant the group: • Create default • Create procedure • Create rule • Create table • Create view • Create encryption key To revoke the permissions, select the permissions and click **Revoke**, then **Apply**.
Object Permissions	Permissions to access database objects – use the Grant Permission and Revoke Permission wizards to grant or revoke permissions for specific database operations such as insert, delete, update, reference, and decrypt for specific database objects such as tables, procedures, encryption keys, and so on. See *Granting and Revoking Permissions to Groups* on page 351.

See also
- *Users Properties* on page 354
- *Creating a Group* on page 350
- *Creating a User* on page 352
- *Transferring Ownership of a Database Object* on page 353

Granting and Revoking Permissions to a Group
Grant or revoke permissions on database objects, options, or commands.

1. In the Administration Console view, select **ASE Servers > Security > Groups**.
2. Find the row containing your column encryption key and click the **Name** field.
3. Click on the arrow, and select **Properties**.
4. Select **Object Permissions**.
5. Select an object from the list of database objects: tables, stored procedures, views, and so on.
6. Select the row corresponding to the database object on which to change group permissions.
7. Click either **Grant** or **Revoke.**

8. Continue clicking through and selecting all the database objects, options, and operations on which to grant or revoke permissions to your group.

9. (Optional) Click **Preview** to see the SQL statements for your command.

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**
   - **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a Group
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your group.
3. Click the Name field of the group you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the group.
6. Confirm the deletion.
7. Click **Finish**.

Manage Users
You can grant database object access to users, and change ownership of database objects using the Users properties wizard.

Creating a User
Create a new user in a database.

1. In the Administration Console view, select **ASE Servers > Security > Users**.
2. Select **New**.
   You see the Add User wizard.
3. Select the server and database in which to create a user.
4. Select the login to which the user will be assigned.
5. Enter the name of the user to create.
6. (Optional) Select **Create guest user** to create a guest user with limited privileges.
7. (Optional) Select a group to which the user will be assigned.
8. (Optional) Click **Summary** to verify your choices for creating the user.

See also
- *Creating a Group* on page 350
- *Transferring Ownership of a Database Object* on page 353
- *Groups Properties* on page 350
- *Users Properties* on page 354

Transferring Ownership of a Database Object

Use the Transfer Database Owner wizard to change ownership of database objects.

You can also search for referencing objects in the current, or other databases, that will be affected if the selected object is transferred to a different owner. If referencing objects exist, you can generate the SQL scripts to create these objects with the new owner. You can also compare the scripts to create the object with the old and new owners.

To transfer object ownership with referencing objects, first save the script that creates referencing objects with the new owner, then click through the wizard to transfer the database object ownership. When Sybase Control Center has completed the transfer, run the script to modify ownership of the referencing objects.

1. In the Administration Console view, select **ASE Servers > Security > Users**.
2. Find the row containing your user name, database, and server, and click the **Name** field.
3. Click on the arrow, and select **Transfer Database Object**.
   You see the Transfer Database Object wizard.
4. Select the type of objects to be transferred to a new owner.
5. Select the specific objects. Click **Preserve Permission**s to retain the old permissions for those objects.
6. Select one of these new owner options:
   - **Select the new user name** – when you choose this option, you must also specify additional information:
     - In the Database to Search screen, select the databases to be searched for objects that reference the object for which you are changing the owner.

- In the Object References screen, you can:
  - Save the script that is automatically generated for referencing objects – you must run the saved script, outside of Sybase Control Center, to create new instances of the referenced objects with the updated owners.

    > **Note:** Run the script only after you have changed the owner, that is, after the Transfer Database Object wizard has completed.

  - Compare the two scripts (one that creates the object with the old owner and the other with the new owner) – click the Name field of the row containing the object, and then click on the icon that appears. Upon comparing the two scripts, select **Accept** to retain the object in the list of referencing objects included in the script, or **Reject** to remove the corresponding object entries from the script.
- **Select the new login name** – change the loginame value (in system catalog `sysobjects`) of the selected objects only.

  A login must meet the following conditions to be available for selection:
  - If the current owner is `guest`, the login name must be valid, the login suid must not be in the `sysusers` or `sysaliases` tables, and the login cannot have `sa_role`.
  - If the current owner is `dbo`, the login name must be valid, the login suid must be either in the `sysaliases` table aliased to the `dbo`, or have `sa_role`.
  - If the current owner is anyone else other than `guest` or `dbo`, the login name must be valid and the login suid must be in `sysaliases` table aliased to the current owner.

7. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Group* on page 350
- *Creating a User* on page 352
- *Groups Properties* on page 350
- *Users Properties* on page 354

Users Properties
User permissions to access database objects and commands.

Click **Properties** on your group to initiate the Users Properties wizard to modify these database properties:

Wizard Option	Users Properties
General	• Groups – to change the group for the user, select from the list of groups.
Objects Owned	Select the database objects that your user owns in this database.

Wizard Option	Users Properties
Command Permissions	Permissions to create database objects – select the kind of permissions you want to grant the user: • Create default • Create procedure • Create rule • Create table • Create view • Create encryption key To revoke the permissions, select the permissions and click **Revoke**, then **Apply**.
Object Permissions	Permissions to access database objects – use the Grant Permission and Revoke Permission wizards to grant or revoke permissions for specific database operations such as insert, delete, update, reference, and decrypt for specific database objects such as tables, procedures, encryption keys, and so on. See *Granting and Revoking Permissions to Groups* on page 351.

Login Aliases	• Logins – to change the logins aliases to the user, use:
	    • **Add** – select a login and click **Apply** to add the login to your user alias.
	    • **Remove** – select the login from the list, and click **Remove**, then **Apply**.

See also
- *Groups Properties* on page 350
- *Creating a Group* on page 350
- *Creating a User* on page 352
- *Transferring Ownership of a Database Object* on page 353

Generating a DDL Script
Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
   - **Compiled Objects**
   - **Schema Objects**
   - **Security**

- **Space Management**

2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.

3. Click in the Name field of the specific object for which you want the DDL script.

4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a User

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your user.

3. Click the Name field of the user you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the user.

6. Confirm the deletion.

7. Click **Finish**.

Manage Roles

Manage permissions to multiple login accounts by creating roles and granting roles to logins.

Note: Only a system security officer can create, modify, or delete roles.

Roles

A system Security Officer can define and create roles as a convenient way to grant and revoke permissions to several logins. A role can be granted only to a login account or another role.

Options for Creating or Modifying Roles

The following options are available when creating or modifying roles.

- Choose permission access for object types or command type
- Expire passwords
- Set mutually exclusive roles
- Set role hierarchy
- Assign logins to roles
- Set passwords

Permissions

Permissions granted to roles override permissions granted to users or groups. For example, if John is granted the role of system security officer and individual permissions of sales

accounts, John will still be able to access permission of the sales accounts if the individual permissions are revoked because his role permissions override his user permissions.

Hierarchical Roles

A System Security Officer can define role hierarchies such that a role can be assigned to another role. For example, the chief financial officer role might contain both the financial analyst and the salary administrator roles.

Mutually Exclusive Roles

Roles can be defined to be mutually exclusive. The supported exclusive types are:

- Membership – one user cannot be granted two different roles. For example, the system administrator and system security officer roles can be defined as mutually exclusive for membership; that is, one user cannot be granted both roles.
- Activation – one user cannot activate, or enable, two different roles. For example, a user might be granted both the senior auditor and the equipment buyer roles, but is not permitted to have both roles enabled simultaneously.

Expiring a Role Password

Use the Administration Console to change a password for a role.

1. In the Perspective Resources view, select the servers on which the roles are defined, then select **Administration Console**.
   The servers you choose here are used for some of the options in the steps below.
2. Select **ASE Servers > Security > Roles > Configure Role Passwords**.
3. On the Servers Selection window, select the individual servers on which to expire role passwords, or select all servers.
4. Select **Expiration**.
5. Select **Expire role passwords** to specify that the password must be changed for the role. Choose one of:

   - **Expire passwords** – to expire passwords for specific roles or roles matching specified characters.
   - **Expire stale passwords** – to expire passwords that have not been changed by a specified date.

Creating a Role

Login accounts can be granted one or more roles. Roles can also granted to other roles.

1. In the Perspective Resources view, select the servers on which the role is to be created, then select **Administration Console**.
   The servers you choose here are used for some of the options in the steps below.
2. Select **ASE Servers > Security > Roles**.

3. Select **New**.

4. In the Introduction window, select the individual servers on which to create a role, or select all servers.

5. Click **Role Name**.

6. Specify the name of the role to create.

7. (Optional) Click **Set password** and enter a password for the role.

8. (Optional) Click **Summary** to verify your selected options.

See also

• *Role Properties* on page 358

Role Properties

Use the Administration Console option on your Adaptive Server to view or modify role properties.

The Role Properties wizard allows you to perform these actions:

Wizard Option	Role Properties
General	Password – the system security officer can set or expire a password for a role.
Logins	Login – add or remove logins assigned to a role. See *Managing Logins Assigned to Roles* on page 358.
Hierarchy	Create roles that are hierarchically mapped or aliased to another role. See *Creating Role Hierarchy* on page 359.
Exclusivity	Control privileges of roles by defining the roles as mutually exclusive. See *Creating Mutually Exclusive Roles* on page 359 .
Command Permissions	Grant or revoke command permissions for the selected role. See *Setting Command Permissions* on page 359.
Object Permissions	Grant or revoke object permissions for a selected role on a selected object type. See *Setting Object Permissions* on page 360.

Managing Logins Assigned to a Role

Use the Administration Console to assign one or more logins to a role.

1. From the Administration Console, select **ASE Servers > Security > Roles**.

2. Select the role to which to assign logins, or view currently assigned logins.

3. Select **Properties**.

4. From the Roles Properties window, select **Logins**.

- To assign logins, click **Add** and select one or more logins from the available list. Optionally, select **Active By Default** to indicate the role must be automatically activated on login.
- To remove logins, select a login account and click **Remove**.

Setting Command Permissions for a Role
Manage login account privileges by granting command permissions to a selected role.

1. In the Perspective Resources view, select the servers on which the roles have been defined, then select **Administration Console**.
2. Select **ASE Servers > Security > Roles**.
3. Select one or more roles on which to set command permissions, then select **Properties**.
4. From the Roles Properties window, select **Command Permissions**.
5. Select a database on which the selected roles will have permission to execute commands.

   - To grant command permissions for the selected roles, click **Grant** and select one or more commands.
   - To revoke command permissions for the selected roles, select a command and click **Revoke**.

Managing Mutually Exclusive Roles
Use mutually exclusive roles to prevent users from being granted different roles, or activating two different roles.

Use mutually exclusive roles to control or restrict permissions or privileges.

1. From the Administration Console, select **ASE Servers > Security > Roles**.
2. Choose the role in which to add or remove mutually exclusive roles and select **Properties**.
3. From the Roles Properties window, select **Exclusivity**.

   A list of roles that are mutually exclusive to the selected role is displayed.

   - To add mutually exclusivity roles, click **Add** , then select one or more roles.
   - (Optional) Select **Membership** to indicate that one user cannot be granted two different roles.
   - (Optional) Select **Activation** to indicate that one user cannot activate, or enable, two different roles.
   - To remove mutually exclusivity roles, select a role and click **Remove**.

Creating Role Hierarchy
Roles can be granted hierarchically to manage permissions or privileges for one or more logins.

1. In the Perspective Resources view, select the servers on which the roles have been defined, then select **Administration Console**.

2. Select **ASE Servers > Security > Roles**.

3. Select the role to which to assign additional roles, then select **Properties** .
   The selected role will be the top-level role. Additional roles can be assigned to the top-level role.

4. Click **Hierarchy**.

5. From the Role Properties window, click **Assign**.

6. From the Add a Role Assignment window, select one or more roles.
   The screen shows a folder. You can expand the folder to see the roles assigned to the top-level role. The top-level role is given permission and privileges of the lower level roles.

Setting Object Permissions for a Role
Use object access permissions to regulate the use of specific commands that access specific database objects.

1. In the Perspective Resources view, select the servers on which the roles have been defined, and select **Administration Console**.

2. Select **ASE Servers > Security > Roles**.

3. Select one or more roles on which to set object permissions, then select **Properties** .

4. Select **Object Permissions**.

5. (Optional) To view the current permission for an object, click on the drop-down menu and select the object type.

6. To grant permissions, select a database from the drop-down menus and click **Grant**.
   a) Select the object type to which to grant access and click **Objects and Options**.
   b) Select the specific object from the columns list and click **Permissions**.
   c) Select the type of access permissions to grant.

7. To revoke command permissions from the selected roles, select an object and click **Revoke**.

   - To revoke all permissions from the selected object type, click **Revoke all permission**.

   - To revoke individual permission access type, click the checked cells under each permission type.

Segments

Monitor the segments used by Adaptive Server databases.

Monitor Segments

Determining the Space Used by a Table on a Segment

Find reserved space figures for tables on a segment.

You can sort tables by their reserved sizes, which simplifies planning for a reorganization or rebuild.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Segments**.

   You can also display the Segments window by clicking a **Segments** link on another window in the Adaptive Server monitor.
3. (Optional for Adaptive Server cluster configurations) Select **Global** for information about segments on the global database. Select **Local** for information about segments on local, temporary databases.
4. In the Segments table, select the segment to monitor.
5. Click the **Used Tables** tab.
   The tab displays the name and amount of space reserved, in kilobytes, for each table.
6. To sort a table by reserved size, select the table and click **Space Reserved**.

For more information on segments, see the chapter on creating and using segments in the *Adaptive Server System Administration Guide, Volume 2.*

See also
- *Extending a Segment in Adaptive Server* on page 361
- *Displaying Information About Segments Used by a Database* on page 258
- *Segment Statistics and Details* on page 362

Extending a Segment in Adaptive Server

Sybase Control Center allows you to extend a segment on a specific device.

1. In the Adaptive Server monitor, select **Segments**.
   You see the Segments table.
2. Select the segment to configure.
3. Right-click the selected row and select **Extend.**
4. Select the **Extend** menu item from the context menu.
   You see the **Extend** dialog, which includes the name, size, and unused size of the device.
5. Click the device name onto which to extend the segment.

   Note: The device list is empty if the selected segment is using all the devices configured on the server. In this case, add a new device to the server in order to extend the segment.

6. Click **OK**.

See also
- *Determining the Space Used by a Table on a Segment* on page 361
- *Displaying Information About Segments Used by a Database* on page 258
- *Segment Statistics and Details* on page 362

Segment Statistics and Details
Interpret the Segment information for Adaptive Server.

The Segment Statistics and Details screen displays information about Adaptive Server segments. In Adaptive Server cluster configurations, selecting the **Global** tab displays information about segments on global databases. Selecting the **Local** tab displays information about segments on local, temporary databases, grouped by cluster instances.

The Segments screen displays information about all segments for this Adaptive Server. The charts on this screen are populated by data from the collection_ase_all_client_kpis, covering the current trend period.

The Segments table lists each segment used by this Adaptive Server and gives the name of the database that uses the segment, the database's size and unused space on the segment, and the number of thresholds.

The tabs at the bottom of the screen show information about the segment selected in the Segments table.

Table 42. Tabs on the Segments Screen

Details	Displays two charts:
	- A pie chart showing current space usage on the selected segment. Includes used and unused space, in megabytes, and as percentages of the available space on the segment. The title above the chart indicates the total available space. - Space Usage – a line graph showing changes in space usage on the selected segment over the current trend period.
Devices Used	Displays devices included in the selected segment and the size of each device, in megabytes. Click a name in the Device column to switch to the Devices monitoring view's information for that device.
Used Tables	Displays tables allocated on the selected segment and the reserved size of each table, in kilobytes.
Used Indexes	Displays indexes allocated on the selected segment and the table associated with each index.

For more information on segments, see the chapter on creating and using segments in the Adaptive Server *System Administration Guide*, Volume 2.

See also
*   *Determining the Space Used by a Table on a Segment* on page 361
*   *Extending a Segment in Adaptive Server* on page 361
*   *Displaying Information About Segments Used by a Database* on page 258
*   *Setting Up Statistics Collection* on page 120

Manage Segments
Sybase Control Center allows you to create, delete, and generate data definition language for segments.

Displaying Segments
Sybase Control Center allows you to display a summary of available segments—labels that point to one or more database devices—in your databases.

In the Administration Console, select and expand **ASE Servers > Space Management > Segments**.
You see a list of existing segments and their properties:
*   **Name** – the name of the database device.
*   **Server** – the name of the server in which the database device resides.
*   **Database** – indicates in which database the segment resides. This column includes both system-provided databases (such as model) and user-created databases.
*   **Last Chance** – indicates whether a last-chance stored procedure such as **sp_thresholdaction** is added to the segment. See *Managing Free Space with Thresholds* in the *System Administration Guide*.
*   **Size** – displays the size of the database, in megabytes.
*   **Used** – displays the amount of memory used by the database, in megabytes.
*   **Free** – displays the amount of unused memory in the database.

Creating a Segment
Use a wizard to create a new segment—a label that points to one or more database devices—in a database.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Segments**.
2. Click the arrow and select **New**.
3. Complete these wizard pages:
    *   **Introduction** – select the server and database in which to create the segment.
    *   **Segment Name** – enter the name of the segment to create.

- **Device Selection** – select the database device to use for the segment.

4. (Optional) Click **Preview** to view the SQL statement that Adaptive Server will use for this wizard.

5. Click **Finish**.

Segment Properties

Use the Segment Properties wizard to modify database devices, tables, and thresholds.

Click **Properties** on your segment—the label that points to one or more database devices—to initiate the Properties wizard.

Wizard Option	Segment Properties
General	Displays the segment's summary information that appears in the segments list view, as well as the segment's hysteresis value. See *Managing Free Space with Thresholds* in the *System Administration Guide* for information about the hysteresis value. Specify how to show the current size: • Pages • Kilobytes • (Default) Megabytes • Gigabytes
Devices	Displays the database devices used by the segment, and their sizes, in megabytes. You can also: • Add a new database device to the segment. See *Adding a Database Device to a Segment* on page 365. • Remove an existing database device that the segment uses. See *Removing a Database Device from a Segment* on page 365. • View the properties of the database device – when you click Properties, the Database Device Properties wizard appears.
Contains	Displays: • Tables that use the segment – the list includes both the table name and its owner. • Indexes that use the segment – the list includes both the index name and the table the index uses.
Thresholds	Displays thresholds that are added to the segment in the form of system procedures and their owners. You can also add and remove thresholds. See *Adding a Threshold to a Segment* on page 366.

See also
- *Adding a Database Device to a Segment* on page 365
- *Removing a Database Device from a Segment* on page 365
- *Adding a Threshold to a Segment* on page 366
- *Removing a Threshold from a Segment* on page 366

Adding a Database Device to a Segment
Use the Segment Properties wizard to add an existing database device to a segment.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Segments**.
2. In the **Name** field, select the segment to modify.
3. Click the arrow on the segment and select **Properties**.
4. In the Devices dialog, click **Add**.
5. Choose an existing database device from the list to the add to the segment.

See also
- *Removing a Database Device from a Segment* on page 365
- *Adding a Threshold to a Segment* on page 366
- *Removing a Threshold from a Segment* on page 366
- *Segment Properties* on page 364

Removing a Database Device from a Segment
Using the Segment Properties wizard in the Sybase Control Center, you can remove a database device that is associated with a segment.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Segments**.
2. In the **Name** field, select the segment to modify.
3. Click the arrow on the segment and select **Properties**.
4. In the Devices wizard page, click **Remove**.
5. Select the database device to remove.

See also
- *Adding a Database Device to a Segment* on page 365
- *Adding a Threshold to a Segment* on page 366
- *Removing a Threshold from a Segment* on page 366
- *Segment Properties* on page 364

Adding a Threshold to a Segment
Sybase Control Center allows you to use a wizard to add database devices for segments—labels that point to one or more database devices—to use.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Segments**.
2. Choose the segment, click the arrow, and click **Properties.**
3. Click **Thresholds**. Any existing threshold-related stored procedures for this segment appear in the table, listed by procedure name, owner, and the amount of free space, in megabytes.
4. Click **Add** to view the Add New Threshold wizard page.
5. Choose a stored procedure, and specify its free-space threshold in pages, kilobytes, megabytes, or gigabytes. The default is in megabytes.
6. Click **Apply** after each new stored procedure you add, and **OK** when you are finished.

See also

Removing a Threshold from a Segment
Sybase Control Center allows you to use a wizard to remove database devices associated with segments.

1. In the left pane of the Administration Console, expand **ASE Servers > Space Management > Segments**.
2. Choose the segment, click on the arrow, and click **Properties**.
3. Click **Thresholds**. Any existing threshold-related stored procedures for this segment appear in the table, listed by procedure name, owner, and the amount of free space, in megabytes.
4. Select the stored procedure to delete, and click **Remove**.
5. Click **Apply** after each stored procedure you delete, and **OK** when you are finished.

See also

Generating a DDL Script

Use Sybase Control Center to generate DDL to create a database or any of its objects.

Sybase Control Center includes an option that lets you generate DDL scripts for databases, tables, caches, devices, dump devices, segments, groups, roles, users, encryption keys, and compiled objects such as stored procedures, extended stored procedures, and functions.

1. In the left pane of the Administration Console, expand **ASE Servers**, then select one of:
    - **Compiled Objects**
    - **Schema Objects**
    - **Security**
    - **Space Management**
2. Select the database object for which to create DDL. For example, to obtain DDL for a specific stored procedure, select **Procedures > Stored Procedures**
   Sybase Control Center displays the list of all objects of the selected type defined in your Adaptive Server.
3. Click in the Name field of the specific object for which you want the DDL script.
4. Click the drop-down arrow, then select the option to generate DDL.

   You can save the DDL in an external file on your local file system.

Deleting a Segment

Sybase Control Center helps you delete database objects, or the database itself.

Note: You can delete only the segments you created; the **Delete** option is not available for segments created by other users.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your segment.
3. Click the Name field of the segment you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the segment.
6. Confirm the deletion.
7. Click **Finish**.

Server Configuration

Display and monitor Adaptive Server configuration values.

Monitor Server Configuration

Monitor Adaptive Server configuration values.

Displaying Configuration Values

Display the values and descriptions of parameters in the Adaptive Server configuration file.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Server Configuration**.

3. (Optional - for Adaptive Server cluster configurations only) Select a cluster instance from **Show Configuration for Cluster Instance**. Select **All** for information on all cluster instances.

4. Select a configuration category from **Show Configuration Parameters for**. For example, select **Cache Manager**.

5. (Optional) Enter filter text in **Show Configuration Parameter Matching**. For example, enter sql.
   The Server Configuration table displays parameters whose names match or include your filter text. The filter narrows the selection made with **Show Configuration Parameters for**, so if you select **All** and then filter on sql, you find more parameters than when you select **Monitoring** and then filter on sql.

6. From the Server Configuration table, select a parameter.
   A description of the parameter appears at the bottom of the screen.

For more information on configuration parameters, see the chapter on setting configuration parameters in the Adaptive Server *System Administration Guide*, Volume 1.

See also

Modifying Server Configuration Parameters

Use Sybase Control Center to configure server configuration parameters.

1. In the Adaptive Server monitor, select **Server Configuration**. Alternately, in the Administration Console, select the server, click the drop-down arrow, and select **Configure**.

   Note: You must have **sa_role** to select **Configure** from the Administration Console.

   You see a table with fields parameter names, value, default value, maximum value, minimum value, and restart required. Editable columns are indicated by a "pencil" icon. To reset a value to the previously configured value, click the "Reset" icon that appears after you have edited a field.

2. Select the server configuration parameter to configure. For example, increase or decrease the size of the procedure cache by selecting the **procedure cache size** parameter, or of the statement cache by selecting **statement cache size** parameter.

   a) Click **Configure Value** for the selected row.
   b) Enter the new value for the configuration parameter.
      If the value is invalid, you see an error message. As some parameters require a restart, the changed value is in the **Pending Value** column till the server is restarted.
   c) Click **Save All** to update the server with the new values or **Reset All** to restore the original values for the resource.

See also
* *Displaying Configuration Values* on page 368
* *Server Configuration Statistics and Details* on page 369

Server Configuration Statistics and Details
Interpret Adaptive Server server configuration information.

The Server Configuration screen displays information about all configuration parameters for this Adaptive Server. You can modify the configuration parameters using either the **sp_configure** stored procedure, or by editing the configuration file. The parameters are stored in the Adaptive Server configuration file on each server.

You can also use the Sybase Control Center to change the current values of the server configuration parameters and change the size of the procedure cache and statement cache by editing the **Value** fields corresponding to the specific row.

For each parameter, the Server Configuration table gives the name; displays the current, default, minimum, and maximum values; and indicates whether you have to restart the server to make a change to this parameter take effect.

See also
* *Modifying Server Configuration Parameters* on page 368
* *Displaying Configuration Values* on page 368

Statistics

Availability and performance statistics in Sybase Control Center for Adaptive Server can help you determine whether your systems are running as efficiently as possible.

Availability statistics are concerned with present conditions; they help you determine whether a system component you are monitoring (a server or a storage device, for example) is running and functioning properly. Performance statistics are concerned with behavior of the same components over time. You can use them to spot trends, identify problems or potential problems, and make plans.

Sybase Control Center includes predefined key performance indicators (KPIs) for Adaptive Server that are grouped into collections. KPIs such as Server Status, which serves as an

availability statistic when it is fresh, also have long-term value as historical performance statistics.

Availability statistics appear on the heat chart and on the screens of the Adaptive Server monitor for each Adaptive Server. The heat chart includes these KPIs:

- Server Status (up or down)
- Server CPU Utilization
- Number of Blocked Processes
- Number of Suspended Processes

These KPIs are part of the default collection: collection_ase_availability.

Performance statistics appear on the statistics chart and on the screens of the Adaptive Server monitor for each Adaptive Server. In the statistics chart, you expand folders in the Statistics tab to drill down to specific KPIs. See the KPI values as tables or graphs, and compare them by displaying several KPIs together. These are the folders that contain KPIs you can use to display data:

- Server Overview
- Devices
- Engines
- Segments
- DataCaches

Adaptive Server cluster configurations have these additional folders on the statistics chart:

- Cluster Instances
- Cluster Workload
- Logical Clusters
- tempDBs Activity

To make specific KPIs available to the statistics chart and to the Adaptive Server monitor screens that use them, in addition to the availability statistics scheduled by default, set up collection jobs in the scheduler for:

- collection_ase_histmon – The KPIs in this collection are only available in the statistics chart and for setting alerts.
- collection_ase_all_client_kpis – This collection is necessary to gather statistics for historical charts in the Adaptive Server Monitor. The statistics are also available in the statistic chart and for setting alerts.

Several configuration options affect the collection and appearance of Adaptive Server data in Sybase Control Center:

- Collection repeat interval – the frequency of data collection. Set the repeat interval on the collection job in the scheduler. This option is set when the collection is scheduled, but may be modified later.
- Screen refresh interval – the period between screen refreshes. Refreshing the screen redraws it with the latest available data. This option can be configured from the Settings window.
- Chart trend period – the period over which data appears in historical charts. This option can be configured from the Settings window.

See also
- *Executing SQL Statements* on page 193

Interpreting Statistics
Understand the scope and freshness of Adaptive Server data in Sybase Control Center.

Each Adaptive Server statistic presented in Sybase Control Center has a scope:

- Delta – the number of occurrences since the last screen refresh. For example, the user log cache statistics on the Transactions screen are delta values.
- Rate – the number of occurrences over the given period of time. Device I/O is given as a rate.
- Percentage or ratio – an amount, number, or rate stated as a proportion to a whole. Percentage statistics include CPU utilization, space usage on devices and segments, and cache hit rates. Ratios include cache volatility.
- Count – a simple value; for example, the size of a database or cache in megabytes, or the number of partitions in a cache.
- Cumulative – the number of occurrences since Adaptive Server started, or since the counter wrapped. On the Processes screen, you can set the Top 5 User Processes chart to display the five processes that use the most CPU, disk I/O, or network resources as rates or as cumulative values.

Note: When a server has been running for a long time, its statistical counters can wrap, which means they restart from zero. This most often affects cumulative statistics. Information about when or how many times a counter has wrapped is not available.

Most Adaptive Server statistics not otherwise labeled are presented as deltas since the last screen refresh; rates, percentages, and cumulative numbers are labeled as such.

Sybase Control Center displays statistics promptly. However, there are several factors that affect the freshness of the data on the screens:

- The screen refresh interval, which you can set on the Settings screen
- The collection repeat interval, which you can set in the scheduler for each Adaptive Server collection
- Network latency

For more information about interpreting Adaptive Server statistics, see:

- *Adaptive Server Performance and Tuning Series: Improving Performance with Statistical Analysis*
- *Adaptive Server Performance and Tuning Series: Monitoring Adaptive Server with sp_sysmon*
- *Adaptive Server Performance and Tuning Series: Basics*

SQL Activity

Monitor SQL queries on Adaptive Server.

Monitoring SQL Queries

Display details about recently executed SQL queries.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **SQL Activity**.
   The SQL Activity screen appears. It lists queries executed during the current trend period, along with details including each query's server process identifier (spid), the login account that executed the query, the kernel process identifier (KPID), batch identifier, and execution statistics.

   Note: In Adaptive Server cluster configurations, information in the Active SQLs table is grouped by cluster instances.

3. Select a SQL query from the Active SQLs in Batch table.
   The SQL statement appears at the bottom of the screen.

   Note: The SQL Activity screen displays SQL text for the most recent collection interval only.

See also
- *Executing SQL Statements* on page 372

Executing SQL Statements

Use Sybase Control Center to execute SQL statements on one or more Adaptive Servers.

You can use the Execute SQL view to run 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, SCC displays an error.

1. In the Administration Console, select one or more servers, click the drop-down arrow, and select **Execute SQL**.
2. Enter the SQL statements.
3. Click **Execute**.

The query runs on all the servers you selected and results appear in the bottom portion of the view. The view includes a results tab for each server. On the tabs,

- A green check indicates a successful query.
- A red X indicates an error. A tab with a red X also displays an error message.

See also

Tables

Manage tables used by Adaptive Server databases.

Manage Tables

Use the Administration Console to create or modify tables and table objects.

Tables consist of columns and rows that contain data on a database. Adaptive Server uses the following types of tables:

- A system table stores information that allows the database to perform its services.
- A user table stores and provides access to user data.
- A proxy table accesses data on remote servers.

Note: System table definitions are not usually updated.

To plan a table's design:

- Decide what columns you need in the table, and the datatype, length, precision, and scale, for each.
- Create any new user-defined datatypes before you define the table where they are to be used.
- Decide which column, if any, should be the IDENTITY column.
- Decide which columns should and which should not accept null values.
- Decide what integrity constraints or column defaults, if any, you need to add to the columns in the table.
- Decide whether you need defaults and rules, and if so, where and what kind.
- Consider the relationship between the NULL and NOT NULL status of a column and defaults and rules.
- Decide what kind of indexes you need and where.

Creating a User Table

Create a user table to store and provide access to user data.

Note: Only a database owner or a user with create table permission can create a table.

1. In the Perspective Resources view, select the servers on which the table is to be created, then click the drop-down arrow next to a server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables** and select **User Tables**.

3. Select **New**.
   You see the Add User Table wizard.

4. In the Introduction window, select the server, database, and owner for the new table.

5. Enter a name for the user table.

6. Enter the SQL statements for the new table and related table objects.

7. (Optional) Click **Summary** to verify your selected options.

8. Click **Finish** to create the user table.

See also
* *Creating a Proxy Table* on page 374

Creating a Proxy Table
Create a proxy table to access data on remote servers.

A proxy table is a user table that allows you to access data in a remote table, view, remote procedure call, directory, or file. A proxy table has all the attributes of a user table, such as columns, indexes, and triggers, but it does not contain any data locally.

Note: Only a database owner or a user with create table permission can create a table.

1. In the Perspective Resources view, select the servers on which the table is to be created, then click the drop-down arrow next to a server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables** and select **Proxy Tables**

3. Select **New**.
   You see the Add Proxy Table wizard.

4. In the Introduction window, select the server, database, and owner for the new table.

5. Enter a name for the proxy table.

6. Enter the SQL statements for the new table and related table objects.

7. (Optional) Click **Summary** to verify your selected options.

8. Click **Finish** to create the proxy table.

See also
* *Creating a User Table* on page 373

Creating a Column

Add a column to an existing table.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:
   - **User Tables**
   - **Proxy Tables**

3. Select the table for the new column.

4. Select **Properties**.

5. From the Table Properties window, select **Columns > New**.
   You see the Add Column wizard.

6. In the Name window, enter the name of the column.

7. In the SQL Editor window, enter the SQL statements for the new column and related objects.

8. (Optional) Click **Summary** to verify your selected options.

Creating an Index

An index provides quick access to data in a table, based on the values in specified columns.

An index is created on one or more table columns and points to the place where the column data is stored on disk. Indexes speed up data retrieval and are useful for enforcing referential integrity. A table can have more than one index.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:
   - **User Tables**
   - **Proxy Tables**

3. Select the table for the new index.

4. Select **Properties**.

5. From the Table Properties window, select **Indexes > New > Index**.
   You see the Add Index wizard.

6. Enter a name for the index.

7. Select the columns to include in the index.

8. (Optional) Click on **Add index column expression**.
   a) Select ascending or descending as the order of the index expression.

    b) Enter a name for the expression.

9. Select the database segment in which to place the index.

10. Select a data cache for the index.

11. (Optional) In the Key Type window, select:

- **Make this index unique**.
  If the index is unique, you can ignore duplicate keys in the Duplicate Keys/Row window.
- **Make this index clustered**.
  If the index is clustered, specify how you want the server to handle requests to insert duplicate rows in a table in the Duplicate Keys/Row window.

12. (Optional) In the Duplicate Key window:

- Click **Ignore duplicate keys** to ignore duplicate keys rather than abort the transaction.

- Choose whether to allow or ignore duplicate rows in a table.

13. (Optional) In the Space Management window:

    a) Specify the percentage amount to fill a page when the index is created.
    b) Specify the number of rows allowed on pages.
    c) Specifying a ratio of empty pages to filled pages.

14. (Optional) Specify the cache strategy when creating the index.

- **Most recently used replacement** – reads new pages into the LRU end of the chain of buffers in cache. The pages are used and immediately flushed when a new page enters the MRU end. This strategy is advantageous when a page is needed only once for a query. It tends to keep such pages from flushing out other pages that can potentially be reused while still in cache.
- **Large buffer prefetch**– if memory pools for large I/O are configured for the cache used by a table or an index, the optimizer can prefetch data or index pages by performing large I/Os of up to eight data pages at a time. This prefetch strategy can be used on the data pages of a table or on the leaf-level pages of a nonclustered index. By default, prefetching is enabled for all tables, indexes, and text or image objects. Setting the prefetch option to off disables prefetch for the specified object.
- **Data already sorted**– if data is already sorted, this option saves index creation time.

15. (Optional) Specify whether to create a local partitioned index.

16. Click **Finish** to create the index.

Creating a Trigger

Create a trigger on a table to enable checks whenever data is inserted, updated, or deleted.

A trigger is a special type of procedure attached to a table column that goes into effect when a user changes the table. Triggers execute immediately after data modification statements are completed.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table for the new trigger.

4. Select **Properties**.

5. From the Table Properties window, select **Triggers > New**.
   You see the Add Trigger wizard.

6. Enter the name of the trigger, then select the owner of the trigger.

7. In the Trigger type window:

   - Select the events, that when executed, will call the trigger.
   - Select **Update of columns**, then select the columns to be updated. If changes are made to any of the selected columns, the trigger executes.

8. Enter the SQL statements for the new trigger and related table objects.

9. (Optional) Click **Summary** to verify your selected options.

Restoring Table Data

Restore table data from an archive or stand-by database.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables** and select **User Tables**.

3. Select the table for which to restore data.

4. Select **Restore Data**.
   You see the Restore Table Data wizard.

5. Select the database where the source table data is located.

6. Select the table to be used as a source for restoring data.

7. (Optional) Click **Preview** to verify your selection options.

8. Click **Copy Data** to start the restore process.

Creating a Foreign Key

Create a foreign key to constrain a column based on values in a reference table.

A foreign key is a column or combination of columns that have values that match the primary key. A foreign key does not need to be unique. It is often in a many to-one relationship to a primary key. A foreign key may be null; if any part of a composite foreign key is null, the entire foreign key must be null.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table for the new foreign key.

4. Select **Properties**.

5. From the Table Properties window, select **Foreign Keys > New**.
   You see the Add Foreign Key wizard.

6. In the Referenced Table window, select the database that contains the table that the foreign key references.

7. In the Referenced Column window, select the column that the foreign key references.

8. Select the referenced table, then specify a name for the foreign key.

9. Select columns from the **Foreign Key Column** to match to columns in **Primary Key Column**.
   Foreign key values should be copies of the primary key values. No value in the foreign key should exist unless the same value exists in the primary key.

10. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Check Constraint* on page 378
- *Creating a Unique Constraint* on page 383
- *Creating a Primary Key* on page 382
- *Binding Defaults and Rules to a Column* on page 380

Creating a Check Constraint

Creating a check constraint specifies a condition that any value must pass before it is inserted into the table.

A check constraint specifies a condition that any value must pass before it is inserted into the table. You can can create a check constraint at the table or column level. Column-level check constraints reference a single column. Table-level check constraints apply to the entire table.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select a table for the new check constraint.

4. Select **Properties**.

5. From the Table Properties window, select **Check Constraint > New**.
   You see the Add Check Constraint wizard.

6. Enter a name for the check constraint.

7. Enter an expression that defines the constraint.
   For example, `salary > 0`. The comparable command line syntax using **alter table** is:

```
alter table sample.dbo.employee
add constraint test_const
CHECK  (salary > 0)
```

8. (Optional) Click **Summary** to verify the check constraint expression.

See also
- *Creating a Unique Constraint* on page 383
- *Creating a Primary Key* on page 382
- *Creating a Foreign Key* on page 377
- *Binding Defaults and Rules to a Column* on page 380

Checking Table Consistency
Check and repair the logical and physical consistency of a table.

1. In the Perspective Resources view, select the server on which the table resides, then click
   the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects >
   Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**
   - **System Tables**

3. Select the table for which to check consistency.

4. Select **Check Consistency**.
   You see the Check Consistency wizard.

5. Select **Choose DBCC options**.

6. (Optional) Select **Check overall consistency**, then optionally click **Ignore non-clustered
   indexes**.
   Check overall consistency checks that:
   - Index and data pages are linked correctly.
   - Indexes are sorted properly.
   - Pointers are consistent.
   - All indexes and data partitions are correctly linked.

- Data rows on each page have entries in the row-offset table.
- Partition statistics for partitioned tables are correct.

7. (Optional) Select **Check allocation**, then optionally click **Fix allocation errors**. Check allocation checks the table to ensure that:
   - All pages are correctly allocated.
   - Partition statistics on the allocation pages are correct.
   - No page is allocated that is not used.
   - All pages are correctly allocated to the partitions in the specified table and that allocated pages are used.
   - No unallocated page is used.

8. (Optional) Select **Reindex**.
   This option allows the system administrator or table owner to check the integrity of indexes attached to a user table and to rebuild suspect indexes.

9. (Optional) Select **Fix text**.
   Fix text is used only for tables that contain text data. Select if you are changing to a new multibyte character set from either a single-byte or a multibyte character set.

10. Select the type of allocation report.

11. Click **Finish** to execute the selected commands.

Binding Defaults and Rules to a Column

Specify constraints on column data by binding defaults or rules to a column.

Make sure that any default value bound to a column or user-defined datatype is compatible with the rule. A default that conflicts with the rule is not inserted.

You cannot bind a rule to a text, image, or timestamp column.

Rules bound to columns take precedence over rules bound to user-defined datatypes.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select **Properties**.

4. Click **Columns**.

5. Select a column, then select **Properties**.

6. Click **Rules and Defaults**.

7. Choose one of:

   - **Default – None.**

- **Default – Binding** to bind an existing default to the column.
- **Default – Value** to bind a default user, defined value to the column.
- **Rule Binding** to bind an existing rule to a column.

8. Click **Apply** to apply your rules or defaults.

See also
- *Creating a Check Constraint* on page 378
- *Creating a Unique Constraint* on page 383
- *Creating a Primary Key* on page 382
- *Creating a Foreign Key* on page 377

Placing a Table on a Segment
Using a segment to put a table on a specific database device can improve Adaptive Server performance and give increased control over placement, size, and space usage of database objects.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

- **User Tables**
- **Proxy Tables**

3. Select the table to be placed on a segment.

4. Select **Properties**.

5. Click **Usage**.

6. Select the segment to on which to place the table.

7. (Optionally) Select the type of unit measurement to view used and available space of the segment.

Setting the Table Locking Scheme
Choose or alter a locking scheme based on required performance.

Conversions between all pages locking and data-only locking schemes can be expensive in time and I/O and require sufficient free space. Convert the locking scheme by creating copies of the tables and re-creating indexes. You must also dump the affected databases, and update statistics before changing between all pages locking and data-only locking schemes.

Conversions between data page and data row locking are quick and inexpensive, and implemented by updates to system tables. The data page and data row schemes are collectively called data-only locking.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table on which set the locking scheme.

4. Select **Properties**.

5. Click **Lock Scheme**.

6. Select the Lock Scheme.

7. Set the space management parameters:

   - **Max rows per page** – limits the number of rows on a data page.
   - **Expected row size** – sets the rows size, but can increase the amount of storage required. If your tables have many rows that are shorter than the expected row size, setting this value and reorganizing the use of table space or changing the locking scheme increases the storage space required for the table.
   - **Reserve page gap** – leaves empty pages on extents that are allocated to the object when commands that perform extent allocation are executed. Setting **Reserve page gap** to a low value increases the number of empty pages and spreads the data across more extents, so the additional space required is greatest immediately after creating an index or reorganizing the use of table space.
   - **Fill factor** – allows space on the index pages to reduce page splits. Very small fill factor values can cause the storage space required for a table or an index to be significantly greater.

8. (Optional) After converting from all pages locking and either of the data-only locking schemes, check table and database consistency. You must also perform a full database dump before you can back up the transaction log with **dump transaction**.

Creating a Primary Key

Creating a primary key constraint ensures that no two rows in a table have the same values in the specified columns.

A primary key is a column or combination of columns that uniquely identifies a row. It cannot be NULL and it must have a unique index. A table with a primary key is eligible for joins with foreign keys in other tables. Think of the primary key table as the master table in a master-detail relationship. There can be many such master-detail groups in a database.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

- **User Tables**
- **Proxy Tables**

3. Select the table for the primary key.

4. Select **Properties**.

5. From the Table Properties window, select **Index > Unique Constraint**.
   You see the Add Unique Constraint wizard.

6. Specify a name, then select **Primary key**.

7. (Optionally) Click **Make supporting index clustered**.

8. Select the columns to include in the primary key.

9. Select a segment on which to place the primary key.

10. (Optional) Specify a fill factor percentage.

11. (Optional) Specify the maximum number of rows per page for the index.

12. (Optional) Specify the ratio of empty pages to filled pages to provide for expansion.

13. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Check Constraint* on page 378
- *Creating a Unique Constraint* on page 383
- *Creating a Foreign Key* on page 377
- *Binding Defaults and Rules to a Column* on page 380

Creating a Unique Constraint
Use a unique constraint to ensure that no two rows have the same values in the columns of the constraint.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table for the unique constraint.

4. Select **Properties**.

5. Select **Index > Unique Constraint**.
   You see the Add Unique Constraint wizard.

6. Specify a name, then select **Unique constraint**.

7. (Optional) Click **Make supporting index clustered**.

8. Select the columns to include in the unique constraint.

9. Select a segment on which to place the unique constraint.

10. (Optional) Specify a fill factor percentage.

11. (Optional) Specify the maximum number of rows per page for the index.

12. (Optional) Specify the ratio of empty pages to filled pages to provide for expansion.

13. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating a Check Constraint* on page 378
- *Creating a Primary Key* on page 382
- *Creating a Foreign Key* on page 377
- *Binding Defaults and Rules to a Column* on page 380

Incrementally Transferring Data
Incremental data transfer lets you transfer data to Adaptive Server or other products.

Incremental Transfer is available only on Adaptive Server 15.5 or higher. In versions earlier than 15.5, the entire table is transferred. You must mark tables as eligible to participate in incremental transfer. You can designate eligibility either when you create a table, or later by turning on **Enable incremental transfer** in the **Properties > General** window.

Incrementally Transferring Data In
Use incremental transfer to read data files into Adaptive Server.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select **Incremental Transfer In**.
   You see the Incremental Transfer In wizard.

4. In the Introduction window, specify the file name of data to be read into Adaptive Server. You can optionally specify an absolute path.
   Only Adaptive Server Enterprise data file format is supported.

5. (Optional) Click **Summary** to verify the file name and path.

Incrementally Transferring Data Out
Transfer table data that has changed since a prior transmission from tables that are marked for incremental transfer.

Enable incremental transfer in the properties window of the selected table.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select **Incremental Transfer Out**.
   You see the Incremental Transfer Out wizard.

4. In the Introduction window, specify a destination file name. Optionally include an absolute path.

5. Specify the data format for the destination file.

6. Specify the order in which the column data is to written.

7. (Optional) Specify whether to encrypt the columns.

8. (Optional) Specify the tracking ID.

9. (Optional) Specify whether to resend previously transferred data, then choose either to resend data using a sequence ID to determine the starting timestamp or resend the entire table.

10. (Optional) Click **Summary** to verify your selected options.

Bulk Copying Data
You can use bulk copy to copy data into or out of a table.

Bulk copying data in or out of a table provides a convenient, high-speed method for transferring data between a database table or view and an operating system file. When copying in from a file, bulk copy inserts data into an existing database table; when copying out to a file, bulk copy overwrites any previous contents of the file.

Bulk Copying Data Into a Table
Use bulk copy to insert data into an existing database table.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select one or more tables in which to copy data.
   If you select one table, you can select a different data file for each partition. If you select multiple tables, you can only provide one data file for all partitions for the table.

4. Select **Bulk Copy In**.

You see the Bulk Copy In wizard.

5. In the Specify Data File window, enter the location.

6. Select the format for copying data.

7. In the Specify Copy Format window, choose the field and row parameters for the file to be copied into the table.

See also

• *Bulk Copying Data Out of a Table* on page 386

Bulk Copying Data Out of a Table
Use bulk copy to copy table data to an external file.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   • **User Tables**
   • **Proxy Tables**

3. Select one or more tables from which to copy data.
   If you select one table, you can select a different data file for each partition. If you select multiple tables, you can only provide one data file for all partitions for the table.

4. Select **Bulk Copy Out**.
   You see the Bulk Copy Out wizard.

5. In the Specify Data File window, enter the location.

6. Select the format for copying data.

7. In the Specify Copy Format window, choose the field and row parameters for the file to be copied from the table.

See also

• *Bulk Copying Data Into a Table* on page 385

Setting Permissions
Grant or revoke permissions on tables for users, groups, and roles.

You can grant and revoke permissions on a table based on the grantee type; users, groups or roles, then select a specific grantee. You can grant or revoke permission for specific columns belonging to a table.

Revoking Table Permissions
Revoke table access permissions from users, groups, or roles.

Table owners and database owners can revoke database object permissions from a table.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table that contains the permissions to revoke.

4. Select **Properties.**

5. Select **Permissions**.

6. Select the grantee, then click **Revoke** to revoke access permissions to the object.
   In the Revoke Permissions wizard, each type of permission and the current granted permissions are shown in cells.

7. Choose one of:

   - **Revoke all permission.**

   - Individual cells to revoke the currently granted permissions. The cell changes to show an "x", indicating that the permission type is no longer granted.

8. (Optional) Click **Preview** to see the SQL statements for your command.

See also

- *Granting Table Permissions* on page 387

Granting Table Permissions

Grant table access permission to users, groups, or roles.

Table owners and database owners can grant database object permissions on a table.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table on which to set permissions.

4. Select **Properties**.

5. Select **Permissions**.

6. Click **Grant** to grant access permissions for the selected object.
   You see the Grant Permission wizard.

7. Select the type of grantee:

- • **Users**
- • **Groups**
- • **Roles**

8. Select one or more grantees.

9. Select the columns to set permissions.

10. Select the types of permissions allowed for the selected grantees.

11. (Optional) Click **Summary** to verify your selected options.

See also
- • *Revoking Table Permissions* on page 386

Managing Partitions

Use partitioning to divide large tables and indexes into smaller, more manageable pieces.

Partitions

Partitions are database objects that have unique IDs and can be managed independently. Each partition can reside on a separate segment.

Adaptive Server supports horizontal partitioning, which means you can distribute a selection of table rows among storage devices. Assign individual table or index rows to a partition according to a partitioning strategy. By default, Adaptive Server creates every table and index on a single, round-robin partition. You can also choose a semantics-based strategy that assigns rows to using hash, list, or range partition strategies.

Semantics-based partitioning is a separately licensed feature.

Hash Partitioning

With hash partitioning, Adaptive Server uses a hash function to specify the partition assignment for each row. You select the partitioning key columns, but Adaptive Server chooses the hash function that controls the partition assignment. Hash partitioning is a good choice for:

- • Large tables with many partitions, particularly in decision-support environments
- • Efficient equality searches on hash key columns
- • Data that has no particular order, for example, alphanumeric product code keys

If you choose an appropriate partition key, hash partitioning distributes data evenly across all partitions. However, if you choose an inappropriate key, for example, a key that has the same value for many rows—the result may be skewed data, with an unbalanced distribution of rows among the partitions.

Range Partitioning

Rows in a range-partitioned table or index are distributed among partitions according to values in the partitioning key columns. The partitioning column values of each row are compared with a set of upper and lower bounds that determine the partition to which the row belongs.

Every partition has an inclusive upper bound and every partition except the first has a noninclusive lower bound.

Range partitioning is particularly useful for high-performance applications in both OLTP and decision-support environments. Select ranges carefully so that rows are assigned equally to all partitions—knowledge of the data distribution of the partition key columns is crucial to balancing the load evenly among the partitions. Range partitions are ordered; that is, each succeeding partition must have a higher bound than the previous partition.

List Partitioning

As with range partitioning, list partitioning distributes rows semantically; that is, according to the actual value in the partitioning key column. A list partition has only one key column. The value in the partitioning key column is compared with sets of user-supplied values to determine the partition to which each row belongs. The partition key must match exactly one of the values specified for a partition.

The value list for each partition must contain at least one value, and value lists must be unique across all partitions. You can specify as many as 250 values in each list partition. List partitions are not ordered.

Round-Robin Partitioning

In round-robin partitioning, Adaptive Server does not use partitioning criteria. Round-robin-partitioned tables have no partition key. Adaptive Server assigns rows in a round-robin manner to each partition so that each partition contains a more or less equal number of rows and load balancing is achieved. Because there is no partition key, rows are distributed randomly across all partitions.

Using a Hash Partition

Create a new partition or change an existing partition using a strategy of a system-generated hashing function.

Note: Semantic partitioning must be enabled in order to create hash, list, or range partitions.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:
   - **User Tables**
   - **Proxy Tables**
3. Select the table for the partition.
4. Select **Properties**.
5. Select **Partitions**.
6. Choose to create a new partition or change an existing partition.

7. In the Select Partition Strategy window, choose the partitioning strategy **Hash**.

8. (Optionally) Specify the number of partitions.

9. In the Select Partition Key Columns window, use the arrow buttons to select partition key columns.
Partition key columns are table columns that determine how the table is to be partitioned.

10. In the Partition Specification window, specify the name of the partition where the partition will reside.

11. (Optional) Click **Summary** to verify your selected options.

See also
- *Updating Partition Statistics* on page 396
- *Deleting Partition Statistics* on page 397
- *Using a Range Partition* on page 390
- *Using a List Partition* on page 391
- *Using a Round-Robin Partition* on page 392
- *Enabling Semantic-based Partitioning* on page 392

Using a Range Partition
Create a new partition or change an existing partition according to whether one or more values in a row fall within a range of predefined values for the partition.

Note: Semantic partitioning must be enabled in order to create hash, list, or range partitions.

---

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

- **User Tables**
- **Proxy Tables**

3. Select the table for the partition.

4. Select **Properties**.

5. Select **Partitions**.

6. Choose to create a new partition or change an existing partition.

7. In the Select Partition Strategy window, choose **Range**.

8. In the Select Partition Key Columns window, use the arrow buttons to select partition key columns.
Partition key columns are table columns that determine how the table is to be partitioned.

9. In the Partition Specification window, specify the name of the partition, the range of values, and where the partition will reside.

10. (Optional) Click **Summary** to verify your selected options.

See also
- *Updating Partition Statistics* on page 396
- *Deleting Partition Statistics* on page 397
- *Using a Hash Partition* on page 389
- *Using a List Partition* on page 391
- *Using a Round-Robin Partition* on page 392
- *Enabling Semantic-based Partitioning* on page 392

Using a List Partition
Create a new partition or change an existing partition according to whether one value in the row matches one of a set of predefined values unique for each partition.

Note: Semantic partitioning must be enabled in order to create hash, list, or range partitions.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables** and select **User Tables**.

3. Select the table for the partition.

4. Select **Properties**.

5. Select **Partitions**.

6. Choose to create a new partition or change an existing partition.

7. In the Select Partition Strategy window, choose **List**.

8. In the Select Partition Key Columns window, use the arrow buttons to select one partition key column.
List partitions use only one key column. The value in the partitioning key column is compared with values supplied in the partition specification window to determine the partition to which each row belongs.

9. In the Partition Specification window, specify the name of the partition, specify a list of discrete values, and where the partition will reside.

10. (Optional) Click **Summary** to verify your selected options.

See also
- *Updating Partition Statistics* on page 396
- *Deleting Partition Statistics* on page 397
- *Using a Hash Partition* on page 389
- *Using a Range Partition* on page 390
- *Using a Round-Robin Partition* on page 392
- *Enabling Semantic-based Partitioning* on page 392

Using a Round-Robin Partition

Create a new partition or change an existing partition using the round-robin strategy so that each partition contains an approximately equal number of rows.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table for the partition.
4. Select **Partitions**.
5. Choose to create a new partition or change an existing partition.
6. In the Select Partition Strategy window, choose the partitioning strategy **Round Robin**.
7. (Optionally) Specify the number of partitions.
8. In the Partition Specification window, specify the name of the partition and where the partition will reside.
   This partitioning strategy is random as no partitioning criteria are used. Round-robin-partitioned tables have no partition keys.
9. (Optional) Click **Summary** to verify your selected options.

See also

- *Updating Partition Statistics* on page 396
- *Deleting Partition Statistics* on page 397
- *Using a Hash Partition* on page 389
- *Using a Range Partition* on page 390
- *Using a List Partition* on page 391
- *Enabling Semantic-based Partitioning* on page 392

Enabling Semantic-based Partitioning

Enable semantic-based partitions to use hash, list, or range partition strategies.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the Perspective Resources window, select the servers on which the user table is to be created, then select **Administration Console**.
3. Click **ASE Servers**.
4. Click the server name, then select **Configure**.

5. In the Server Configuration window, turn on **enable semantic partitioning**, then click **Save All**.

See also
- *Updating Partition Statistics* on page 396
- *Deleting Partition Statistics* on page 397
- *Using a Hash Partition* on page 389
- *Using a Range Partition* on page 390
- *Using a List Partition* on page 391
- *Using a Round-Robin Partition* on page 392

Managing Table Statistics
Manage tables, indexes, and columns statistics which are used to estimate query costs.

Updating Table Statistics
Update column-related statistics such as histograms and densities.

The Adaptive Server cost-based optimizer uses statistics about the tables, indexes, and columns named in a query to estimate query costs. It chooses the access method that offers the lowest cost as determined by the optimizer. Cost estimates rely on accurate statistics.

Updating statistics at the table level applies only to user tables.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables** and select **User Tables**.

3. Select a table for which to update statistics.

Only one table can be selected for update.

4. Select **Update Statistics**.
The Update Statistics wizard appears.

5. In the Options window, choose one of the following:

- **Generate statistics for the leading column in each index**.
- **Generate statistics for all columns in each index**.

6. (Optional) Click **Use sampling**, then enter the sampling percentage.
The percentage to use when sampling depends on your needs. Test various percentages until you receive a result that reflects the most accurate information on a particular data set.

7. (Optional) Click **Use step number**, then enter the step numbers.

Increasing the number of steps beyond what is needed for good query optimization can affect performance, largely due to the amount of space that is required to store and use the statistics.

8. (Optional) In the Automatic Update window, click on **Update statistics only when datachange threshold is reached**, and enter a threshold value.
The datachange threshold determines when the amount of change in a table or partition has reached the predefined threshold.

9. Click **Finish** to apply the update statistics option selections.

See also
* *Updating Partition Statistics* on page 396
* *Updating Index Statistics* on page 396
* *Updating Column Statistics* on page 394

Deleting Table Statistics
The Delete Statistics wizard allows you to remove statistics for a table.

Delete statistics applies only to user and proxy tables.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   * **User Tables**
   * **Proxy Tables**
3. Select the table for which to update statistics.
4. Select **Delete Statistics**.
5. Click **Finish** to delete statistics on the selected table.

Updating Column Statistics
Update table statistics on individual columns.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   * **User Tables**
   * **Proxy Tables**
3. Select a table that contains the column.
4. Select **Properties**.

5. Click **Columns**.

6. Select the column for which to update statistics.

7. (Optional) Click **Use sampling**, then enter the sampling percentage.
The percentage to use when sampling depends on your needs. Test various percentages until you receive a result that reflects the most accurate information on a particular data set.

8. (Optional) Click **Use step number**, then enter the step numbers.
Increasing the number of steps beyond what is needed for good query optimization can affect performance, largely due to the amount of space that is required to store and use the statistics.

9. (Optional) In the Automatic Update window, click on **Update statistics only when datachange threshold is reached**, and enter a threshold value.
The datachange threshold determines when the amount of change in a table or partition has reached the predefined threshold.

10. Click **Finish** to apply the update statistics option selections.

See also

Deleting Column Statistics
The Delete Statistics wizard allows you to remove statistics from an individual column of a table.

Delete statistics applies only to user and proxy tables.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

- **User Tables**
- **Proxy Tables**

3. Select a table that contains the column.

4. Select **Properties**.

5. Click **Columns**.

6. Select the column for which to update statistics.

7. Select **Delete Statistics**.

8. Click **Finish** to delete statistics on the selected table.

Updating Index Statistics

Update table statistics for individual indexes.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

    *   **User Tables**
    *   **Proxy Tables**
3. Select a table that contains the index.
4. Select **Properties**.
5. Click **Indexes** .
6. Select the index for which to update statistics.
7. (Optional) Click **Use sampling**, then enter the sampling percentage.
    The percentage to use when sampling depends on your needs. Test various percentages until you receive a result that reflects the most accurate information on a particular data set.
8. (Optional) Click **Use step number**, then enter the step numbers.
    Increasing the number of steps beyond what is needed for good query optimization can affect performance, largely due to the amount of space that is required to store and use the statistics.
9. (Optional) In the Automatic Update window, click on **Update statistics only when datachange threshold is reached**, and enter a threshold value.
    The datachange threshold determines when the amount of change in a table or partition has reached the predefined threshold.
10. Click **Finish** to apply the update statistics option selections.

See also
*   *Updating Partition Statistics* on page 396
*   *Updating Table Statistics* on page 393
*   *Updating Column Statistics* on page 394

Updating Partition Statistics

Update table statistics for partitions.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.
2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

- **User Tables**
- **Proxy Tables**

3. Select a table that contains the partition.

4. Select **Properties**.

5. Click **Partitions** .

6. Select the partition for which to update statistics.

7. Select the type of statistics update.

8. (Optional) Click **Use sampling**, then enter the sampling percentage.
   The percentage to use when sampling depends on your needs. Test various percentages until you receive a result that reflects the most accurate information on a particular data set.

9. (Optional) Click **Use step number**, then enter the step numbers.
   Increasing the number of steps beyond what is needed for good query optimization can affect performance, largely due to the amount of space that is required to store and use the statistics.

10. Click **Finish** to apply the update statistics option selections.

See also

Deleting Partition Statistics
Delete table statistics from partitions.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

   - **User Tables**
   - **Proxy Tables**

3. Select the table on the partition.

4. Select **Properties**.

5. Click **Partitions**.

6. Select the partition from which to delete statistics.

7. Select **Delete Statistics**.

8. Click **Finish** to delete statistics on the selected partition.

See also

- *Updating Partition Statistics* on page 396
- *Using a Hash Partition* on page 389
- *Using a Range Partition* on page 390
- *Using a List Partition* on page 391
- *Using a Round-Robin Partition* on page 392
- *Enabling Semantic-based Partitioning* on page 392

Table Properties

Use the Tables Properties wizard to modify device usage, compression, permissions, cache, and the locking scheme.

Click **Properties** on your table to initiate the Properties wizard.

Wizard Option	Table Properties
General	- Name – you can specify a different table name. - Using cache – select the cache to bind to the table. - Identity gap – specify how ID numbers are allocated in memory. For example, a value of 10 indicates ID numbers are allocated in memory in blocks of 10. - Data compression – specify the type of data compression. - LOB compression – specify the level of compression. - Enable incremental transfer – allows you to transfer data incrementally, and, if required, to a different product. The incremental transfer feature must be available on the selected server.
Usage	- Usage – assigns space allocations for a table on a particular segment and segments to a device. Objects cannot grow beyond the space available in the segment's device. See *Placing a Table on a Segment* on page 381. - Show – select the units of measurement.
Permission	Permission options – grant or remove table permissions for users, groups, or roles. See *Setting Permissions for a Table* on page 387.

Wizard Option	Table Properties
Lock Scheme	Lock Scheme – specify the locking scheme to set how much data is locked at one time. See *Setting the Table Locking Scheme* on page 381.
	For more information about locking scheme, see Granularity of locks and locking schemes in *Performance and Tuning: Locking*.
Data	Displays the table data or table contents.
Referenced By	Displays the name, object type, and owner of objects that reference the specified table.
References	Displays the name, object type, and owner of objects of the specified reference.
Columns	Displays each column belonging to the table. Clicking the column name opens the properties window for the selected column. See *Column Properties* on page 394.
Indexes	Displays each index belonging to the table. Clicking the index name opens the properties window for the selected index. See *Index Properties* on page 401.
Triggers	Displays each trigger belonging to the table. Clicking the trigger name opens the properties window for the selected trigger. See *Trigger Properties* on page 402.
Foreign Keys	Displays each foreign key belonging to the table. Clicking the foreign key name opens the foreign key window for the selected index. See *Foreign Key Properties* on page 402.
Check Constraints	Displays each check constraint belonging to the table. Clicking the check constraint name opens the properties window for the selected check constraint. See *Check Constraint Properties* on page 402.
Partitions	Displays each partition for the table. Clicking the partition name opens the properties window for the selected partition. See *Partition Properties* on page 403.

See also
- *Column Properties* on page 400
- *Index Properties* on page 401
- *Trigger Properties* on page 402
- *Foreign Key Properties* on page 402
- *Check Constraint Properties* on page 402
- *Partition Properties* on page 403

Column Properties

Use the Columns Properties wizard to change permissions, create check constraints, specify encryption keys, and bind rules and defaults to columns.

Wizard Option	Column Properties
General	• Name – specify a different table name. • Datatype – change the datatype of the column, and depending on the datatype, the width and scale. • Primary key – constrains the values in the indicated column or columns so that no two rows have the same value, and so that the value cannot be NULL. • Allow nulls – specifies that Adaptive Server assign a null value if a user does not provide a value. • Identity – indicates that the column has the IDENTITY property. Each table in a database can have one IDENTITY column with a datatype of either: exact numeric and a scale of 0, or of the integer datatypes, including signed or unsigned bigint, int, smallint, or tinyint. • Object storage specifier – specifies whether a Java-SQL column is stored separately from the row (off row) or in storage allocated directly in the row (in row). • Data compression – supported only on user tables, and in version 15.7 and later.
Rules and Defaults	• Default – specify a default value that appears in the column if no value is entered for an insertion or update. • Rule binding – bind rules to columns to provided criteria against which Adaptive Server checks data entered for an insertion or update. See *Binding Defaults and Rules to a Column* on page 380.
Check Constraints	Creates filters that data must pass through before the data can be inserted into a table. See *Creating a Check Constraint* on page 378.
Permissions	You can grant and revoke permissions on a column or a table. See *Setting Permissions* on page 387.
Encryption	You can specify an encryption key for column encryption and optionally a default value when you do not have decrypt permission. See the **Encrypted Columns Users Guide**.

See also
- *Table Properties* on page 398
- *Index Properties* on page 401
- *Trigger Properties* on page 402
- *Foreign Key Properties* on page 402

- *Check Constraint Properties* on page 402
- *Partition Properties* on page 403

Index Properties

Use the Index Properties wizard to modify cache bindings, specify a device segment, and change index values.

Wizard Option	Index Properties
General	- Name – specify a different index name. - Unique – prohibits duplicate index values. - Clustered – physical order of rows on the current database device to be the same as the indexed order of the rows. - Suspect – indicates the integrity of the index is suspect. - Using cache – specifies the current cache binding. - Bind to – change the cache binding.
Columns	Displays the columns used in the index.
Miscellaneous	- Segment – change the segment on which the index is placed. - Duplicate keys – indicates if duplicate keys are allowed. - Duplicate rows – indicates if duplicate rows are allowed. - Data presorted – indicates the index data has been presorted. - Cache strategy – you can specify the Adaptive Server strategy for determining where in cache to place data pages when reading in new data. You can also choose to prefetch index pages by performing large I/Os of up to eight data pages simultaneously. - Rows per page – limits the number of rows on data pages and the leaf-level pages of indexes. - Reserve page gap – specifies a ratio of filled pages to empty pages to be left during extent I/O allocation operations. - Fill factor – specifies how full Adaptive Server makes each page when it creates a new index on existing data.
Index Partitions	Displays the name, segment, and creation date.

See also
- *Table Properties* on page 398
- *Column Properties* on page 400
- *Trigger Properties* on page 402
- *Foreign Key Properties* on page 402
- *Check Constraint Properties* on page 402
- *Partition Properties* on page 403

Trigger Properties

The Trigger Properties wizard shows the selected trigger options and objects referenced by the trigger.

Wizard Option	Trigger Properties
General	Shows the select trigger options.
Referenced by	Displays the name, type, and owner of the objects that are referenced by the specified trigger.

See also

- *Table Properties* on page 398
- *Column Properties* on page 400
- *Index Properties* on page 401
- *Foreign Key Properties* on page 402
- *Check Constraint Properties* on page 402
- *Partition Properties* on page 403

Foreign Key Properties

The Foreign Key Properties wizard shows current foreign key options and the matching primary keys.

Wizard Option	Foreign Properties
General	Shows the selected foreign key options.
Columns	Shows the defined foreign keys and the primary keys to which the foreign keys applies. See *Creating a Foreign Key* on page 377.

See also

- *Table Properties* on page 398
- *Column Properties* on page 400
- *Index Properties* on page 401
- *Trigger Properties* on page 402
- *Check Constraint Properties* on page 402
- *Partition Properties* on page 403

Check Constraint Properties

The Check Constraint Properties wizard shows the check constraint definitions.

Wizard Option	Check Constraint Properties
General	• Name, Owner, Creation date – shows the check constraint properties. • Check Constraint – shows the check constraint expression or condition that values must pass before being inserted into the table. See *Creating a Check Constraint* on page 378

See also
- *Table Properties* on page 398
- *Column Properties* on page 400
- *Index Properties* on page 401
- *Trigger Properties* on page 402
- *Foreign Key Properties* on page 402
- *Partition Properties* on page 403

Partition Properties
The Partition Properties wizard shows the partition name, strategy, and type.

Wizard Option	Partition Properties
General	Shows the partition properties including the name, strategy, and type of partition. You can change the segment on which the index is placed and the type of data compression. See *Creating a Partition* on page 389.

See also
- *Table Properties* on page 398
- *Column Properties* on page 400
- *Index Properties* on page 401
- *Trigger Properties* on page 402
- *Foreign Key Properties* on page 402
- *Check Constraint Properties* on page 402

Deleting a Table
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your table.
3. Click the Name field of the table you want to delete.
4. Select **Delete** from the menu.

5. Choose to delete the table.

6. Confirm the deletion.

7. Click **Finish**.

Thread Pools

Monitor threads or manage thread pools into groups of CPU resources.

Monitor Threads

Monitor Adaptive Server threads.

To monitor threads, you must start Adaptive Server in threaded mode.

Identifying the Threads in a Thread Pool

Find the threads belonging to a certain thread pool, and the associated kernel task name.

Thread pools can only be configured in Adaptive Server 15.7 and later versions.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.

2. In the left pane, select **Threads**.

3. Select the thread from the list of threads.
You see the thread pool name, and all the threads belonging to the pool.

4. Select **Tasks**.
You see the kernel task name associated with the thread pool.

See also
* *Creating a Thread Pool* on page 405
* *Thread Pool Properties* on page 406

Thread Statistics and Details

Interpret Adaptive Server thread information.

In a cluster environment, threads created on instances of a cluster are displayed as an Instance Thread in the right hand pane of the Threads screen.

Table 43. Tabs

Details	Displays information about affinity and number of ticks in the selected thread, including:
	• Total number of ticks • Number of idle ticks • Number of sleeping ticks • Number of busy ticks Also displays page faults and operating system context switches with the current thread, including: • Number of minor and major page faults • Operating system thread ID and alternative thread ID • Number of voluntary and forced context switches
Thread CPU Utilization	Displays graphs depicting user and system CPU utilization.
	Note: A data collection job must be scheduled to display the graphs.
Tasks	Displays a list of all the kernel task names and IDs associated with thread pools.

See also

Manage Thread Pools
Use thread pools to group CPU resources and execute Adaptive Server tasks associated with that thread pool.

Creating a Thread Pool
Group Adaptive Server engines into thread pools.

Prerequisites
The kernel mode must be set to threaded to create a thread pool. You can change the kernel mode on the Server Configuration screen, or by executing the SQL command **sp_configure** at the server level from the Administration Console.

sp_configure "kernel mode", 0, threaded

You must restart Adaptive Server for the change to take affect.

Task

Thread pools are groups of resources, such as engines, that execute user tasks, run specific jobs such as signal handling, and process requests from a work queue. Adaptive Server supports both system-defined and user-defined thread pools.

Note: Thread pools can only be configured in Adaptive Server 15.7 and later.

1. In the left pane of the Administration Console, expand **ASE Servers > Performance**, then select **Thread Pools**.
2. Select **New**.
   You see the Add Thread Pools wizard.
3. Select an Adaptive Server that is configured to run in threaded mode.
4. (Optional) For a cluster environment, click **Create this as local thread pool**, then select the cluster instance on which the thread pool will be created.
5. Specify the name of the thread pool you want to create.

   Note: You cannot name thread pools starting with a **syb_** prefix since that is reserved for system thread pools.

6. Specify the number of threads. The maximum number of threads you can configure cannot exceed the value of **max online engines**.
7. Specify the thread pool idle time out in microseconds.
8. (Optional) Provide a description for the thread pool.
9. (Optional) Click **Preview** to see the SQL statements for your command.
10. (Optional) Click **Summary** to verify your selected options.

See also
- *Creating Execution Classes* on page 310
- *Creating Engine Groups* on page 295
- *Thread Pool Properties* on page 406
- *Execution Classes Properties* on page 310
- *Engine Groups Properties* on page 295
- *Thread Statistics and Details* on page 404

Thread Pool Properties
View properties of thread pools in an Adaptive Server.

In the Thread Pools screen, select **Properties** on your thread pool name to initiate the Properties wizard and modify these properties:

Wizard Option	Thread Pool Properties
General	• Name and type – Adaptive Server supports two types of threads: Engine (or multiplexed) or Run to completion (RTC) threads. User created thread pools are always multiplexed. • Instance name (cluster environment only) – only displayed when the thread pool is created on an instance of the cluster rather than the entire cluster. • Thread count – you can increase the thread count up to a maximum value of **max online engines** configuration parameter. • Idle Time Out – set to: 　• 0 – threads change to sleep mode if no work is available. 　• -1 – threads never change to sleep mode even if no work is available. • Description – add a description for the thread pool.
Execution Classes	Sybase Control Center displays the execution classes that are associated with each user or system thread pool. You cannot modify their properties.

See also
- *Execution Classes Properties* on page 310
- *Engine Groups Properties* on page 295
- *Thread Statistics and Details* on page 404
- *Creating a Thread Pool* on page 405
- *Creating Execution Classes* on page 310
- *Creating Engine Groups* on page 295

Transactions

Monitor Transactions
Monitor active Adaptive Server transactions.

Identifying a Transaction's Process
Get information about a currently running transaction, including the process that initiated the transaction.

1. In the Perspective Resources view, select the server to monitor, click the drop-down arrow, and select **Monitor**.
2. In the left pane, select **Transactions**.

   You can also display the Transactions screen by clicking a SPID link on Running Processes tab of the Databases screen.

3. Locate the transaction in the Transactions table.

4. In the SPID column, click the SPID of the process associated with your transaction. (The ID number is a link.)
The Processes screen appears, displaying information about your transaction's parent process.

See also
- *Process Statistics and Details* on page 324
- *Transaction Statistics and Details* on page 408

Transaction Statistics and Details
Interpret information about Adaptive Server transactions.

The Transactions table displays information about all active transactions on the selected Adaptive Server. Details include the name of the transaction, the login of the user who owns it, the application that launched the transaction, the process that initiated the transaction (SPID column), the transaction's start time, the name of the host where the transaction is running, and the database it is running against. (If a transaction affects more than one database, the table does not display them all—it shows the transaction's current database and the process that started the transaction.)

Note: In Adaptive Server cluster configurations, information in the Transactions table is grouped by cluster instances.

When you select a transaction in the table, the User Log Cache Usage tab at the bottom of the screen displays statistics about the user log cache for the selected transaction. Details include bytes written, number of flushes and full flushes, maximum cache usage (in bytes), and current usage (in bytes).

See also
- *Process Statistics and Details* on page 324
- *Identifying a Transaction's Process* on page 407

User-Defined Datatypes
From the Perspective Resources view, use the Administration Console option on your Adaptive Server to create user-defined datatypes and view their properties.

Manage User-Defined Datatypes
Add, delete, or modify user-defined datatypes using the Sybase Control Center Administration Console.

Adding a User-Defined Datatype
Create a user-defined datatype using the Sybase Control Center Administration Console.

1. In the Administration Console, select **Server > Schema Objects > User Defined Datatypes**.
2. Select **New**.
   You see the Add User-defined Datatype wizard.
3. On the Introduction screen, select the server and database in which to create the datatype, and the datatype owner.
4. On the User Defined Datatype Name screen, enter the name of the datatype.
5. On the System Datatype screen, select the system datatype that the user-defined datatype is based on, and whether the datatype allows null or identity values.
   Depending on the system datatype, specify the size for your user-defined datatype.
6. On the Options screen, bind the datatype to a rule or default. Select **In Future Only** if you do not want existing columns to acquire the new rule or default.
7. (Optional) Click **Summary** to see the database options you have selected.

User-Defined Datatypes Properties

Use the Properties wizard to access information about user-defined datatypes, including bound rules and defaults.

Click **Properties** on your datatype to initiate the Properties wizard.

Wizard Option	User-Defined Datatypes Properties
General	• View the name, type, database, and owner of the user-defined datatype.

Advanced Options	• View or modify these options:
	  • The system datatype that the user-defined datatype is based on
	  • Whether the datatype allows null values
	  • Whether the datatype allows identity values
	  • Defaults and rules bound to the datatype

Referenced By	• View the name, type, owner, and properties of objects referenced by this user-defined datatype. Click an object, then click **Properties** to see the object properties.
References	• View the name, type, owner, and properties of objects that this user-defined datatype references. To view object properties, select an object, then click **Properties**.

Deleting a User-Defined Datatype

Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.
2. Navigate to your user-defined datatype.
3. Click the Name field of the user-defined datatype you want to delete.
4. Select **Delete** from the menu.
5. Choose to delete the user-defined datatype.
6. Confirm the deletion.
7. Click **Finish**.

Views

Manage views using the Sybase Control Center Administration Console.

Manage Views

Create, delete, modify, and administer views using the Administration Console.

Creating a View

Create a view using the Sybase Control Center Administration Console.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Views**.
2. Select **New**.
   You see the Create View wizard.
3. On the Introduction screen, select the server, database, and owner of the new view.
4. Enter the name of the view.
5. On the SQL Editor screen, provide the SQL statements for the view.
6. (Optional) Click **Preview** to see the SQL statements for your command.
7. (Optional) Click **Summary** to verify your selected options.

See also
- *View Properties* on page 410

View Properties

Use the view Properties wizard to access information on column datatype and permissions, and on database objects that reference, and are referenced by, the view.

Click **Properties** on your view to initiate the Properties wizard.

Wizard Option	View Properties
General	• View the name, type, database, owner, and creation date of the view.
SQL	• View the SQL statements for creating the view.
Columns	• View the name and type of all columns in the view.
Data	• View the data for each row in the view.

Permissions	• Grant and revoke permissions on a view to users, groups, or roles. Choose **with grant** to allow the grantee to further grant permissions to other users. You can grant permission to:
	  • Select the view
	  • Insert a row in the view
	  • Update a row in the view
	  • Delete a row in the view
	  • Decrypt a row in the view
	
	Note: If restricted decrypt permission is set, only a system security officer can grant decrypt permission.
	
	• See *Granting Permissions on Database Objects* on page 313 and *Revoking Permissions on Database Objects* on page 314.

Referenced By	• View the name, type, owner, and properties of objects referenced by this view. Click an object, then click **Properties** to see the object properties.
References	• View the name, type, owner and properties of objects that this view references. To view object properties, select an object, then click **Properties**.

See also
• *Creating a View* on page 410

Granting Permissions on Database Objects
Grant permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.

2. Click the **Name** field of the object on which to grant permissions.

3. Click the drop-down arrow and select **Properties**.

4. Select **Permissions**.

5. Click **Grant** to grant access permissions for the selected object.
   You see the Grant Permission wizard.

6. Select the type of grantee:

   • **Users**
   • **Groups**
   • **Roles**

7. Select one or more grantees.

8. Select the types of permissions allowed for the selected grantees.

9. (Optional) Click **Summary** to verify your selected options.

Revoking Permissions on Database Objects
Revoke permissions on objects such as stored procedures, functions, extended stored procedures, and views.

1. In the left pane of the Administration Console, expand **ASE Servers > Compiled Objects**, then select **Procedures**.

2. Click the **Name** field of the object on which to revoke permissions.

3. Click the drop-down arrow and select **Properties**.

4. Select **Permissions**.

5. Select the grantee, then click **Revoke** to revoke access permissions to the object.
   In the Revoke Permissions wizard, each type of permission and the current granted permissions are shown in cells.

6. Choose one of:

   • **Revoke all permission.**

   • Individual cells to revoke the currently granted permissions. The cell changes to show an "x", indicating that the permission type is no longer granted.

7. (Optional) Click **Preview** to see the SQL statements for your command.

Bulk Copying Data Out of a Table
Use bulk copy to copy table data to an external file.

1. In the Perspective Resources view, select the server on which the table resides, then click the drop-down arrow next to the server name and select **Administration Console**.

2. In the left pane of the Administration Console, expand **ASE Servers > Schema Objects > Tables**, then choose one of the following:

- **User Tables**
- **Proxy Tables**

3. Select one or more tables from which to copy data.
   If you select one table, you can select a different data file for each partition. If you select multiple tables, you can only provide one data file for all partitions for the table.

4. Select **Bulk Copy Out**.
   You see the Bulk Copy Out wizard.

5. In the Specify Data File window, enter the location.

6. Select the format for copying data.

7. In the Specify Copy Format window, choose the field and row parameters for the file to be copied from the table.

Deleting a View
Sybase Control Center helps you delete database objects, or the database itself.

1. In the Administration Console view, select **ASE Servers**.

2. Navigate to your view.

3. Click the Name field of the view you want to delete.

4. Select **Delete** from the menu.

5. Choose to delete the view.

6. Confirm the deletion.

7. Click **Finish**.

Troubleshoot Sybase Control Center for Adaptive Server

Solve problems that occur in Sybase Control Center for Adaptive Server.

Properties for Engine Groups Incorrectly Displayed

Problem: When you view the properties for an instance level engine group from the execution class General properties screen, the instance level engine group does not correctly identify the engine.

Solution: Open the instance level engine group properties screen directly from the Administration Console > Performance > Engine Groups screen.

Same Name Engine Groups Are Not Selectable

Problem: If you create engine group with the same name on multiple instances and then create an execution class or edits the properties page of an execution class to include a same named engine group, the engine groups with the same name will be displayed, but only one will be selectable.

Solution: Do not create create engine group with the same name on multiple instances.

Error: Unable to Format the Date String

Problem: While using Adaptive Server 15.5, the error log contains error messages about date string format.

If you use Adaptive Server 15.5, you see error messages in the Sybase Control Center server log such as this:

```
2009-11-17 09:13:14,493 ERROR [RMI TCP
Connection(12)-10.33.55.77] Unable to format the date string
00000 08:39:50.23 using the format yyyy/MM/dd HH:mm:ss.SSS
```

Solution:

1. Shut down the Unified Agent managing the Adaptive Server.
2. Rename the agent's UAF-2_5/plugins/com.sybase.ase/lib/
   ASEAgentPlugin.jar.

3. Copy `ASEAgentPlugin.jar` from `<scc-installation-directory>/
   SCC-3_2/plugins/ASEMAP` to `UAF-2_5/plugins/com.sybase.ase/
   lib`.
4. Start the Unified Agent that is managing the Adaptive Server.

"Number of Transactions" KPI Is Not Updated

Problem: Cannot update the KPI Number of Transactions

Solution: The Number of Transactions KPI is populated only if the monitored Adaptive Server version is 15.0.3 or later.

If your server version is 15.0.3 or later and the Number of Transactions KPI is not being updated, verify that you have installed the latest version of the Adaptive Server `installmaster` script on your server.

Invalid Connection Profile

Problem: During the login authentication step, a security error warns that the connection profile for this Adaptive Server is invalid.

Solution: Check the connection information stored by Sybase Control Center for this Adaptive Server. You can make sure the information is valid by using it with iSQL or Sybase Central to connect to the server.

See also
- *Registering an Adaptive Server* on page 114

Cannot Monitor Adaptive Server or Display Statistics Chart

Problem: In the Perspective Resources view, the **Monitor** and Statistics Chart context menu items for a monitored Adaptive Server are grayed out.

Solution: Make sure your user account is authenticated on the Adaptive Server. To monitor the Adaptive Server, you must also make sure that your account has been granted **mon_role** on the Adaptive Server.

See also
- *Role Assignment in Sybase Control Center for Adaptive Server* on page 119

Data on Screens or Charts Is Missing

Problem: Data on screens or charts is missing.

Solution 1: Schedule collections
Schedule collection jobs for Adaptive Server statistics collections. If collections are scheduled, make sure the start and stop times are set correctly.

- In the Adaptive Server monitor, the Overview, Devices, Engines, and Segments screens use data from collection_ase_all_client_kpis.
- Also in the Adaptive Server monitor, the Replication Agent screen uses data from collection_ase_rat.
- The statistics chart uses data from collection_ase_histmon and collection_ase_all_client_kpis.
- The heat chart uses data from collection_ase_availability, which is the default collection.

Solution 2: Check revalidation frequency
If appropriate collections are scheduled but data is missing on caches, devices, engines, or segments that were recently added to a registered Adaptive Server, check the value of the **revalidation_frequency** parameter. If the repository has not refreshed (revalidated) its list of monitored resources since the new resources were added, data on the new resources is not collected or displayed. Wait for the next revalidation to see the data.

Solution 3: Reset system clocks
If appropriate collections are scheduled but data is missing, truncated, or incomplete, compare the clock settings on the machines where Sybase Control Center and the client browser are running. Clocks that are out of sync by a few minutes cause screens and charts to display incomplete data. If the clocks' time difference exceeds the value of the chart trend period, displays that use the chart trend period contain no data. (This problem occurs when the time on one or both system clocks is incorrect. It is not caused by time zone differences—servers and clients can operate successfully in different time zones.)

Solution 4: Reset collection interval or trend period
If appropriate collections are scheduled but lines on graphs are missing, truncated, or incomplete, the values of the collection repeat interval and the chart trend period might be too close together. When these options are set to similar values, graphs sometimes contain only a single data point. Because Sybase Control Center needs at least two data points to draw a curve on a graph and a single data point is not displayed, graphs are empty. For example, if the collection interval is 12 minutes and the trend period is 15 minutes, graphs will display two data points for only a few minutes at a time, so they will appear to be blank more often than not.

To resolve the problem, decrease the collection repeat interval (set on the collection job in the scheduler) or increase the chart trend period (set on the Settings screen in the Adaptive Server monitor) so that multiple collection intervals occur within the trend period.

See also
- *Setting Adaptive Server Parameters in the Configuration File* on page 131
- *Setting Up Statistics Collection* on page 120

Adaptive Server Is Responding Slowly

Problem: A monitored Adaptive Server is responding slowly. How do you tell whether the problem lies in the network or the server?

Solution: On the Adaptive Server monitor for the Adaptive Server in question, select **Engines**. On the Engines screen, select an engine from the Engines table and check the **Engine CPU Utilization** graph. If the graph shows high activity for the period of slow response, the engine might be overloaded. If the graphs for all engines on this server show low activity, a network problem is more likely.

See also
- *Displaying Engine CPU Utilization* on page 294

Error: No Result Set for this Query

Problem: The agent log contains one or more instances of the error "No result set for this query."

Solution: This error occurs when queries executed by Sybase Control Center for Adaptive Server cannot be completed. To solve the problem, try increasing the space available to the tempdb on Adaptive Server. This example uses the **alter database** command to increase the size of tempdb by 20MB:

```
alter database tempdb on tempdb_dev=20
```

For more information on increasing the size of the tempdb, see:

- The **alter database** command in the Adaptive Server *Reference Manual: Commands*
- The chapter on temporary databases in the Adaptive Server *Performance and Tuning Series: Physical Database Tuning*

If the problems persists, contact Sybase technical support.

Collection Job for Adaptive Server Fails

Problem: A collection job for Adaptive Server may fail when the **number of open databases** is too low.

Solution: Modify the value of **number of open databases** by using either the Server Configuration screen of the Adaptive Server monitor, or these steps:

1. Log in to the Adaptive Server.
   ```
   isql -S<server_name> -U<sa user name> -P<sa password>
   ```
2. Run this command to display the current configuration value:
   ```
   sp_configure 'number of open databases'
   ```
3. Run this command to change the current configuration value:
   ```
   sp_configure 'number of open databases', <number>
   ```
   Add 10 to the current configuration value and substitute this number for *<number>*.

See also
- *Modifying Server Configuration Parameters* on page 368

Cannot Authenticate Server Configured with a Multibyte Character Set

Problem: If the Adaptive Server is configured to use a language that requires a multibyte character set such as Chinese, an attempt to authenticate the server fails if the correct character set is not specified in the connection profile for the server.

Solution: The character set for the connection profile can be specified in either of these ways:

- On the Resource Registration screen while registering the resource.
- On the Connection page of the Properties dialog for the Adaptive Server resource.

For example, if your Adaptive Server is using the Chinese language, it may be using character set **gb18030**. In this case, specify **gb18030** as the character set.

See also
- *Registering an Adaptive Server* on page 114

Database Objects Are Not Updated

Problem: Changes made to database objects are sometimes not visible in Sybase Control Center dialogs or screens.

Solution: Click **Refresh** on the Sybase Control Center screens to see the updated values for the database objects.

You may see this problem when you:

* Click **Finish** on a wizard, and do not see the updates (that should be generated by the wizard action) on your current screens.
* Create or update database objects outside of Sybase Control Center.

Features Are Not Enabled Although You Have sa_role

Problem: Some features are not enabled even though you are using a login account that has sa_role on the managed server.

Solution: If your login account was granted **sa_role** after you opened the ASE Monitor view, exit from the ASE Monitor and reauthenticate with the Adaptive Server. This causes Sybase Control Center to reconnect to the Adaptive Server and the new connection acquires the updated login privileges.

Alerts Are Configured But Do Not Fire

Problem: An alert is configured and the condition for the alert occurs, but the alert does not fire.

Solution: The collection for the KPI that the alert requires is not scheduled. If the alert is defined on one of the KPIs displayed in the historical charts in the ASE Monitor then the all_client_kpis collection must be scheduled. If the KPI is not one of the KPIs used by the ASE Monitor charts then the ase_histmon collection must be scheduled.

Error: No Data Was Found For Statistic

Problem: I selected a KPI in the Statistic Chart and clicked Graph Statistic but I get the error "No data was found for statistic myserver: ."

Solution: The error is because the KPI belongs to a collection that is not scheduled for this server. Schedule the collection.

See also
* *Key Performance Indicators for Adaptive Server* on page 125

Cannot Find Error Information For Monitor View

Problem: The ASE Monitor view is not responding and error information is needed.

Solution: Errors are reported in both the Sybase Control Center server log file in $SCC_HOME/SCC-3_2/log/agent.log, and the log file for the Adaptive Server component of Sybase Control Center in $SCC_HOME/SCC-3_2/plugins/ASEMAP/log/ASEMAP.log

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 102
* *Setting Up Security* on page 83

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 98
* *Configuring Memory Usage* on page 75

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 120

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 183
* *Graphing Performance Counters* on page 148

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 153

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 75

SCC Out of Memory 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 75

Glossary: Sybase Control Center for Adaptive Server

Glossary of Sybase Control Center terms related to Adaptive Server.

See the glossary in the Adaptive Server documentation for a complete list of Adaptive Server terms.

Adaptive Server – a server in the Sybase client/server architecture that manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory. Sybase Control Center can manage multiple Adaptive Servers.

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.

blocking – waiting for a lock; a task that needs to acquire a lock on a row, page, or table must wait, or block, if another process holds an incompatible lock on its target object.

cache – See data cache, procedure cache, or statement cache.

chart trend peiod – the period, in minutes, over which data is displayed in historical charts. Set the chart trend period on the Settings screen of the Adaptive Server monitor. Contrast with 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.

collection repeat interval – the period, in seconds, between successive repititions 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.

data cache – also called buffer cache and named cache. An area of memory within Adaptive Server that contains the images of database pages and the data structures required to manage the pages. Each cache is given an unique name that is used for configuration purposes. By default, Adaptive Server has a single cache named "default data cache." Caches configured by users are called user-defined caches.

device – in Adaptive Server, any piece of a disk or file in the file system used to store databases and their component objects.

engine – an instance of the Adaptive Server executable that can communicate with other Adaptive Server engines in shared memory. An Adaptive Server running on a uniprocessor machine always has one engine, engine 0. An Adaptive Server running on a multiprocessor machine can have one or more engines.

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.

heat chart – a graphical view of resource availability and selected performance and status metrics for all the registered resources in the current perspective.

index – a database object that consists of key values from data tables and pointers to the pages that contain those values. Indexes speed up access to data rows by pointing Adaptive Server to the location of a table column's data on disk.

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.

lock – a concurrency control mechanism that protects the integrity of data and transaction results in a multiuser environment. Adaptive Server applies table, page, and row locks to:

• Prevent two or more users from changing the same data at the same time
• Prevent processes from reading data that is in the process of being changed

managed resource – see resource.

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.

procedure cache – memory used for stored procedures, batch query plans, triggers, the statement cache, datachange tracking, query compilation, and other objects used during query execution.

query plan – the ordered set of steps required to carry out a SQL query, complete with the access methods chosen for each table. Query plans are chosen by the Adaptive Server optimizer.

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 server, agent, or other entity that can be monitored or administered by Sybase Control Center. Resources SCC can manage include Adaptive Server, Replication Server, Replication Agent, Mirror Replication Agent, Sybase IQ, and certain subcomponents.

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 Adaptive Server component of Sybase Control Center. Refreshing a screen redraws it with the most recent available data. Set the screen refresh interval on the Settings screen of the Adaptive Server monitor. See also collection repeat interval.

segment – space allocated on one or more database devices. Segments can be used to control the placement of tables and indexes on specific database devices.

semaphore – a simple internal locking mechanism that prevents a second task from accessing the data structure currently in use. Adaptive Server uses semaphores to protect transaction

logs, user log caches, and I/O devices. A semaphore is relevant only in symmetric multiprocessing (SMP) environments.

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.

singleton installation – a Sybase Control Center installation that runs a single SCC agent or server. Contrast with instance; see also shared-disk mode.

statement cache – memory used to store computed query plans. The statement cache is part of the procedure cache.

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.

Transact-SQL – the SQL dialect used in Sybase Adaptive Server.

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.

wait event – a condition that causes an Adaptive Server process to pause and wait for another event. Common wait events are waiting for disk I/O to complete, waiting on the Adaptive Server scheduler runnable queue for a CPU to become available, and waiting for another process's lock on a table to be released.

Index

T

V

W

X