



Sybase Control Center for Event Stream Processor

3.2.7 SP02

DOCUMENT ID: DC60012-01-0327-03

LAST REVISED: April 2013

Copyright © 2013 by Sybase, Inc. All rights reserved.

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase trademarks can be viewed at the Sybase trademarks page at <http://www.sybase.com/detail?id=1011207>. Sybase and the marks listed are trademarks of Sybase, Inc. ® indicates registration in the United States of America.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Java and all Java-based marks are trademarks or registered trademarks of Oracle and/or its affiliates in the U.S. and other countries.

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names mentioned may be trademarks of the respective companies with which they are associated.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

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

Contents

About Sybase Control Center for SAP Sybase Event Stream Processor	1
Stream Processor	1
User Interface Overview	2
Toolbar Icons	2
Status Icons	3
Common Display Options	4
Keyboard Shortcuts	7
Displaying the Versions of SCC Components	8
Style and Syntax Conventions	9
Accessibility Features	10
Sybase Control Center Accessibility Information	11
Get Started	13
Quick Start for an Evaluation	13
Get Started in a Production Environment	24
Deploying an Instance from a Shared Disk Installation	67
Enabling and Disabling Shared-Disk Mode	68
Shared-Disk Mode	69
sccinstance Command	70
Launching Sybase Control Center	74
Registering the ODBC Driver in Windows	75
Starting and Stopping Sybase Control Center in Windows	75
Starting and Stopping Sybase Control Center in UNIX	78
Configuring Memory Usage	82
scc Command	86
Logging in to Sybase Control Center	90
Logging out of Sybase Control Center	91
Setting Up Security	91
Security	92

Configuring Authentication for Windows	93
Configuring a Pluggable Authentication Module (PAM) for UNIX	94
Configuring an LDAP Authentication Module	95
Mapping Sybase Control Center Roles to LDAP or OS Groups	106
Encrypting a Password	108
Configuring Ports	109
Configuring the E-mail Server	111
Configuring the Automatic Logout Timer	112
Configuring Retrieval Thresholds for the Administration Console	113
User Authorization	114
Assigning a Role to a Login or a Group	115
Removing a Role from a Login or a Group	115
Adding a Group	116
Removing a Group	116
Adding a Login Account to a Group	117
Removing a Login Account from a Group	117
Adding a Login Account to the System	118
Removing a Login Account from the System	119
Modifying a User Profile	120
Logins, Roles, and Groups	120
Configure Sybase Control Center	123
Configuring Policies for Monitoring and Administering Event Stream Processor	124
Configuring Event Stream Processor for Monitoring .	126
Configuring Event Stream Processor for Administration	127
Registering an ESP Node	128
Importing Resources for Batch Registration	129
Authenticating an ESP Node	131
Update the Authentication Type	132
Clearing Authentication Parameters	136

Registering and Authenticating a Sybase Control Center Agent	136
Viewing Sybase Control Center Agent Connection Information	137
Parameters Required for Starting an ESP Node	138
Creating a Perspective	139
Adding a Resource to a Perspective	139
Authenticating a Login Account for a Managed Resource	140
Setting Up Statistics Collection	141
About Statistics	142
Event Stream Processor Data Collections	143
Key Performance Indicators for Event Stream Processor	147
Changing the Screen Refresh Interval	150
Creating an Alert	151
Event Stream Processor Alerts	153
Alert Types and Severities for Event Stream Processor	158
Alert-Triggered Scripts	159
Alert-Triggered Script Examples	160
Substitution Parameters for Scripts	161
Manage and Monitor Event Stream Processor	165
Displaying Resource Availability: the Heat Chart	165
Graphing Performance Counters: the Statistics Chart	166
Managing Workspaces	167
Updating Access Control	168
Nodes	169
Viewing Overview Statistics and Alerts for a Node	169
Viewing All Statistics for a Node	171
Viewing Statistics for Projects on a Node	172
Viewing Statistics for Streams on a Node	174

Viewing Statistics for Connections on a Node ...	175
Viewing Statistics for Adapters on a Node	176
Viewing Statistics for Publishers on a Node	177
Viewing Statistics for Subscribers on a Node ...	178
Starting a Node	179
Stopping a Node	180
Viewing Schema for a Stream	181
Clusters	181
Viewing Overview Statistics and Alerts for a Cluster	182
Viewing Topology for a Cluster	183
Viewing Nodes Within a Cluster	184
Viewing Statistics for Projects on a Cluster	186
Viewing Statistics for Streams on a Cluster	188
Viewing Statistics for Connections on a Cluster	189
Viewing Statistics for Adapters on a Cluster	190
Viewing Statistics for Publishers on a Cluster ...	191
Viewing Statistics for Subscribers on a Cluster .	192
Viewing Schema for a Stream	194
Projects	194
Viewing Projects	194
Starting a Project	195
Stopping a Project	196
Managing Projects	197
Adapters	198
Viewing Adapters	198
Starting an Adapter	199
Stopping an Adapter	200
Viewing File Activity for the Sybase IQ Output Adapter	201
Administration Console	202
Browsing and Managing Resources	202
Searching and Filtering Resources	203
Job Scheduling	204

Executing and Stopping a Data Collection Job .	205
Deleting a Data Collection Job	205
Resuming and Suspending a Data Collection Job	206
Adding a New Schedule to a Job	207
Viewing or Deleting a Schedule	208
Modifying the Data Collection Interval for a Job	208
Resuming and Suspending the Scheduler	209
Viewing the Job Execution History	209
Alerts	210
Types, Severities, and States	211
Viewing Alerts	212
Modifying an Alert	212
Testing an Alert-Triggered Script	213
Deleting an Alert	214
Alert Subscriptions	214
Alert Notifications	217
Log Files for Event Stream Processor	218
Viewing the SCC Agent for Event Stream Processor Log File	219
Viewing the Node Log File	220
Viewing the Project Log File	221
Manage Sybase Control Center	223
Resources	223
Unregistering a Resource	223
Adding a Resource to a Perspective	224
Removing a Resource from a Perspective	224
Modifying a Resource's Name and Connection Properties	225
Searching for Resources in the Resource Explorer	226
Perspectives	226
Creating a Perspective	227
Removing a Perspective	227

Renaming a Perspective	227
Views	228
Managing a View	228
Arranging View Layout in a Perspective	229
Instances	230
Enabling and Disabling Shared-Disk Mode	230
Deploying an Instance from a Shared Disk	
Installation	231
Refreshing or Converting an Instance	232
Removing an Instance	233
Shared-Disk Mode	234
sccinstance Command	235
Repository	240
Scheduling Backups of the Repository	240
Modifying the Backup Schedule	241
Forcing an Immediate Backup	242
Restoring the Repository from Backups	243
Configuring Repository Purging	244
Logging	245
Viewing the Sybase Control Center for Event	
Stream Processor Log	246
Modifying the Event Stream Processor Log	
Configuration	246
Viewing Sybase Control Center Server Logs ...	247
Viewing the Sybase Control Center Client Log	
.....	248
Changing the Logging Level	248
Logging or Message Levels	249
Changing Logging Configuration	250
Sybase Control Center Console	251
Console Commands	251
Troubleshoot Sybase Control Center for Event Stream	
Processor	255
Problems with Basic Sybase Control Center	
Functionality	255

Cannot Log In	255
Sybase Control Center Fails to Start	255
Browser Refresh (F5) Causes Logout	256
Alerts Are Not Generated	256
Performance Statistics Do Not Cover Enough Time	256
Resetting the Online Help	256
Data Collections Fail to Complete	257
Memory Warnings at Startup	258
SCC Out of Memory Errors	258
Statistics Do Not Display	258
Troubleshooting Tips	259
Glossary: Sybase Control Center for Event Stream	
Processor	261
Index	265

Contents

About Sybase Control Center for SAP Sybase Event Stream Processor

Sybase[®] Control Center for SAP[®] Sybase Event Stream Processor is a Web-based tool for managing and monitoring ESP Server nodes, clusters, projects, and other components of the Event Stream Processor environment.

Sybase Control Center supports SAP Sybase Event Stream Processor version 5.1 or later.


The SCC architecture allows multiple administrators using Web clients to monitor and control all the Event Stream Processor components in an enterprise through one or more SCC servers. SCC for Event Stream Processor 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 SCC for Event Stream Processor to track a variety of performance metrics, gathering statistics that over time will give you powerful insight into patterns of use. 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 Event Stream Processor environment has performed in the past helps you ensure that Event Stream Processor meets your needs in the future.

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


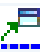





Toolbar Icons

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

Table 1. 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.

Icon	Name	Description
	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 2. 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




Icon	Status	Description
	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 3. 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

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











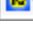
The Perspective Resources view, Resource Explorer, Administration Console, Alert Monitor, heat chart, and other views in Sybase Control Center—including those in product modules—use a tabular grid format to display information about managed resources. You can use options provided by the grid format to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

Table 4. Column Sorting Options

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

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

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

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


Filter by Column

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

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


You can filter on multiple columns; for example, in a listing of servers, use the Status column to display only running servers, then use the Version column to display 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 at the top of the SCC window.

Keyboard Shortcuts

Frequently used key sequences for the Sybase Control Center Web interface.

Key Sequence	Action
Ctrl-Alt-F12	Pull down the first menu in the topmost view or in the SCC menu bar. Repeat to toggle between the two first menus.
Ctrl-Alt-Shift-F12	Pull down the first menu (Application) in the SCC menu bar.
Space	Select the highlighted option—equivalent to clicking the mouse.
Escape	<ul style="list-style-type: none"> • Release a drop-down menu • Exit an editable cell or field • Close a window
Arrow keys	<ul style="list-style-type: none"> • Highlight the next list item or menu option in the indicated direction. For example, the down arrow highlights the next item down in a menu; the right arrow highlights an item in the menu to the right. • In a tree hierarchy, the right arrow expands a node; the left arrow collapses it.
Tab	<ul style="list-style-type: none"> • In a view, highlight the next item in the tab order. (Tab order progresses through the accessible fields in a left-to-right, top-to-bottom fashion, starting at the upper left.) • In a two-pane view, jump from the tree hierarchy in the left pane to the right pane. • In a view that includes a table or grid display, press Tab twice to highlight the table, then press the down-arrow to enter it.
Shift-Tab	<ul style="list-style-type: none"> • In a view, highlight the previous item in the tab order. • In a two-pane view, jump from the right pane back to the tree hierarchy in the left pane.
Home	Highlight the first item in the active view (or the active section of a view), for example the first row in a table.
End	Highlight the last item in the active view (or the active section of a view), for example the last row in a table.

Key Sequence	Action
In the SCC menu bar, View > Select > <your view>	Select an open view and bring it to the front.
Ctrl-Alt Arrow key	Move the selected view in the indicated direction.
Ctrl-Alt +	Increase the size of displayed text.
Ctrl-Alt -	Decrease the size of displayed text.
F11	Enable or disable the browser's full-screen mode.
In the SCC menu bar, Appli-cation > Display > Full Screen	Enable or disable Sybase Control Center's full-screen mode.

Sybase Control Center is built on Adobe Flex. For complete information about Adobe Flex keyboard shortcuts, see http://livedocs.adobe.com/flex/3/html/help.html?content=accessible_5.html.

Displaying the Versions of SCC Components

View a list of components installed in Sybase Control Center and their versions.

Check the versions of the product modules in your SCC installation to determine whether your installation is up to date. SCC release bulletins list supported product module versions. You can find SCC release bulletins on the Product Documentation web site at <http://sybooks.sybase.com/sybooks/sybooks.xhtml?prodID=10680>

1. Log in to SCC and select **Help > About Sybase Control Center**.
2. Compare the versions of product modules (listed as management agent plug-ins) against the versions published in the most recent *Sybase Control Center Release Bulletin*.

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 5. Style Conventions

Key	Definition
monospaced (fixed-width)	<ul style="list-style-type: none"> • SQL and program code • Commands to be entered exactly as shown • File names • Directory names
<i>italic monospaced</i>	In SQL or program code snippets, placeholders for user-specified values (see example below).
<i>italic</i>	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 6. 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 effectively with versions of JAWS for Windows screen reading software before version 11, 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 Event Stream Processor for Monitoring*

To enable users to monitor ESP node and cluster activity using Sybase Control Center, map a native OS, preconfigured login, LDAP, or SAP BI group to the espMonitorRole role.

5. *Registering an ESP Node*

Make Sybase Control Center aware of an ESP node (acting as a cluster manager, controller, or both) and its connection information by registering it as a resource.

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: the Heat Chart*

Get Started

Use the heat chart to view the status and availability of servers in the current perspective.

8. *Viewing Overview Statistics and Alerts for a Node*

View alerts for and monitor the performance of the selected ESP node by viewing overview statistics such as CPU, memory, and thread usage.

9. *Viewing Overview Statistics and Alerts for a Cluster*

View alerts for and monitor the performance of the selected ESP cluster by viewing overview statistics such as CPU, memory, and thread usage history.

See also

- *Get Started in a Production Environment* on page 24

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 provides a default login account, `sccadmin`, for initial configuration and setting up permanent authentication. The password is set during installation.

Note: The `sccadmin` account and the preconfigured user login module on which it is 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 as soon as you have set up and tested authentication, and that you change the password on the sccadmin account if you do not plan to set up and test authentication right away.

-
5. (Optional) Change the password or disable sccadmin—see the *Sybase Control Center Installation Guide* for instructions.

Configuring Event Stream Processor for Monitoring

To enable users to monitor ESP node and cluster activity using Sybase Control Center, map a native OS, preconfigured login, LDAP, or SAP BI group to the espMonitorRole role.

The ESP node uses your corresponding authentication provider to determine which groups a user belongs to and then uses the `csi_role_mapping.xml` file to map these groups to the appropriate roles. This file is located in the `ESP-5_1\security` directory, which by default is installed in the `Sybase` directory. You may choose to map an existing group to espMonitorRole or create a new group.

1. Open `ESP-5_1\security\csi_role_mapping.xml`.

The `<Mapping>` element represents a mapping for a logical role. The `<LogicalName>` element represents the role that SCC checks for this mapping. There can only be one `<LogicalName>` per mapping. Do not modify this element.

2. Set the LDAP group within the `<MappedName>` element. This element represents the group you wish to map to the logical role. You can include more than one `<MappedName>` element for a mapping.

You can have the same `<MappedName>` element for two different `<LogicalName>` elements.

Here is an example of a mapping where the group "Administrators" maps to espMonitorRole:

```
<Mapping>
<LogicalName>espMonitorRole</LogicalName>
<MappedName>Administrators</MappedName>
</Mapping>
```

Here is an example of a mapping where several groups are mapped to espMonitorRole:

```
<Mapping>
<LogicalName>espMonitorRole</LogicalName>
<MappedName>IT</MappedName>
<MappedName>Developers</MappedName>
<MappedName>Operators</MappedName>
</Mapping>
```

Registering an ESP Node

Make Sybase Control Center aware of an ESP node (acting as a cluster manager, controller, or both) and its connection information by registering it as a resource.

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

Table 7. New Resource Type Details

Field	Description
Resource Name	(Required) Name of the resource to register. Set a name which will make the ESP node easily identifiable.
Resource Type	Select ESP Node .
Description	(Optional) A brief description to help you identify the resource.

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

Table 8. New Resource Connection Details

Field	Description
Host Name	(Required) Name of host on which the ESP node runs. Default value is localhost.
Port Number	(Required) Port number for the ESP node.
SSL is Enabled (Y/N)	(Required) Specify Y if SSL is enabled on the ESP Server, and N if it is not.

5. Click **Next**.
6. Enter authentication information based on the authentication type you specified while installing Event Stream Processor. This information is stored in the `ProductModule.xml` file, and must match the authentication type of the cluster in Event Stream Processor.

Authentication Type	Fields
Native OS	User name and password.
LDAP	User name and password.
Kerberos	KDC, Realm, Service Name, User name, and Ticket Cache Location. See the <i>Event Stream Processor Administrators Guide</i> for more information.

Authentication Type	Fields
RSA	User name, Password, and Keystore location.

- Or if you prefer not to authenticate now, click **I do not want to supply authentication information**.
7. Click **Next**.
 8. (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.
 9. (Optional) Click **Open the resource explorer to view this new resource**. (This option is not present when the Resource Explorer is open.)
The resource is added to the Resource Explorer even if you choose not to view it.
 10. Click **Finish**.
The ESP node is registered. If you have chosen to authenticate the resource and the authentication is successful, the other active nodes that belong to the same cluster as this node are registered and authenticated automatically because in Event Stream Processor, authentication is performed on a cluster-wide basis rather than on a per-node basis. These other nodes are named "<cluster_id>_<node_id>", by default, and "<cluster_id>_<node_id>(manager)" is for both manager type and dual (manager and controller) type nodes. You can edit the names of these nodes under resource properties.

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: the Heat Chart

Use the heat chart to view the status and availability of servers 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.

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.

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.

Get Started

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**.

Viewing Overview Statistics and Alerts for a Node

View alerts for and monitor the performance of the selected ESP node by viewing overview statistics such as CPU, memory, and thread usage.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Overview**.

Name	Statistics
Server	<p>Node name - name of the ESP node.</p> <p>Host - host that server is currently running on.</p> <p>Node version - version of the ESP node.</p> <p>Platform - platform of the machine that the node is running on.</p> <p>Node type - whether node is a manager or controller.</p> <p>State - valid values are running or unknown.</p>

Name	Statistics
Alerts	<p>This is the header table containing all alerts that have fired for the selected ESP node after you open the ESP Node Monitor view. If you log off without closing this view, alerts that have fired since you logged back on display.</p> <p>Time - when the alert is triggered.</p> <p>Alert Name - name of the alert. This is based on the KPI.</p> <p>Resource - the resource for which the alert is triggered.</p> <p>Severity - alert severity rating. Possible severity ratings are Normal, Warning, or Critical, and are based on ranges of values you specified when setting the alert threshold.</p> <p>Value - the KPI value. The alert is triggered when the KPI value falls within the range of values you specified when setting the alert threshold.</p> <p>Threshold - the range of values you assigned to alert severity ratings. For example, if the low value for the Normal rating is 0 and the high value 100, the threshold is the range of 0 to 100.</p>

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The data on the graph starts displaying from the time you open the ESP Node Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.
Memory Usage History	Line graph displaying total memory usage over time, in kilobytes (KB). The data on the graph starts displaying from the time you open the ESP Node Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.

Statistics	Description
Thread Usage History	Line graph displaying number of threads used over time. The data on the graph starts displaying from the time you open the ESP Node Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.

Viewing Overview Statistics and Alerts for a Cluster

View alerts for and monitor the performance of the selected ESP cluster by viewing overview statistics such as CPU, memory, and thread usage history.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Overview**.

Name	Statistics
Alerts	<p>This is the header table containing all alerts that have fired for the nodes in the cluster after you open the ESP Cluster Monitor view. If you log off without closing this view, alerts that have fired since you logged back on display.</p> <p>Time - when the alert is triggered.</p> <p>Alert Name - name of the alert. This is based on the KPI.</p> <p>Resource - the resource for which the alert is triggered.</p> <p>Severity - alert severity rating. Possible severity ratings are Normal, Warning, or Critical, and are based on ranges of values you specified when setting the alert threshold.</p> <p>Value - the KPI value. The alert is triggered when the KPI value falls within the range of values you specified when setting the alert threshold.</p> <p>Threshold - the range of values you assigned to alert severity ratings. For example, if the low value for the Normal rating is 0 and the high value 100, the threshold is the range of 0 to 100.</p>

These statistics display:

Statistics	Description
CPU History	<p>Line graph displaying the average CPU user usage, CPU system usage, and total CPU usage across the cluster. The data on the graph starts displaying from the time you open the ESP Cluster Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.</p>

Statistics	Description
Memory Usage History	Line graph displaying average total memory usage across the cluster, in kilobytes (KB). The data on the graph starts displaying from the time you open the ESP Cluster Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.
Thread Usage History	Line graph displaying the average number of threads used across the cluster. The data on the graph starts displaying from the time you open the ESP Cluster Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.

Get Started in a Production Environment

Perform a complete setup 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. *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.
4. *Logging in to Sybase Control Center*
 Enter the Sybase Control Center Web console.
5. *Setting Up Security*
 Configure login authentication and map roles.
6. *Configuring the E-mail Server*

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

7. *Configuring the Automatic Logout Timer*

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

8. *User Authorization*

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

9. *Configure Sybase Control Center*

Set up Sybase Control Center for Sybase Event Stream Processor.

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

Get Started

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 32
- *Starting and Stopping Sybase Control Center in UNIX* on page 35

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 reenables, 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, SAP Sybase ESP clusters, or SAP 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	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

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:

Get Started

```
$ 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, which includes the management UI) and the Sybase Control Center agent that runs on each product server managed by SCC. When you install SCC and the SCC agent in the same directory by selecting both options in the installer, you start and stop them together—by executing a single command or controlling a single service. This topic applies both to singleton installations (which do not use a shared disk) and to 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, the SCC agent, or both when they are installed together:**

- (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.)

- 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, the SCC agent, or both when they are installed together:**

- 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:**

- Open the Windows Control Panel.
- 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 Sybase Control Center, the SCC agent, or both) from the Windows command line:**
 - a) To start the service, enter the **net start** command.

For a singleton installation:

```
net start "sybase control center 3.2.8"
```

```
The Sybase Control Center 3.2.8 service is starting.....  
The Sybase Control Center 3.2.8 service was started  
successfully.
```

For an instance, include the instance name (Boston-1 in this example) in parentheses:

```
net start "sybase control center 3.2.8 (Boston-1)"
```

```
The Sybase Control Center 3.2.8 (Boston-1) service is  
starting.....  
The Sybase Control Center 3.2.8 (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 3.2.8"
```

```
The Sybase Control Center 3.2.8 service is stopping.....  
The Sybase Control Center 3.2.8 service was stopped  
successfully.
```

For an instance, include the instance name (Boston-1 in this example) in parentheses:

```
net stop "sybase control center 3.2.8 (Boston-1)"
```

```
The Sybase Control Center 3.2.8 (Boston-1) service is  
stopping.....
```

```
The Sybase Control Center 3.2.8 (Boston-1) service was
stopped successfully.
```

See also

- *Deploying an Instance from a Shared Disk Installation* on page 25
- *Configuring Memory Usage* on page 39

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, which includes the management UI) and the Sybase Control Center agent that runs on each product server managed by SCC.. When you install SCC and the SCC agent in the same directory by selecting both options in the installer, you start and stop them together—by executing a single command or controlling a single service. This topic applies to both singleton installations (which do not use a shared disk) 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 (or both, when they are installed together) 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 (or both, when they are installed together) 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 (or both, when they are installed together) .**

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 38 or step *5.d* on page 38.

Get Started

- 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 output 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 <ul style="list-style-type: none"> • <code>jvmopt=-Xmx</code> – if you are running Sybase Control Center as a Windows service • <code>SCC_MEM_MAX</code> – if you are running SCC as a UNIX service • <code>SCC_MEM_MAX</code> – if you are starting SCC from the command line 	<ul style="list-style-type: none"> • You need to prevent Sybase Control Center from using more than a given amount of memory • Sybase Control Center fails to start and may display an error: <code>Could not create the Java Virtual machine.</code> • An <code>OutOfMemory</code> error says Sybase Control Center is out of heap space • A warning message about system memory appears during the start process • The machine where Sybase Control Center is installed has less than 4GB of memory. (Starting Sybase Control Center on a machine with less than 4GB of memory triggers the startup warning message about system memory.) 	<p>On machines with less than 4GB of memory, set maximum memory to 256MB or more.</p> <p>Default value: none. (On machines with 4GB or more of memory, maximum memory is set dynamically and is effectively limited only by the amount of system memory available.)</p>

Modify this value	When	Guidelines
<p>Permanent memory</p> <ul style="list-style-type: none"> • <code>jvmopt=-XX:MaxPermSize</code> – if you are running Sybase Control Center as a Windows service • <code>SCC_MEM_PERM</code> – if you are running SCC as a UNIX service • <code>SCC_MEM_PERM</code> – if you are starting SCC from the command line 	<p>An OutOfMemory error says Sybase Control Center is out of permanent generation space</p>	<p>Increase by 32MB increments. If you reach a value equal to twice the default and still see the OutOfMemory error, contact Sybase technical support.</p> <p>Default value: 128MB</p>

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 Sybase Control Center 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

- *Starting and Stopping Sybase Control Center in Windows* on page 32
- *Logging in to Sybase Control Center* on page 42

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 a Sybase Control Center Windows Service* on page 41
- *Changing a Memory Option for an SCC UNIX Service* on page 41

Changing a Memory Option for a Sybase Control Center 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 Sybase Control Center properties file:

```
<Sybase Control Center-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 40
- *Changing a Memory Option for an SCC UNIX Service* on page 41

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
```

Get Started

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

See also

- *Changing a Memory Option on the Command Line* on page 40
- *Changing a Memory Option for a Sybase Control Center Windows Service* on page 41

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.

Only one login session per account is permitted at a time; multiple users cannot be logged in to the same account simultaneously.

Note: When logging in to a newly installed Sybase Control Center for which secure authentication has not been configured, use the sccadmin account—the password is set during installation. 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@sap`, not `fred@sap.com`.

See also

- *Configuring Memory Usage* on page 39

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 *Quick Start for an Evaluation* on page 13.

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. Configuration is required only if you have upgraded from an older version of Sybase Control Center and no longer want to use the older version's authentication settings; if you do not want to use Windows for authentication; or if you want to create login accounts manually. Sybase recommends that you allow SCC to create accounts automatically.

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.

4. *Configuring an LDAP Authentication Module*

Configure an LDAP authentication module for Sybase Control Center by editing the security configuration 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 62

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 `role-mapping.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_config.xml` 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:

- Preconfigured user login – defines a user name, password, and a list of roles. The default user name is `sccadmin`; its password is configured during installation and its native role is SCC Administrator, which maps to `sccAdminRole`. You can create additional accounts by adding preconfigured user login modules to `csi_config.xml`. However, Sybase does not recommend the use of preconfigured user login modules for authentication in production environments.
- 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 after an upgrade from SCC 3.2.5 or earlier.
- 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. Configuration is required only if you have upgraded from an older version of Sybase Control Center and no longer want to use the older version’s authentication settings; if you do not want to use Windows for authentication; or if you want to create login accounts manually. Sybase recommends that you allow SCC to create accounts automatically.

This task is optional. However, if you choose not to create Sybase Control Center accounts automatically, you must enter them manually. Even when SCC users authenticate through LDAP or the local operating system, SCC 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. Click to select or deselect the box labeled **Automatically add SCC login records for authenticated logins**.
4. Click to select or deselect 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.
- Grant privileges to login accounts that require more than basic user access. You can grant privileges by assigning Sybase Control Center roles directly to the login accounts, or by assigning the login accounts to groups and mapping Sybase Control Center roles to the groups. The group approach is generally more efficient.

Configuring a Pluggable Authentication Module (PAM) for UNIX

Set up Sybase Control Center to support username and password login using accounts on the UNIX operating system.

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

Platform	Action
Solaris	Append the contents of the <code><SCC-install-dir>/utility/<sunos>/pam.conf</code> file (provided with Sybase Control Center) to the <code>/etc/pam.conf</code> file on your Solaris platform.
Linux	Copy the <code><SCC-install-dir>/utility/<linux>/sybase-csi</code> file (provided with Sybase Control Center) to the <code>/etc/pam.d</code> directory on your Linux platform. Note: The <code>sybase-csi</code> 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) Change account creation options.
 - a) Log in to Sybase Control Center using an account with administrative privileges (`sccAdminRole`).
 - b) Select **Application > Administration > Security**.
 - c) Click to select or deselect the box labeled **Automatically add SCC login records for authenticated logins**. (By default, this option is enabled for SCC 3.2.6 and later.)
 - d) Click to select or deselect the box labeled **Automatically grant sccUserRole to newly created logins**. (By default, this option is enabled for SCC 3.2.6 and later.)
 - e) 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-csi` file provided with Sybase Control Center. Instead, in your `/etc/pam.d` directory, create a `sybase-csi` file that contains:

```
# sybase-csi 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. Sybase Control Center needs the accounts for purposes of setting authorization (user privileges).
- Grant privileges to 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

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 106
- *Adding a Login Account to the System* on page 118

Configuring an LDAP Authentication Module

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

1. Open the `<SCC-install-dir>\conf\csi_config.xml` file.
2. Uncomment the LDAP module in the configuration file by removing the surrounding `<!--` and `-->` characters (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 of this task for values that work for ActiveDirectory. Configuration properties you can use in the LDAP module are described in a subtopic.

```
<authenticationProvider controlFlag="sufficient"
name="com.sybase.security.ldap.LDAPLoginModule">
  <options name="BindDN" value="cn=Directory Manager"/>
  <options name="BindPassword" value="secret"/>
  <options name="DefaultSearchBase" value="dc=example,dc=com"/>
  <options name="ProviderURL" value="ldap://localhost:10389"/>
  <options name="ServerType" value="openldap"/>
</authenticationProvider>
<provider name="com.sybase.security.ldap.LDAPAttributer"
type="attributer"/>
```

Note: Change only values shown in bold. If BindPassword is encrypted (which Sybase recommends), the line that defines it must include `encrypted="true"`. The line should look similar to this:

```
<options name="BindPassword" encrypted="true"
value="lsnjikfwregfqr43hu5io..." />
```

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:

Windows:

```
keytool -import -keystore %SYBASE_JRE7%\lib\security\cacerts -
file <your cert file and path>
-alias ldapcert -storepass changeit
```

UNIX:

```
keytool -import -keystore $SYBASE_JRE7/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

Map Sybase Control Center roles to LDAP groups.

See also

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 57

LDAP Configuration Properties

Use these properties in your `csi_config.xml` file to control the Sybase Control Center LDAP service.

Note: These characters have special meaning when they appear in a name in LDAP: , (comma), = (equals), + (plus), < (less than), > (greater than), # (number or hash sign), ; (semicolon), \ (backslash), / (forward slash), LF (line feed), CR (carriage return), " (double quotation mark), ' (single quotation mark), * (asterisk), ? (question mark), & (ampersand), and a space at the beginning or end of a string. LDAP providers do not handle these special characters in any of the names or DN's in any of the configuration properties. Additionally, some of the properties, as identified below, cannot use these special characters in common names.

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

Property	Default Value	Description
DefaultSearchBase	None	<p>The LDAP search base that is used if no other search base is specified for authentication, roles, attribution, and self registration:</p> <ol style="list-style-type: none"> 1. <code>dc=<domainname>,dc=<tld></code> For example, a machine in the sybase.com domain would have a search base of <code>dc=sybase,dc=com</code>. 2. <code>o=<company name>,c=<country code></code> For example, this might be <code>o=Sybase,c=us</code> for a machine within the Sybase organization. <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
SecurityProtocol	None	<p>The protocol to be used when connecting to the LDAP server.</p> <p>To use an encrypted protocol, use <code>ssl</code> instead of <code>ldaps</code> in the URL.</p>
AuthenticationMethod	Simple	<p>The authentication method to use for all authentication requests into LDAP. Legal values are generally the same as those of the <code>java.naming.security.authentication</code> JNDI property. Choose one of:</p> <ul style="list-style-type: none"> • <code>simple</code> — For clear-text password authentication. • <code>DIGEST-MD5</code> — 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	<p>For most LDAP servers: (& (uid={uid}) (objectclass=person))</p> <p>or</p> <p>For Active Directory e-mail lookups: (& (userPrincipalName={uid}) (objectclass=user)) [ActiveDirectory]</p> <p>For Active Directory Windows user name lookups: (& (sAMAccountName={uid}) (objectclass=user))</p>	<p>The filter to use when looking up the user.</p> <p>When performing a user name based lookup, this filter is used to determine the LDAP entry that matches the supplied user name.</p> <p>The string "{uid}" in the filter is replaced with the supplied user name.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
AuthenticationScope	onelevel	<p>The authentication search scope. The supported values for this are:</p> <ul style="list-style-type: none"> • onellevel • subtree <p>If you do not specify a value or if you specify an invalid value, the default value is used.</p>

Property	Default Value	Description
AuthenticationSearchBase	None	<p>The search base used to authenticate users. If this property is not configured, the value for DefaultSearchBase is used.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
BindDN	None	<p>The user DN to bind against when building the initial LDAP connection.</p> <p>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.</p> <p>However, the LDAP attributer may use this DN to create users in the LDAP server. When the self-registration feature is used, this user may need permissions to create a user record. This behavior may occur if you do not set useUserCredentialsToBind to true. In this case, the LDAP attributer uses this DN to update the user attributes.</p>

Property	Default Value	Description
BindPassword	None	<p>The password for BindDN, which is used to authenticate any user. BindDN and BindPassword separate the LDAP connection into units.</p> <p>The AuthenticationMethod property determines the bind method used for this initial connection.</p> <p>Sybase recommends that you encrypt passwords, and provides a password encryption utility. If you encrypt BindPassword, include <code>encrypted=true</code> in the line that sets the option. For example:</p> <pre data-bbox="713 579 1177 656"><options name="BindPassword" encrypted="true" value="1snjikf-wregfqr43hu5io..." /></pre> <p>If you do not encrypt BindPassword, the option might look like this:</p> <pre data-bbox="713 749 1177 800"><options name="BindPassword" value="s3cr3T" /></pre>
RoleSearchBase	None	<p>The search base used to retrieve lists of roles. If this property is not configured, LDAP uses the value for DefaultSearchBase.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul data-bbox="713 1034 1185 1190" style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property. <hr/>

Property	Default Value	Description
RoleFilter	<p>For SunONE/iPlanet: (<code>&</code>; (object-class=ldapsubentry) (object-class=nsroledefinition))</p> <p>For Netscape Directory Server: ((object-class=groupofnames) (object-class=groupofuniquenames))</p> <p>For ActiveDirectory: ((object-class=groupofnames) (object-class=group))</p>	<p>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 and this property is not initialized, no roles are available.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
RoleMemberAttributes	For Netscape Directory Server and OpenLDAP Server: member,unique-member	<p>A comma-separated list of role attributes from which LDAP derives the DNs of users who have this role.</p> <p>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 has a default value only when the Netscape server type is chosen.</p>
RoleNameAttribute	cn	The attribute of the role entry used as the role name. This is the role name displayed in the role list or granted to the authenticated user.
RoleScope	onelevel	<p>The role search scope. Supported values include:</p> <ul style="list-style-type: none"> • onelevel • subtree <p>If you do not specify a value or if you specify an invalid value, LDAP uses the default value.</p>

Property	Default Value	Description
SkipRoleLookup	false	<p>Set this property to true to grant the roles looked up using the attributes specified by the property UserRoleMembershipAttributes without cross-referencing them with the roles looked up using the RoleSearchBase and RoleFilter.</p> <p>LDAP configuration validation succeeds even when an error is encountered when listing all the available roles. The error is logged to the server log during validation but not reported in SCC, allowing the configuration to be saved. This has an impact when listing the physical roles for role mapping as well as in SCC. To successfully authenticate the user, set the SkipRoleLookup property to true.</p>
UserRoleMembershipAttributes	<p>For iPlanet/SunONE: nsRoleDN</p> <p>For Active Directory: memberOf</p> <p>For all others: none</p>	<p>Defines a user attribute that contains the DNs of all of the roles a user is a member of.</p> <p>These comma-delimited values are cross-referenced with the roles retrieved in the role search base and search filter to generate a list of user's roles.</p> <p>If the SkipRoleSearch property is set to true, these comma-delimited values are not cross-referenced with the roles retrieved in the role search base and role search filter. See <i>SkipRoleLookup</i>.</p> <hr/> <p>Note: If you use nested groups with Active Directory, you must set this property to tokenGroups.</p>
UserFreeformRoleMembershipAttributes	None	<p>The free-form 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 the department attribute in the user's LDAP record has the values {sales, consulting}, the user is granted the roles sales and consulting.</p>
Referral	ignore	<p>The behavior when a referral is encountered. Valid values are dictated by LdapContext, but might include follow, ignore, or throw.</p>

Get Started

Property	Default Value	Description
DigestMD5Authentication-Format	DN For OpenLDAP: User name	The DIGEST-MD5 bind authentication identity format.
UseUserAccountControlAttribute	For Active Directory: true	When this property is set to true, the UserAccountControl attribute detects disabled user accounts, account expirations, password expirations, and so on. Active Directory also uses this attribute to store the above information.
EnableLDAPConnectionTrace	False	Enables LDAP connection tracing. The output is logged to a file in the <code>temp</code> directory. The location of the file is logged to the server log.
ConnectTimeout	0	Specifies the timeout, in milliseconds, for attempts to connect to the LDAP server. The property value sets the JNDI <code>com.sun.jndi.ldap.connect.timeout</code> property when attempting to establish a connection to a configured LDAP server. If the LDAP provider cannot establish a connection within the configured interval, it aborts the connection attempt. An integer value less than or equal to zero results in the use of the network protocol's timeout value.
ReadTimeout	0	Controls the length of time, in milliseconds, the client waits for the server to respond to a read attempt after the initial connection to the server has been established. The property values sets the JNDI <code>com.sun.jndi.ldap.read.timeout</code> property when attempting to establish a connection to a configured LDAP server. If the LDAP provider does not receive an LDAP response within the configured interval, it aborts the read attempt. The read timeout applies to the LDAP response from the server after the initial connection is established with the server. An integer value less than or equal to zero indicates no read timeout is specified.

Property	Default Value	Description
LDAPPoolMaxActive	8	Caps the number of concurrent LDAP connections to the LDAP server. A non-positive value indicates no limit. If this option is set for multiple LDAP providers, the value set by the first LDAP provider loaded takes precedence over all the others. When LDAPPoolMaxActive is reached, any further attempts by the LDAP provider classes to borrow LDAP connections from the pool are blocked indefinitely until a new or idle object becomes available in the pool. Connection pooling improves the LDAP provider's performance and resource utilization by managing the number of TCP connections established with configured LDAP servers.
controlFlag	optional	When you configure multiple authentication providers, use controlFlag for each provider to control how the authentication providers are used in the login sequence. controlFlag is a generic login module option rather than an LDAP configuration property.

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.

Prerequisites

- Required: Configure an LDAP authentication module.
- Optional: Create these LDAP groups and assign Sybase Control Center users to them:
 - sybase – confers sccUserRole. Assign all SCC users to the sybase group.
 - SCC Administrator – confers sccAdminRole. Assign only SCC administrators to this group.

Task

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).

The sybase and SCC Administrator groups are convenient because they are predefined in `role-mapping.xml`. If you add sybase and SCC Administrator groups to your LDAP

Get Started

system and populate them with SCC users and administrators, you can skip to the next task—you do not need to complete the steps below.

The table lists default mappings of LDAP and OS groups to SCC roles. Login modules are defined in `csi_config.xml`.

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

There are two ways to accomplish the mapping:

- (Recommended) Add a “sybase” group and an “SCC Administrator” 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 one or both groups.
- Configure Sybase Control Center to use existing groups in LDAP or the operating system by editing the `role-mapping.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/role-mapping.xml
```

3. Locate the `sccUserRole` section of the file:

```
<Mapping>
  <LogicalName>sccUserRole</LogicalName>
  <MappedName>SCC Administrator</MappedName>
  <MappedName>SCC Agent Administrator</MappedName>
  <MappedName>sybase</MappedName>
</Mapping>
```

4. Add a `MappedName` line for the LDAP or OS group you are using to authenticate SCC users. The `sccUserRole` section should look similar to this:

```
<Mapping>
  <LogicalName>sccUserRole</LogicalName>
```

```

    <MappedName>SCC Administrator</MappedName>
    <MappedName>SCC Agent Administrator</MappedName>
    <MappedName>sybase</MappedName>
    <MappedName>my_SCC_group</MappedName>
  </Mapping>

```

5. Locate the sccAdminRole section of the file:

```

<Mapping>
  <LogicalName>sccAdminRole</LogicalName>
  <MappedName>SCC Administrator</MappedName>
</Mapping>

```

6. Add a MappedName line for the LDAP or OS group you are using to authenticate SCC administrators. The sccAdminRole section should look similar to this:

```

<Mapping>
  <LogicalName>sccAdminRole</LogicalName>
  <MappedName>SCC Administrator</MappedName>
  <MappedName>my_SCC_admin_group</MappedName>
</Mapping>

```

7. Save the file and exit.
8. (LDAP only) Ensure that the roles defined in the LDAP repository match the roles defined in role-mapping.xml.
9. In the <SCC-install-dir>\conf\csi_config.xml 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_config.xml.
10. Start Sybase Control Center.

See also

- *Configuring an LDAP Authentication Module* on page 47

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 configuration 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 SHA-256 hash algorithm for passwords used in the PreConfiguredLoginModule in csi_config.xml.

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 -csi**. Enter your new password at the resulting prompt.

Get Started

passencrypt encrypts the password you enter (which does not appear on the screen) and displays the password in encrypted form.

3. Copy the encrypted password.
4. Paste the encrypted password where needed.

Configuring Ports

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

Prerequisites

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

Task

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

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

Port Name	Description	Service Names	Property Names	Default Port
jiniRmid	JINI remote method invocation daemon Present on SCC server and SCC agent	Jini	rmiPort	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

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.

Get Started

```
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 application's 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 42

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 application's 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.

Get Started

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 Sybase Control Center* on page 123

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 65
- *Adding a Login Account to a Group* on page 65
- *Logins, Roles, and Groups* on page 66

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 64
- *Adding a Login Account to a Group* on page 65
- *Logins, Roles, and Groups* on page 66

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 64
- *Adding a Group* on page 65
- *Logins, Roles, and Groups* on page 66

Logins, Roles, and Groups

Sybase Control Center includes predefined login accounts and roles.

A login account identifies a user who can connect to Sybase Control Center. An account has roles that control the tasks the user is allowed to perform. Users can be authenticated through native SCC accounts, but a safer approach is to delegate authentication to the operating system or to an LDAP directory service.

Sybase Control Center comes with a predefined login account. Sybase recommends using the predefined account only for installing and setting up Sybase Control Center. This account is not intended for use in a production environment.

Table 9. Predefined Login Account

Login Name	Description
sccadmin	Can use all the administration features in Sybase Control Center. Use for configuration and test.

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 10. Predefined Roles

Role	Description
sccUserRole	Provides nonadministrative access to Sybase Control Center. Required for all users and assigned automatically to every authenticated 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

- *Assigning a Role to a Login or a Group* on page 64
- *Adding a Group* on page 65
- *Adding a Login Account to a Group* on page 65

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
```

Get Started

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 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78

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 reenables, 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 69
- *sccinstance Command* on page 70

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, SAP Sybase ESP clusters, or SAP 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 68
- *sccinstance Command* on page 70

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	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

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
```

Get Started

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 68
- *Shared-Disk Mode* on page 69

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 251

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 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78
- *Configuring Memory Usage* on page 82
- *scc Command* on page 86

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, which includes the management UI) and the Sybase Control Center agent that runs on each product server managed by SCC. When you install SCC and the SCC agent in the same directory by selecting both options in the installer, you start and stop them together—by executing a single command or controlling a single service. This topic applies both to singleton installations (which do not use a shared disk) and to instances of SCC agents and servers running from a shared disk.

Get Started

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, the SCC agent, or both when they are installed together:**
 - 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, the SCC agent, or both when they are installed together:**

- 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:**

- Open the Windows Control Panel.
- Select **Administrative Tools > Services**.
- 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.”
- To start or stop the service, right-click the **Sybase Control Center** entry in the Services list and choose **Start** or **Stop**.
- To configure automatic starting, double-click the service.
- To set the service to automatically start when the machine starts, change the **Startup type** to Automatic.
- To restart the service in case of failure, choose the **Recovery** tab and change the First, Second, and Subsequent failures to Restart Service.
- Click **Apply** to save the modifications and close the dialog.

- **Start or stop the Sybase Control Center service (controlling Sybase Control Center, the SCC agent, or both) from the Windows command line:**

- To start the service, enter the **net start** command.

For a singleton installation:

```
net start "sybase control center 3.2.8"
```

```
The Sybase Control Center 3.2.8 service is starting.....
The Sybase Control Center 3.2.8 service was started
successfully.
```

For an instance, include the instance name (Boston-1 in this example) in parentheses:

```
net start "sybase control center 3.2.8 (Boston-1) "
```

```
The Sybase Control Center 3.2.8 (Boston-1) service is
```

```
starting.....  
The Sybase Control Center 3.2.8 (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 3.2.8"  
  
The Sybase Control Center 3.2.8 service is stopping.....  
The Sybase Control Center 3.2.8 service was stopped  
successfully.
```

For an instance, include the instance name (Boston-1 in this example) in parentheses:

```
net stop "sybase control center 3.2.8 (Boston-1)"  
  
The Sybase Control Center 3.2.8 (Boston-1) service is  
stopping.....  
The Sybase Control Center 3.2.8 (Boston-1) service was  
stopped successfully.
```

See also

- *Registering the ODBC Driver in Windows* on page 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78
- *Configuring Memory Usage* on page 82
- *scc Command* on page 86

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, which includes the management UI) and the Sybase Control Center agent that runs on each product server managed by SCC.. When you install SCC and the SCC agent in the same directory by selecting both options in the installer, you start and stop them together—by executing a single command or controlling a single service. This topic applies to both singleton installations (which do not use a shared disk) 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 (or both, when they are installed together) 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 (or both, when they are installed together) 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.

Get Started

- 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 (or both, when they are installed together) .**

- 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 81 or step *5.d* on page 82.
- 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 output 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**.)

Get Started

- 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 75
- *Starting and Stopping Sybase Control Center in Windows* on page 75
- *Configuring Memory Usage* on page 82
- *scc Command* on page 86

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
<p>Maximum memory</p> <ul style="list-style-type: none"> • <code>jvmopt=-Xmx</code> – if you are running Sybase Control Center as a Windows service • <code>SCC_MEM_MAX</code> – if you are running SCC as a UNIX service • <code>SCC_MEM_MAX</code> – if you are starting SCC from the command line 	<ul style="list-style-type: none"> • You need to prevent Sybase Control Center from using more than a given amount of memory • Sybase Control Center fails to start and may display an error: Could not create the Java Virtual machine. • An OutOfMemory error says Sybase Control Center is out of heap space • A warning message about system memory appears during the start process • The machine where Sybase Control Center is installed has less than 4GB of memory. (Starting Sybase Control Center on a machine with less than 4GB of memory triggers the startup warning message about system memory.) 	<p>On machines with less than 4GB of memory, set maximum memory to 256MB or more.</p> <p>Default value: none. (On machines with 4GB or more of memory, maximum memory is set dynamically and is effectively limited only by the amount of system memory available.)</p>
<p>Permanent memory</p> <ul style="list-style-type: none"> • <code>jvmopt=-XX:MaxPermSize</code> – if you are running Sybase Control Center as a Windows service • <code>SCC_MEM_PERM</code> – if you are running SCC as a UNIX service • <code>SCC_MEM_PERM</code> – if you are starting SCC from the command line 	<p>An OutOfMemory error says Sybase Control Center is out of permanent generation space</p>	<p>Increase by 32MB increments. If you reach a value equal to twice the default and still see the OutOfMemory error, contact Sybase technical support.</p> <p>Default value: 128MB</p>

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

See also

- *Registering the ODBC Driver in Windows* on page 75
- *Starting and Stopping Sybase Control Center in Windows* on page 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78
- *scc Command* on page 86

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 a Sybase Control Center Windows Service* on page 84
- *Changing a Memory Option for an SCC UNIX Service* on page 85
- *Starting and Stopping Sybase Control Center in Windows* on page 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78
- *scc Command* on page 86

Changing a Memory Option for a Sybase Control Center 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 Sybase Control Center properties file:

```
<Sybase Control Center-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 84
- *Changing a Memory Option for an SCC UNIX Service* on page 85
- *Starting and Stopping Sybase Control Center in Windows* on page 75

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 84
- *Changing a Memory Option for a Sybase Control Center Windows Service* on page 84
- *Starting and Stopping Sybase Control Center in UNIX* on page 78

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	ScsSADataserver 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
tds	Tabular Data Stream™ port (used to communi- cate 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 75
- *Starting and Stopping Sybase Control Center in Windows* on page 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78
- *Configuring Memory Usage* on page 82
- *Configuring Ports* on page 109
- *Logging or Message Levels* on page 249

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.

Only one login session per account is permitted at a time; multiple users cannot be logged in to the same account simultaneously.

Note: When logging in to a newly installed Sybase Control Center for which secure authentication has not been configured, use the sccadmin account—the password is set during installation. 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@sap`, not `fred@sap.com`.

See also

- *Logging out of Sybase Control Center* on page 91

Logging out of Sybase Control Center

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

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

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

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

If no automatic logout period is configured,

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

See also

- *Logging in to Sybase Control Center* on page 90

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 *Quick Start for an Evaluation* on page 13.

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. Configuration is required only if you have upgraded from an older version of Sybase Control Center and no longer want to use the older version's authentication settings; if you do not want to use Windows for authentication; or if you want to create login accounts manually. Sybase recommends that you allow SCC to create accounts automatically.

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

Get Started

Set up Sybase Control Center to support username and password login using accounts on the UNIX operating system.

4. *Configuring an LDAP Authentication Module*

Configure an LDAP authentication module for Sybase Control Center by editing the security configuration 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 `role-mapping.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_config.xml` 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:

- Preconfigured user login – defines a user name, password, and a list of roles. The default user name is `sccadmin`; its password is configured during installation and its native role is SCC Administrator, which maps to `sccAdminRole`. You can create additional accounts by adding preconfigured user login modules to `csi_config.xml`. However, Sybase does not recommend the use of preconfigured user login modules for authentication in production environments.
- 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 after an upgrade from SCC 3.2.5 or earlier.
- 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 45
- *Configuring an LDAP Authentication Module* on page 95
- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 106

Configuring Authentication for Windows

Authentication through the Windows operating system is enabled by default. Configuration is required only if you have upgraded from an older version of Sybase Control Center and no longer want to use the older version's authentication settings; if you do not want to use Windows for authentication; or if you want to create login accounts manually. Sybase recommends that you allow SCC to create accounts automatically.

This task is optional. However, if you choose not to create Sybase Control Center accounts automatically, you must enter them manually. Even when SCC users authenticate through LDAP or the local operating system, SCC 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. Click to select or deselect the box labeled **Automatically add SCC login records for authenticated logins**.
4. Click to select or deselect 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.
- Grant privileges to 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 95
- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 106
- *Adding a Login Account to the System* on page 118

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.

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/<sunos>/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-csi file (provided with Sybase Control Center) to the /etc/pam.d directory on your Linux platform. <hr/> Note: The sybase-csi 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) Change account creation options.
 - a) Log in to Sybase Control Center using an account with administrative privileges (`sccAdminRole`).
 - b) Select **Application > Administration > Security**.
 - c) Click to select or deselect the box labeled **Automatically add SCC login records for authenticated logins**. (By default, this option is enabled for SCC 3.2.6 and later.)
 - d) Click to select or deselect the box labeled **Automatically grant sccUserRole to newly created logins**. (By default, this option is enabled for SCC 3.2.6 and later.)
 - e) 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-csi` file provided with Sybase Control Center. Instead, in your `/etc/pam.d` directory, create a `sybase-csi` file that contains:

```
# sybase-csi 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. Sybase Control Center needs the accounts for purposes of setting authorization (user privileges).
- Grant privileges to login accounts that require more than basic user access. You can grant privileges by assigning Sybase Control Center roles directly to the login accounts, or by assigning the login accounts to groups and mapping Sybase Control Center roles to the groups. The group approach is generally more efficient.

Configuring an LDAP Authentication Module

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

1. Open the `<SCC-install-dir>\conf\csi_config.xml` file.

Get Started

2. Uncomment the LDAP module in the configuration file by removing the surrounding `<!--` and `-->` characters (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 of this task for values that work for ActiveDirectory. Configuration properties you can use in the LDAP module are described in a subtopic.

```
<authenticationProvider controlFlag="sufficient"
name="com.sybase.security.ldap.LDAPLoginModule">
  <options name="BindDN" value="cn=Directory Manager"/>
  <options name="BindPassword" value="secret"/>
  <options name="DefaultSearchBase" value="dc=example,dc=com"/>
  <options name="ProviderURL" value="ldap://localhost:10389"/>
  <options name="ServerType" value="openldap"/>
</authenticationProvider>
<provider name="com.sybase.security.ldap.LDAPAttributer"
type="attributer"/>
```

Note: Change only values shown in bold. If BindPassword is encrypted (which Sybase recommends), the line that defines it must include `encrypted="true"`. The line should look similar to this:

```
<options name="BindPassword" encrypted="true"
value="lsnjikfwregfqr43hu5io..."/>
```

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:

Windows:

```
keytool -import -keystore %SYBASE_JRE7%\lib\security\cacerts -
file <your cert file and path>
-alias ldapcert -storepass changeit
```

UNIX:

```
keytool -import -keystore $SYBASE_JRE7/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

Map Sybase Control Center roles to LDAP groups.

See also

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 106

LDAP Configuration Properties

Use these properties in your `csi_config.xml` file to control the Sybase Control Center LDAP service.

Note: These characters have special meaning when they appear in a name in LDAP: , (comma), = (equals), + (plus), < (less than), > (greater than), # (number or hash sign), ; (semicolon), \ (backslash), / (forward slash), LF (line feed), CR (carriage return), " (double quotation mark), ' (single quotation mark), * (asterisk), ? (question mark), & (ampersand), and a space at the beginning or end of a string. LDAP providers do not handle these special characters in any of the names or DNs in any of the configuration properties. Additionally, some of the properties, as identified below, cannot use these special characters in common names.

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

Property	Default Value	Description
DefaultSearchBase	None	<p>The LDAP search base that is used if no other search base is specified for authentication, roles, attribution, and self registration:</p> <ol style="list-style-type: none"> 1. <code>dc=<domainname>,dc=<tld></code> For example, a machine in the sybase.com domain would have a search base of <code>dc=sybase,dc=com</code>. 2. <code>o=<company name>,c=<country code></code> For example, this might be <code>o=Sybase,c=us</code> for a machine within the Sybase organization. <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
SecurityProtocol	None	<p>The protocol to be used when connecting to the LDAP server.</p> <p>To use an encrypted protocol, use <code>ssl</code> instead of <code>ldaps</code> in the URL.</p>
AuthenticationMethod	Simple	<p>The authentication method to use for all authentication requests into LDAP. Legal values are generally the same as those of the <code>java.naming.security.authentication</code> JNDI property. Choose one of:</p> <ul style="list-style-type: none"> • <code>simple</code> — For clear-text password authentication. • <code>DIGEST-MD5</code> — 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
<p>AuthenticationFilter</p>	<p>For most LDAP servers: (&(uid={uid})(objectclass=person)) or For Active Directory e-mail lookups: (&(userPrincipalName={uid})(objectclass=user)) [ActiveDirectory] For Active Directory Windows user name lookups: (&(sAMAccountName={uid})(objectclass=user))</p>	<p>The filter to use when looking up the user.</p> <p>When performing a user name based lookup, this filter is used to determine the LDAP entry that matches the supplied user name.</p> <p>The string "{uid}" in the filter is replaced with the supplied user name.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
<p>AuthenticationScope</p>	<p>onelevel</p>	<p>The authentication search scope. The supported values for this are:</p> <ul style="list-style-type: none"> • onellevel • subtree <p>If you do not specify a value or if you specify an invalid value, the default value is used.</p>

Property	Default Value	Description
AuthenticationSearchBase	None	<p>The search base used to authenticate users. If this property is not configured, the value for DefaultSearchBase is used.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property.
BindDN	None	<p>The user DN to bind against when building the initial LDAP connection.</p> <p>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.</p> <p>However, the LDAP attributer may use this DN to create users in the LDAP server. When the self-registration feature is used, this user may need permissions to create a user record. This behavior may occur if you do not set useUserCredentialsToBind to true. In this case, the LDAP attributer uses this DN to update the user attributes.</p>

Property	Default Value	Description
BindPassword	None	<p>The password for BindDN, which is used to authenticate any user. BindDN and BindPassword separate the LDAP connection into units.</p> <p>The AuthenticationMethod property determines the bind method used for this initial connection.</p> <p>Sybase recommends that you encrypt passwords, and provides a password encryption utility. If you encrypt BindPassword, include encrypted=true in the line that sets the option. For example:</p> <pre data-bbox="713 579 1180 656"><options name="BindPassword" encrypted="true" value="1snjikf-wregfqr43hu5io..." /></pre> <p>If you do not encrypt BindPassword, the option might look like this:</p> <pre data-bbox="713 749 1180 800"><options name="BindPassword" value="s3cr3T" /></pre>
RoleSearchBase	None	<p>The search base used to retrieve lists of roles. If this property is not configured, LDAP uses the value for DefaultSearchBase.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul data-bbox="713 1034 1185 1190" style="list-style-type: none"> • Do not use special characters, as listed above, in common names or distinguished names in the value of this property. • Do not use Chinese or Japanese characters in user names or passwords of this property. <hr/>

Property	Default Value	Description
RoleFilter	<p>For SunONE/iPlanet: (<code>&</code>; (object-class=ldapsubentry) (object-class=nsroledefinition))</p> <p>For Netscape Directory Server: (<code> </code> (object-class=groupofnames) (object-class=groupofuniquenames))</p> <p>For ActiveDirectory: (<code> </code> (object-class=groupofnames) (object-class=group))</p>	<p>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 and this property is not initialized, no roles are available.</p> <hr/> <p>Note: When you use this property to authenticate SCC:</p> <ul style="list-style-type: none"> Do not use special characters, as listed above, in common names or distinguished names in the value of this property. Do not use Chinese or Japanese characters in user names or passwords of this property.
RoleMemberAttributes	For Netscape Directory Server and OpenLDAP Server: member,unique-member	<p>A comma-separated list of role attributes from which LDAP derives the DNs of users who have this role.</p> <p>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 has a default value only when the Netscape server type is chosen.</p>
RoleNameAttribute	cn	The attribute of the role entry used as the role name. This is the role name displayed in the role list or granted to the authenticated user.
RoleScope	onelevel	<p>The role search scope. Supported values include:</p> <ul style="list-style-type: none"> onelevel subtree <p>If you do not specify a value or if you specify an invalid value, LDAP uses the default value.</p>

Property	Default Value	Description
SkipRoleLookup	false	<p>Set this property to true to grant the roles looked up using the attributes specified by the property UserRoleMembershipAttributes without cross-referencing them with the roles looked up using the RoleSearchBase and RoleFilter.</p> <p>LDAP configuration validation succeeds even when an error is encountered when listing all the available roles. The error is logged to the server log during validation but not reported in SCC, allowing the configuration to be saved. This has an impact when listing the physical roles for role mapping as well as in SCC. To successfully authenticate the user, set the SkipRoleLookup property to true.</p>
UserRoleMembershipAttributes	<p>For iPlanet/SunONE: nsRoleDN</p> <p>For Active Directory: memberOf</p> <p>For all others: none</p>	<p>Defines a user attribute that contains the DNs of all of the roles a user is a member of.</p> <p>These comma-delimited values are cross-referenced with the roles retrieved in the role search base and search filter to generate a list of user's roles.</p> <p>If the SkipRoleSearch property is set to true, these comma-delimited values are not cross-referenced with the roles retrieved in the role search base and role search filter. See <i>SkipRoleLookup</i>.</p> <hr/> <p>Note: If you use nested groups with Active Directory, you must set this property to tokenGroups.</p>
UserFreeformRoleMembershipAttributes	None	<p>The free-form 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 the department attribute in the user's LDAP record has the values {sales, consulting}, the user is granted the roles sales and consulting.</p>
Referral	ignore	<p>The behavior when a referral is encountered. Valid values are dictated by LdapContext, but might include follow, ignore, or throw.</p>

Property	Default Value	Description
DigestMD5Authentication-Format	DN For OpenLDAP: User name	The DIGEST-MD5 bind authentication identity format.
UseUserAccountControlAttribute	For Active Directory: true	When this property is set to true, the UserAccountControl attribute detects disabled user accounts, account expirations, password expirations, and so on. Active Directory also uses this attribute to store the above information.
EnableLDAPConnection-Trace	False	Enables LDAP connection tracing. The output is logged to a file in the <code>temp</code> directory. The location of the file is logged to the server log.
ConnectTimeout	0	Specifies the timeout, in milliseconds, for attempts to connect to the LDAP server. The property value sets the JNDI <code>com.sun.jndi.ldap.connect.timeout</code> property when attempting to establish a connection to a configured LDAP server. If the LDAP provider cannot establish a connection within the configured interval, it aborts the connection attempt. An integer value less than or equal to zero results in the use of the network protocol's timeout value.
ReadTimeout	0	Controls the length of time, in milliseconds, the client waits for the server to respond to a read attempt after the initial connection to the server has been established. The property values sets the JNDI <code>com.sun.jndi.ldap.read.timeout</code> property when attempting to establish a connection to a configured LDAP server. If the LDAP provider does not receive an LDAP response within the configured interval, it aborts the read attempt. The read timeout applies to the LDAP response from the server after the initial connection is established with the server. An integer value less than or equal to zero indicates no read timeout is specified.

Property	Default Value	Description
LDAPPoolMaxActive	8	Caps the number of concurrent LDAP connections to the LDAP server. A non-positive value indicates no limit. If this option is set for multiple LDAP providers, the value set by the first LDAP provider loaded takes precedence over all the others. When LDAPPoolMaxActive is reached, any further attempts by the LDAP provider classes to borrow LDAP connections from the pool are blocked indefinitely until a new or idle object becomes available in the pool. Connection pooling improves the LDAP provider's performance and resource utilization by managing the number of TCP connections established with configured LDAP servers.
controlFlag	optional	When you configure multiple authentication providers, use controlFlag for each provider to control how the authentication providers are used in the login sequence. controlFlag is a generic login module option rather than an LDAP configuration property.

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.

Prerequisites

- Required: Configure an LDAP authentication module.
- Optional: Create these LDAP groups and assign Sybase Control Center users to them:
 - sybase – confers sccUserRole. Assign all SCC users to the sybase group.
 - SCC Administrator – confers sccAdminRole. Assign only SCC administrators to this group.

Task

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).

The sybase and SCC Administrator groups are convenient because they are predefined in `role-mapping.xml`. If you add sybase and SCC Administrator groups to your LDAP

system and populate them with SCC users and administrators, you can skip to the next task—you do not need to complete the steps below.

The table lists default mappings of LDAP and OS groups to SCC roles. Login modules are defined in `csi_config.xml`.

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

There are two ways to accomplish the mapping:

- (Recommended) Add a “sybase” group and an “SCC Administrator” 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 one or both groups.
- Configure Sybase Control Center to use existing groups in LDAP or the operating system by editing the `role-mapping.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/role-mapping.xml
```

3. Locate the `sccUserRole` section of the file:

```
<Mapping>
  <LogicalName>sccUserRole</LogicalName>
  <MappedName>SCC Administrator</MappedName>
  <MappedName>SCC Agent Administrator</MappedName>
  <MappedName>sybase</MappedName>
</Mapping>
```

4. Add a `MappedName` line for the LDAP or OS group you are using to authenticate SCC users. The `sccUserRole` section should look similar to this:

```
<Mapping>
  <LogicalName>sccUserRole</LogicalName>
```

```
<MappedName>SCC Administrator</MappedName>
<MappedName>SCC Agent Administrator</MappedName>
<MappedName>sybase</MappedName>
<MappedName>my_SCC_group</MappedName>
</Mapping>
```

5. Locate the sccAdminRole section of the file:

```
<Mapping>
  <LogicalName>sccAdminRole</LogicalName>
  <MappedName>SCC Administrator</MappedName>
</Mapping>
```

6. Add a MappedName line for the LDAP or OS group you are using to authenticate SCC administrators. The sccAdminRole section should look similar to this:

```
<Mapping>
  <LogicalName>sccAdminRole</LogicalName>
  <MappedName>SCC Administrator</MappedName>
  <MappedName>my_SCC_admin_group</MappedName>
</Mapping>
```

7. Save the file and exit.
8. (LDAP only) Ensure that the roles defined in the LDAP repository match the roles defined in role-mapping.xml.
9. In the <SCC-install-dir>\conf\csi_config.xml 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_config.xml.
10. Start Sybase Control Center.

See also

- *Configuring an LDAP Authentication Module* on page 95
- *Configuring Authentication for Windows* on page 93
- *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 45
- *Assigning a Role to a Login or a Group* on page 115
- *User Authorization* on page 114
- *Configuring Event Stream Processor for Monitoring* on page 126
- *Configuring Event Stream Processor for Administration* on page 127

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 configuration 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 SHA-256 hash algorithm for passwords used in the PreConfiguredLoginModule in csi_config.xml.

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 -csi**. Enter your new password at the resulting prompt.
passencrypt encrypts the password you enter (which does not appear on the screen) and displays the password in encrypted form.
3. Copy the encrypted password.
4. Paste the encrypted password where needed.

Configuring Ports

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

Prerequisites

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

Task

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

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

Port Name	Description	Service Names	Property Names	Default Port
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

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

```
scc --port port-name=port-number
```

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

Use the first, simpler format unless you want to configure the database services to use different ports. (By default, they all use the same port.)

4. Start Sybase Control Center.
5. Execute **scc --info ports** again to confirm that the port has been reassigned.

Examples

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

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


Set only the database messaging service to port 3639.

```
scc --port Messaging:messaging.db.port=3639
```

Set the HTTP port to 9292.

```
scc --port http=9292
```

Set the Jini RMI daemon to port 9696.

```
scc --port jiniRmid=9696
```

Set the main Sybase Control Center messaging service to port 2001.

```
scc --port msg=2001
```

Set the RMI port to 9991.

```
scc --port rmi=9991
```

Set the Tabular Data Stream port to 9997.

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

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

See also

- *scc Command* on page 86

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 application's menu bar, select **Application > Administration**.
2. Select **General Settings**.
3. Click the **E-mail** tab.
4. Enter the name of the e-mail server through which Sybase Control Center will send alert notifications.
5. Change the default e-mail server port only in consultation with your e-mail administrator.

6. (Optional) Click **Customize e-mail settings** to display options for setting the domain name and e-mail sender for alert e-mail notifications.
7. (Optional) Enter your domain name (for example, mycompany.com).
Most e-mail servers do not require SCC to provide an explicit domain name. Try providing a domain name here if your first attempt to configure e-mail alerts fails.
8. (Optional) Change the default e-mail sender name.
This name appears in the "From" field of SCC e-mail alert messages. Do not use spaces; use hyphens or underscore characters instead.

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

Next

(Optional) Configure automatic logout.

See also

- *Launching Sybase Control Center* on page 74
- *Logging in to Sybase Control Center* on page 90

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 application's 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 74
- *Logging in to Sybase Control Center* on page 90

Configuring Retrieval Thresholds for the Administration Console

(Optional) Set limits on the time the Administration Console waits for data to load or on the number of rows it loads.

Prerequisites

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

Task

Performing some tasks may cause the Administration Console to load a large amount of data, which can be time-consuming and can place a heavy load on your network. This is particularly likely if your perspective includes many resources. The Administration Console mitigates this problem by displaying partial results and by displaying placeholders called message rows when data takes longer than a specified number of seconds to retrieve or exceeds a specified number of rows. The data retrieval options let you specify those numbers.

This data retrieval scheme reduces network traffic because result sets that exceed the specified row count are not transmitted unless you ask for them by expanding a message row. By displaying partial results and message rows for data from slow-responding resources, the scheme also minimizes the time you spend waiting.

1. From the application's menu bar, select **Application > Administration**.
2. Select **General Settings**.
3. Click the **Administration Console** tab.
4. Set the timeout for data retrieval in seconds.

When SCC is not able to return all requested data within this period of time, it displays any data it has received and generates message rows in place of the missing results. The Administration Console replaces message rows with real data as soon as the data arrives.

5. Set the row count.

When a request returns results that exceed the specified row count, SCC displays a message row in place of the expected results. You can expand the message row by selecting it, clicking the drop-down arrow, and selecting **Expand**.

6. Click **OK** (to apply the change and close the properties dialog) or **Apply** (to apply the change and leave the dialog open).

See also

- *Searching and Filtering Resources* on page 203

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 140
- *Configuring Event Stream Processor for Monitoring* on page 126
- *Configuring Event Stream Processor for Administration* on page 127

Assigning a Role to a Login or a Group

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

Prerequisites

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

Task

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

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

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

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

7. Click **OK**.

See also

- *Removing a Role from a Login or a Group* on page 115

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 account 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 115

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 116
- *Adding a Login Account to a Group* on page 117
- *Removing a Login Account from a Group* on page 117

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 116
- *Adding a Login Account to a Group* on page 117
- *Removing a Login Account from a Group* on page 117

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 116
- *Removing a Group* on page 116
- *Removing a Login Account from a Group* on page 117

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 116
- *Removing a Group* on page 116
- *Adding a Login Account to a Group* on page 117

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 manually 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. When you do that, SCC automatically creates a native account for every authenticated user.

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**.

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 93
- *Configuring a Pluggable Authentication Module (PAM) for UNIX* on page 45
- *Configuring an LDAP Authentication Module* on page 95
- *Removing a Login Account from the System* on page 119
- *Modifying a User Profile* on page 120

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

- *Adding a Login Account to the System* on page 118
- *Modifying a User Profile* on page 120

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 date.
8. Click **Apply**.
9. Click the **User Info** tab.
10. Edit the user information.
When this user configures e-mail alert subscriptions, Sybase Control Center automatically populates the subscription dialog with the e-mail address you enter here.
11. Click **Apply**.

See also

- *Adding a Login Account to the System* on page 118
- *Removing a Login Account from the System* on page 119

Logins, Roles, and Groups

Sybase Control Center includes predefined login accounts and roles.

A login account identifies a user who can connect to Sybase Control Center. An account has roles that control the tasks the user is allowed to perform. Users can be authenticated through

native SCC accounts, but a safer approach is to delegate authentication to the operating system or to an LDAP directory service.

Sybase Control Center comes with a predefined login account. Sybase recommends using the predefined account only for installing and setting up Sybase Control Center. This account is not intended for use in a production environment.

Table 11. Predefined Login Account

Login Name	Description
sccadmin	Can use all the administration features in Sybase Control Center. Use for configuration and test.

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 12. Predefined Roles

Role	Description
sccUserRole	Provides nonadministrative access to Sybase Control Center. Required for all users and assigned automatically to every authenticated user.
sccAdminRole	Provides administrative privileges for managing Sybase Control Center.

Monitoring privileges for SCC product modules are assigned automatically.

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

Get Started

Configure Sybase Control Center

Set up Sybase Control Center for Sybase Event Stream Processor.

Note: Before configuring Sybase Control Center for use in a production environment, complete the tasks in the *Get Started* section of the help. Setting up security is particularly important.

1. *Configuring Policies for Monitoring and Administering Event Stream Processor*

Edit the Project, Node, and Cluster policies in the policy.xml file to grant SCC monitoring and administrative access to a native OS, preconfigured login, LDAP, or SAP Business Intelligence (BI) group. Do not edit the policy.xml file if you are using Kerberos or RSA authentication for ESP as you are automatically granted SCC monitoring and administrative access.

2. *Configuring Event Stream Processor for Monitoring*

To enable users to monitor ESP node and cluster activity using Sybase Control Center, map a native OS, preconfigured login, LDAP, or SAP BI group to the espMonitorRole role.

3. *Configuring Event Stream Processor for Administration*

To enable users to start and stop ESP nodes, projects, and adapters from the Sybase Control Center Administration Console, and to view node and project log files, map a native OS, preconfigured login, LDAP, or SAP BI group to the espAdminRole role.

4. *Registering an ESP Node*

Make Sybase Control Center aware of an ESP node (acting as a cluster manager, controller, or both) and its connection information by registering it as a resource.

5. *Importing Resources for Batch Registration*

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

6. *Authenticating an ESP Node*

If you did not authenticate the ESP node while registering it, you can authenticate afterwards from the Perspective Resources window.

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

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

8. *Creating a Perspective*

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

9. *Adding a Resource to a Perspective*

Add one or more resources to the current perspective.

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

Configure Sybase Control Center

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.

11. *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.

12. *Changing the Screen Refresh Interval*

Use the Settings panel to modify the screen refresh interval, in seconds.

13. *Creating an Alert*

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

See also

- *User Authorization* on page 63

Configuring Policies for Monitoring and Administering Event Stream Processor

Edit the Project, Node, and Cluster policies in the `policy.xml` file to grant SCC monitoring and administrative access to a native OS, preconfigured login, LDAP, or SAP Business Intelligence (BI) group. Do not edit the `policy.xml` file if you are using Kerberos or RSA authentication for ESP as you are automatically granted SCC monitoring and administrative access.

The `policy.xml` file must be identical on every node in a cluster. In a multinode cluster where nodes are installed on different hosts, this is often accomplished by placing the ESP `security` directory on a shared drive. If your cluster's `policy.xml` file does not reside on a shared drive, make this change to the `policy.xml` for each node in the cluster.

1. Open `ESP-5_1\security\policy.xml`.
2. In the `<Role>` element, specify the group to which you wish to grant monitoring and administrative access and verify that the following is present:

The sample below grants monitoring and administrative access to a group called "sybase".

```
<Policies>

    <Policy type="Project">
        <Subjects>
            <Role>sybase</Role>
        </Subjects>
        <Resources>
            <!--The group has "read" privileges for "any"
project resource, including meta-data streams and other project
streams-->
            <Resource>*any</Resource>
        </Resources>
        <Actions>
```

```

        <Action>read</Action>
    </Actions>
</Policy>

<Policy type="Node">
    <Subjects>
        <Role>sybase</Role>
    </Subjects>
    <Resources>
        <Resource>Node</Resource>
    </Resources>
    <Actions>
        <Action>read</Action>
        <Action>stop</Action>
    </Actions>
</Policy>

<Policy type="Cluster">
    <Subjects>
        <Role>sybase</Role>
    </Subjects>
    <Resources>
        <Resource>Security</Resource>
        <Resource>Node</Resource>
        <Resource>Workspace</Resource>
        <Resource>Application</Resource>
    </Resources>
    <Actions>
        <Action>read</Action>
        <!--This privilege is required for write
operations, such as reload policy, add workspace/project, and so
on-->
        <Action>write</Action>
        <!--This privilege is required for stop
operations-->
        <Action>stop</Action>
        <!--The privilege is required for start
operations-->
        <Action>start</Action>
    </Actions>
</Policy>
</Policies>

```

To enable users within the group you specified in the `policy.xml` file to monitor and administer ESP, map this group to the `espMonitorRole` and `espAdminRole` roles.

See also

- *Updating Access Control* on page 168

Configuring Event Stream Processor for Monitoring

To enable users to monitor ESP node and cluster activity using Sybase Control Center, map a native OS, preconfigured login, LDAP, or SAP BI group to the `espMonitorRole` role.

The ESP node uses your corresponding authentication provider to determine which groups a user belongs to and then uses the `csi_role_mapping.xml` file to map these groups to the appropriate roles. This file is located in the `ESP-5_1\security` directory, which by default is installed in the `Sybase` directory. You may choose to map an existing group to `espMonitorRole` or create a new group.

1. Open `ESP-5_1\security\csi_role_mapping.xml`.

The `<Mapping>` element represents a mapping for a logical role. The `<LogicalName>` element represents the role that SCC checks for this mapping. There can only be one `<LogicalName>` per mapping. Do not modify this element.

2. Set the LDAP group within the `<MappedName>` element. This element represents the group you wish to map to the logical role. You can include more than one `<MappedName>` element for a mapping.

You can have the same `<MappedName>` element for two different `<LogicalName>` elements.

Here is an example of a mapping where the group "Administrators" maps to `espMonitorRole`:

```
<Mapping>
<LogicalName>espMonitorRole</LogicalName>
<MappedName>Administrators</MappedName>
</Mapping>
```

Here is an example of a mapping where several groups are mapped to `espMonitorRole`:

```
<Mapping>
<LogicalName>espMonitorRole</LogicalName>
<MappedName>IT</MappedName>
<MappedName>Developers</MappedName>
<MappedName>Operators</MappedName>
</Mapping>
```

See also

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 106
- *User Authorization* on page 114
- *Nodes* on page 169
- *Clusters* on page 181

Configuring Event Stream Processor for Administration

To enable users to start and stop ESP nodes, projects, and adapters from the Sybase Control Center Administration Console, and to view node and project log files, map a native OS, preconfigured login, LDAP, or SAP BI group to the espAdminRole role.

The ESP node uses your corresponding authentication provider to determine which groups a user belongs to and then uses the `csi_role_mapping.xml` file to map these groups to the appropriate roles. This file is located in the `ESP-5_1\security` directory, which by default is installed in the `Sybase` directory. You may choose to map an existing group to espAdminRole or create a new group.

1. Open `ESP-5_1\security\csi_role_mapping.xml`.

The `<Mapping>` element represents a mapping for a logical role. The `<LogicalName>` element represents the role that SCC checks for this mapping. There can only be one `<LogicalName>` per mapping. Do not modify this element.

2. Set the LDAP group within the `<MappedName>` element. This element represents the group you wish to map to the logical role. You can include more than one `<MappedName>` element for a mapping.

You can have the same `<MappedName>` element for two different `<LogicalName>` elements.

Here is an example of a mapping where the group "Administrators" maps to espAdminRole:

```
<Mapping>
<LogicalName>espAdminRole</LogicalName>
<MappedName>Administrators</MappedName>
</Mapping>
```

Here is an example of a mapping where several groups are mapped to espAdminRole:

```
<Mapping>
<LogicalName>espAdminRole</LogicalName>
<MappedName>IT</MappedName>
<MappedName>Developers</MappedName>
<MappedName>Operators</MappedName>
</Mapping>
```

See also

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 106
- *User Authorization* on page 114
- *Starting a Node* on page 179
- *Stopping a Node* on page 180
- *Projects* on page 194
- *Adapters* on page 198

Registering an ESP Node

Make Sybase Control Center aware of an ESP node (acting as a cluster manager, controller, or both) and its connection information by registering it as a resource.

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

Table 13. New Resource Type Details

Field	Description
Resource Name	(Required) Name of the resource to register. Set a name which will make the ESP node easily identifiable.
Resource Type	Select ESP Node .
Description	(Optional) A brief description to help you identify the resource.

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

Table 14. New Resource Connection Details

Field	Description
Host Name	(Required) Name of host on which the ESP node runs. Default value is localhost.
Port Number	(Required) Port number for the ESP node.
SSL is Enabled (Y/N)	(Required) Specify Y if SSL is enabled on the ESP Server, and N if it is not.

5. Click **Next**.
6. Enter authentication information based on the authentication type you specified while installing Event Stream Processor. This information is stored in the `ProductModule.xml` file, and must match the authentication type of the cluster in Event Stream Processor.

Authentication Type	Fields
Native OS	User name and password.
LDAP	User name and password.

Authentication Type	Fields
Kerberos	KDC, Realm, Service Name, User name, and Ticket Cache Location. See the <i>Event Stream Processor Administrators Guide</i> for more information.
RSA	User name, Password, and Keystore location.

- Or if you prefer not to authenticate now, click **I do not want to supply authentication information**.
7. Click **Next**.
 8. (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.
 9. (Optional) Click **Open the resource explorer to view this new resource**. (This option is not present when the Resource Explorer is open.)
The resource is added to the Resource Explorer even if you choose not to view it.
 10. Click **Finish**.
The ESP node is registered. If you have chosen to authenticate the resource and the authentication is successful, the other active nodes that belong to the same cluster as this node are registered and authenticated automatically because in Event Stream Processor, authentication is performed on a cluster-wide basis rather than on a per-node basis. These other nodes are named "<cluster_id>_<node_id>", by default, and "<cluster_id>_<node_id>(manager)" is for both manager type and dual (manager and controller) type nodes. You can edit the names of these nodes under resource properties.

See also

- *Authenticating an ESP Node* on page 131
- *Registering and Authenticating a Sybase Control Center Agent* on page 136
- *Registering Unregistered Nodes* on page 185

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, SAP 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 223
- *Unregistering a Resource* on page 223

Authenticating an ESP Node

If you did not authenticate the ESP node while registering it, you can authenticate afterwards from the Perspective Resources window.

Prerequisites

Register the node and add it to the perspective.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Authenticate**.
2. Select either:
 - **Use my current SCC login** - to authenticate using the SCC login user name and password. Select this option if you configured SCC to use the same authentication type as the cluster in Event Stream Processor and if both are authenticating against the same service instance. For example, this would be a convenient choice if you are using Native OS or LDAP authentication. If you select this, skip step 3 and go to step 4 directly.
 - **Specify different credentials** - to authenticate based on authentication type you specified while installing Event Stream Processor. This information is stored in the `ProductModule.xml` file.
3. Enter authentication information based on the authentication type you specified while installing Event Stream Processor. This information is stored in the `ProductModule.xml` file, and must match the authentication type of the cluster in Event Stream Processor.

Authentication Type	Fields
Native OS	User name and password.
LDAP	User name and password.
Kerberos	KDC, Realm, Service Name, and Ticket cache location.
RSA	User name, Password, and Keystore location.

- (Optional) Select **Remember these credentials for future sessions** to have SCC remember the credentials for the resource across sessions.
- Click **OK**.

The node is authenticated, and you can now monitor and administer the node.

See also

- *Registering and Authenticating a Sybase Control Center Agent* on page 136
- *Registering an ESP Node* on page 128
- *Registering Unregistered Nodes* on page 185

Update the Authentication Type

The authentication type set in the `ProductModule.xml` file must match the authentication type that the cluster in Event Stream Processor uses. If the cluster authentication type changes post-installation, update the authentication type in the `ProductModule.xml` file to reflect this change.

See the *Event Stream Processor Administrators Guide* for information on cluster authentication configuration.

The `ProductModule.xml` is located in the `<SCC_DIR>\plugins\ESPMAP` directory. There are four available authentication types you can configure in the file: Native OS, LDAP, RSA, and Kerberos.

Updating the Authentication Type to Native OS Authentication

If the cluster authentication type in Event Stream Processor has changed to Native OS, update the authentication type in the `ProductModule.xml` file to Native OS.

- Open the `ProductModule.xml` file in the `<SCC_DIR>\plugins\ESPMAP` directory.
- Uncomment out the `<scc:ap_definition>` block under Native OS to enable Native OS. For example:

```
<!-- Begin profile for Native OS Username/password
authentication -->
    <scc:ap_definition
name="ESP_Node_NativeOS_Authentication" description="Authenticate
with the ESP Node using Native OS security">
```

```

        <scc:ap_value_defn key="nativeos_username"
required="yes" default="">
        <scc:prompt reskey="RESKEY_AP_PROMPT_USERNAME"
default="Native OS Username"/>
        <scc:tooltip reskey="RESKEY_AP_TT_USERNAME"
default="Enter the username"/>
        </scc:ap_value_defn>
        <scc:ap_value_defn key="nativeos_password"
required="no" obscure="true" default="">
        <scc:prompt reskey="RESKEY_AP_PROMPT_PASSWORD"
default="Password"/>
        <scc:tooltip reskey="RESKEY_AP_TT_PASSWORD"
default="Enter your password"/>
        </scc:ap_value_defn>
    </scc:ap_definition>                                <!-- End profile for
Native OS username/password authentication -->

```

Note: Only one `scc:ap_definitions` block can take effect at a time so comment out blocks for all other authentication types.

Updating the Authentication Type to LDAP Authentication

If the cluster authentication type in Event Stream Processor has changed to LDAP, update the authentication type in the `ProductModule.xml` file to LDAP.

1. Open the `ProductModule.xml` file in the `<SCC_DIR>\plugins\ESPMAP` directory.
2. Uncomment out the `<scc:ap_definition>` block under LDAP authentication to enable LDAP authentication. For example:

```

<!-- The following scc:ap_definition shall be uncommented out if
LDAP authentication is enabled -->
- <!-- Profile for LDAP authentication -->
    <scc:ap_definition name="ESP_Node_LDAP_Authentication"
description="Authenticate with ESP Node by LDAP">
        <scc:ap_value_defn key="ldap_username" required="no"
default="sccadmin">
        <scc:prompt reskey="RESKEY_AP_PROMPT_USERNAME"
default="Username"/>
        <scc:tooltip reskey="RESKEY_AP_TT_USERNAME"
default="Enter the username"/>
        </scc:ap_value_defn>
        <scc:ap_value_defn key="ldap_password" required="no"
obscure="true" default="">
        <scc:prompt reskey="RESKEY_AP_PROMPT_PASSWORD"
default="Password"/>
        <scc:tooltip reskey="RESKEY_AP_TT_PASSWORD"
default="Enter your password"/>
        </scc:ap_value_defn>
    </scc:ap_definition>

```

Note: Only one `scc:ap_definitions` block can take effect at a time so comment out blocks for all other authentication types.

Updating the Authentication Type to RSA Authentication

If the cluster authentication type in Event Stream Processor has changed to RSA, update the authentication type in the `ProductModule.xml` file to RSA.

1. Open the `ProductModule.xml` file in the `<SCC_DIR>\plugins\ESPMAP` directory.
2. Uncomment out the `<scc:ap_definition>` block under RSA authentication to enable RSA authentication. For example:

```
<!-- The following scc:ap_definition shall be uncommented out if
RSA authentication is enabled
-->
- <!-- Profile for RSA authentication -->
    <scc:ap_definition name="ESP_Node_RSA_Authentication"
description="Authenticate with ESP Node by RSA">
    <scc:ap_value_defn key="rsa_username" required="no"
default="sccadmin">
    <scc:prompt reskey="RESKEY_AP_PROMPT_USERNAME"
default="Username"/>
    <scc:tooltip reskey="RESKEY_AP_TT_USERNAME"
default="Enter the username"/>
    </scc:ap_value_defn>
    <scc:ap_value_defn key="rsa_password" required="no"
default="">
    <scc:prompt reskey="RESKEY_AP_PROMPT_PASSWORD"
default="Password"/>
    <scc:tooltip reskey="RESKEY_AP_TT_PASSWORD"
default="Enter your password"/>
    </scc:ap_value_defn>
    <scc:ap_value_defn key="rsa_keystore" required="no"
default="">
    <scc:prompt reskey="RESKEY_AP_PROMPT_PASSWORD"
default="Keystore location"/>
    <scc:tooltip reskey="RESKEY_AP_TT_PASSWORD"
default="Enter your keystore location"/>
    </scc:ap_value_defn>
    </scc:ap_definition>
```

Note: Only one `scc:ap_definitions` block can take effect at a time so comment out blocks for all other authentication types.

Updating the Authentication Type to Kerberos Authentication

If the cluster authentication type in Event Stream Processor has changed to Kerberos, update the authentication type in the `ProductModule.xml` file to Kerberos.

1. Open the `ProductModule.xml` file in the `<SCC_DIR>\plugins\ESPMAP` directory.
2. Uncomment out the `<scc:ap_definition>` block under Kerberos to enable Kerberos authentication. For example:

```
<!-- Begin profile for Kerberos authentication -->
<scc:ap_definition
```



```

name="ESP_Node_Kerberos_Authentication"
description="Authentication with ESP Node by Kerberos
authentication">
    <scc:ap_value_defn key="kerberos_kdc"
required="yes" default="">
        <scc:prompt reskey="RESKEY_AP_PROMPT_KDC"
default="KDC"/>
            <scc:tooltip reskey="RESKEY_AP_TT_KDC"
default="Key Distribution Center"/>
        </scc:ap_value_defn>
    <scc:ap_value_defn key="kerberos_realm" required="yes"
default="">
        <scc:prompt reskey="RESKEY_AP_PROMPT_REALM"
default="Realm"/>
            <scc:tooltip reskey="RESKEY_AP_TT_REALM"
default="Kerberos realm"/>
        </scc:ap_value_defn>
    <scc:ap_value_defn key="kerberos_service_name"
required="yes" default="">
        <scc:prompt
reskey="RESKEY_AP_PROMPT_SERVICE_NAME" default="Service Name"/>
            <scc:tooltip
reskey="RESKEY_AP_TT_SERVICE_NAME" default="Kerberos Service
Name"/>
        </scc:ap_value_defn>
    <scc:ap_value_defn key="kerberos_username"
required="no" default="">
        <scc:prompt reskey="RESKEY_AP_PROMPT_USERNAME"
default="Username"/>
            <scc:tooltip reskey="RESKEY_AP_TT_USERNAME"
default="Kerberos client username"/>
        </scc:ap_value_defn>
    <scc:ap_value_defn key="kerberos_ticket_cache_location"
required="no" default="">
        <scc:prompt
reskey="RESKEY_AP_PROMPT_TICKET_CACHE_LOCATION" default="Ticket
Cache Location"/>
            <scc:tooltip
reskey="RESKEY_AP_TT_TICKET_CACHE_LOCATION" default="Kerberos
server ticket cache location"/>
        </scc:ap_value_defn>
    </scc:ap_definition>
<!-- End profile for Kerberos authentication -->

```

Note: Only one `scc:ap_definitions` block can take effect at a time so comment out blocks for all other authentication types.

Clearing Authentication Parameters

Clear authentication parameters to stop administrative access and unauthenticate the SCC agent for the node.

Prerequisites

- Register and authenticate a node.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

In the Perspective Resources window, select the ESP node, click the arrow, and select **Clear Authentication**.

The authenticated status displays "No". You are no longer able to monitor node or cluster activity.

See also

- *Registering and Authenticating a Sybase Control Center Agent* on page 136

Registering and Authenticating a Sybase Control Center Agent

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

The Sybase Control Center Agent for Event Stream Processor (or SCC agent) runs on the same host as the ESP node and enables Sybase Control Center for Event Stream Processor (which can be installed remotely) to manage the ESP node. The SCC agent is installed automatically as part of the Sybase ESP Server.

To perform certain administrative tasks, including starting a Sybase ESP node, you must register and authenticate the node's SCC agent.

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **ESP Nodes**.
4. In the right pane of the Administration Console, select an ESP node, click the arrow and select **Register Agent**.
5. (Optional) Set the node configuration file path, cluster log configuration file path, and the startup folder path.

If you set these parameters, provide them as fully-qualified, absolute paths and make sure they do not reference any environment variables. Also, if provided, the node configuration file must match the RPC port defined for the node in SCC.

If you do not provide values for these fields, the SCC agent attempts to discover the correct node configuration file by iterating over all subfolders under `$ESP_HOME/cluster/nodes` and looking for ".xml". It derives the other two parameters from this.

6. Enter the SCC agent port (the default port is 9999) and click **OK**.

The SCC agent host name is automatically set as the ESP node's host name and cannot be changed.

7. In the Administration Console, select the same ESP node, click the arrow, and select **Authenticate Agent**.

8. Enter the Sybase Control Center agent user (the default is `uafadmin`) and password.

Next

For instructions on changing the password for the SCC agent's default `uafadmin` account, see the topic on setting passwords in the *Sybase Control Center Installation Guide*.

See also

- *Authenticating an ESP Node* on page 131
- *Creating a Perspective* on page 139
- *Viewing the Node Log File* on page 220
- *Viewing the Project Log File* on page 221
- *Clearing Authentication Parameters* on page 136
- *Registering an ESP Node* on page 128
- *Registering Unregistered Nodes* on page 185

Viewing Sybase Control Center Agent Connection Information

View Sybase Control Center agent connection information in the server properties.

1. In the Perspective Resources view, select a resource.
2. From the application menu bar, select **View > Open > Administration Console**.
3. In the left pane of the Administration Console, select **ESP Nodes**.
4. Select the ESP node from the right pane and either:
 - Click the arrow to the right of the name and select **Properties**, or
 - From the Administration Console menu bar, select **Resource > Properties**.
5. Select **Agent** from the left pane.

Area	Description
Agent Page	<p>Agent registered – Indicates if the Sybase Control Center agent is registered: true or false.</p> <p>Agent authenticated – Indicates if the Sybase Control Center agent is authenticated: true or false.</p> <p>Agent status – Status of the Sybase Control Center agent: Running, Stopped, or Unknown.</p> <p>Agent host – Name of the host machine where the Sybase Control Center agent is running.</p> <p>Agent port – Port number on the host machine where the Sybase Control Center agent is running.</p> <p>Agent user – User name for authentication of the agent. Default is <i>uafadmin</i>.</p> <p>Agent process owner – The user name that owns the agent process.</p> <p>Agent home – The home directory of the Sybase Control Center agent.</p> <p>Agent version – The version of the Sybase Control Center agent.</p> <p>SCC agent plugin version – The agent plugin version of the Sybase Control Center agent.</p> <p>ESP directory – Installation directory of the ESP node with which the Sybase Control Center agent is associated.</p> <p>ESP version – Version of the ESP node with which the Sybase Control Center agent is associated.</p>

Parameters Required for Starting an ESP Node

When registering an SCC agent for an ESP node, there are three optional parameters that you can specify: node configuration file path, log configuration file path, and startup folder path.

The node configuration file path is the path to the node configuration file that the agent should use to start the ESP node. This is the equivalent to the `-cluster-node` parameter for the `esp_server` command.

The log configuration file path is the path to the log configuration file. This is the equivalent to the `-cluster-log-properties` parameter for the `esp_server` command.

The startup folder path is the path to the folder from where the ESP node was started. The agent looks for the startup folder because that is where the cluster log resides. Note that this is not necessarily the location of the `esp_server` binary. For example, if the `esp_server` binary is located under `$ESP_HOME/bin`, but your current directory is `/myhome/myserverrun`, you can:

```
> cd /myhome/myserverrun
> $ESP_HOME/bin/esp_server ..
```

In this example, `/myhome/myserverrun` would be the startup folder.

If you do not provide these parameters, the SCC agent attempts to discover the correct node configuration file and derives the other two parameters from this. The agent iterates over all subfolders under `$ESP_HOME/cluster/nodes` and looks for ".xml". It parses these files looking for those matching the node's port (as specified in node registration in SCC) and name (as specified during installation of Event Stream Processor). If it finds exactly one matching file, that is assumed to be the desired node configuration file.

If multiple matching configuration files are found this way, discovery fails and you will need to specify the node configuration manually. The same happens if the agent does not find any matching configuration files.

Once a node configuration file is found, or if the node configuration was specified in the agent registration, the agent attempts to locate the other parameters, if they are left blank. The startup location is assumed to be the folder where the node configuration file is located, and the log configuration is assumed to reside in the same folder as the node configuration file. Any `*.log.properties` files located in the same folder are used for log configuration, with preference given to the `cluster.log.properties` file. If multiple `.log.properties` files exist in the folder, or if none are found, the server starts without the `--cluster-log-config` option.

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

- *Registering and Authenticating a Sybase Control Center Agent* on page 136

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.

Configure Sybase Control Center

1. From the Sybase Control Center toolbar, click the **Launch Resource Explorer** icon.
2. Select the resources to add to your perspective. Use **Shift-click** or **Control-click** to select multiple resources.
3. Perform one of these actions:
 - Select **Resources > Add Resources to Perspective**.
 - Drag and drop resources from the Resource Explorer onto the Perspective Resources view. You can select and drag multiple resources.

See also

- *Removing a Resource from a Perspective* on page 224
- *Resources* on page 223

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

- *User Authorization* on page 114

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 resource 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**. If you choose **Later**, specify the start date and time.

Configure Sybase Control Center

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.

You cannot change the duration of a schedule (the once/repetitively/until setting) after you create it. To change the schedule duration, delete and recreate the schedule.

12. Click **Finish**.

See also

- *Changing the Screen Refresh Interval* on page 150
- *Job Scheduling* on page 204
- *Viewing or Deleting a Schedule* on page 208

About Statistics

Understand availability and performance statistics in Sybase Control Center.

The statistics you work with in Sybase Control Center can be divided into two types:

- Availability statistics are concerned with present conditions; they help you determine whether a resource you are monitoring (a server or an agent, for example) is running and functioning properly.
- Performance statistics are concerned with behavior of the same resources over time. They describe the flow of data through your environment. You can use performance statistics to spot trends, identify problems like resource bottlenecks, and make plans.

Sybase Control Center includes predefined key performance indicators (KPIs) for each product module; these KPIs are grouped into collections. KPIs such as server status, which serves as an availability statistic when it is fresh, have long-term value as historical performance statistics.

Availability statistics appear on the heat chart and on resource monitoring screens in each product module.

Performance statistics appear on the statistics chart and on resource monitoring screens in each product module.

Some KPIs are included in the default collection for each product module. To make other KPIs available to the heat chart, statistics chart, and resource monitoring views, you must set up collection jobs in the scheduler. See the data collections help topic for information on data collections and the KPIs contained in them.

Several configuration options affect the collection and display of data in Sybase Control Center:

- Collection repeat interval—The frequency of data collection. Set this on the collection job in the scheduler.
- Screen refresh interval—The period between screen refreshes. Refreshing the screen redraws it with the newest available data. Set the screen refresh interval in the product module. (May not be settable in all product modules.)
- Chart trend period—The period over which data is displayed in historical charts. Set the trend period in the product module. (May not be settable in all product modules.)

See also

- *Event Stream Processor Data Collections* on page 143
- *Key Performance Indicators for Event Stream Processor* on page 147

Event Stream Processor Data Collections

Predefined data collections you can schedule for Event Stream Processor. Collected statistics appear on Sybase Control Center monitoring screens and trigger user-configured alerts.

Sybase provides several data collections (without overlap) to collect data using a fine granularity and reduce the data collection workload. These are:

- Monitoring Statistics
- Collect_Overview_Project_Stream
- Collect_Conn_Publisher_Subscriber
- Collect_Adapter

The Monitoring Statistics collection is the default collection.

Table 15. ESP Node Data Collections

Collection	Description	KPIs
Monitoring Statistics	<p>Contains node availability statistics and total CPU usage displayed in the Heat Chart. This is the default collection; it is automatically scheduled when you authenticate an ESP node.</p> <hr/> <p>Note: Sybase strongly recommends that you leave this collection running for each monitored ESP node.</p> <hr/>	<ul style="list-style-type: none"> • Resource State • Total CPU Usage

Collection	Description	KPIs
Collect_Overview_Project_Stream	Schedule the collection to enable statistics and alerts for overall node activity, and projects and streams running on the node.	<p>Overview</p> <ul style="list-style-type: none"> • Total CPU Usage • Total System CPU Usage • Total User CPU Usage • Physical Memory Usage • Virtual Memory Usage • Total Thread Number • Number of Rows Received per Second • Number of Rows Sent per Second • Total Number of Connections • Number of Publishers • Number of Subscribers <p>Project</p> <ul style="list-style-type: none"> • Total CPU Usage • Total System CPU Usage • Total User CPU Usage • Physical Memory Usage • Virtual Memory Usage • Total Thread Number • Number of Publisher Rows Transferred • Number of Subscriber Rows Transferred • Number of Publisher Connections • Number of Subscriber Connections <p>Stream</p> <ul style="list-style-type: none"> • Rows per Sec • Total CPU Usage • Total System CPU Usage • Total User CPU Usage • Stream Depth Queued

Configure Sybase Control Center

Collection	Description	KPIs
		<ul style="list-style-type: none"> Stream Rows Stored
Collect_Conn_Publisher_Subscriber	Schedule the collection to enable statistics and alerts for gateway connections, publishers, and subscribers.	<p>Connection</p> <ul style="list-style-type: none"> Number of Rows Sent per Sec Number of Rows Received per Sec <p>Publisher</p> <ul style="list-style-type: none"> Total CPU Usage Total System CPU Usage Total User CPU Usage Number of Rows Sent per Sec <p>Subscriber</p> <ul style="list-style-type: none"> Total CPU Usage Total System CPU Usage Total User CPU Usage Stream Depth Queued Stream Rows Stored Number of Rows Received per Sec
Collect_Adapter	Schedule the collection to enable statistics and alerts for adapters on projects which are running on the node.	<p>Adapter</p> <ul style="list-style-type: none"> Adapter Total Rows Adapter Good Rows Adapter Bad Rows Adapter Latency

See also

- *About Statistics* on page 142
- *Key Performance Indicators for Event Stream Processor* on page 147

Key Performance Indicators for Event Stream Processor

Key performance indicators (KPIs) provide the statistics that appear on the charts in Sybase Control Center.

You can view these statistics from the **All Statistics** tab in the ESP Node Monitor view or by selecting a node, clicking the arrow, and selecting **Statistics Chart**.

Table 16. Overview Statistics

KPI	Description
Total System CPU Usage	(Percent) The sum of total CPU usage, since the last update, of all projects which are running on the node.
Total System CPU Usage	(Percent) The sum of total system (kernel on Windows) CPU usage, since the last update, of all projects which are running on the node.
Total User CPU Usage	(Percent) The sum of total user CPU usage, since the last update, of all projects which are running on the node.
Physical Memory Usage	(KB) The sum of physical memory usage, since the last update, of all projects which are running on the node.
Total Thread Number	(Count) The sum of the thread number, since the last update, of all projects which are running on the node.
Virtual Memory Usage	(KB) The sum of virtual memory usage, since the last update, of all project which are running on the node.
Total Number of Connections	(Count) Total number of connections on the node, including publishers and subscribers.
Number of Rows Sent per Second	(Count) The number of rows output from the projects running on the node, per second.
Number of Rows Received per Second	(Count) The number of rows input into the projects running on the node, per second.
Number of Publishers	(Count) Number of publishers running on the node.
Number of Subscribers	(Count) Number of subscribers running on the node.

Table 17. Project Statistics

KPI	Description
Total System CPU Usage	(Percent) The total CPU usage of the project since the last update.

KPI	Description
Total System CPU Usage	(Percent) The total system CPU usage for the project since the last update.
Total User CPU Usage	(Percent) The total user CPU usage for the project since the last update.
Physical Memory Usage	(KB) The physical memory usage for the project since the last update.
Total Thread Number	(Count) The thread number for the project since the last update.
Number of Publisher Rows Transferred	(Count) Number of rows input into the project per second.
Number of Subscriber Rows Transferred	(Count) Number of rows output from the project per second.
Virtual Memory Usage	(KB) The virtual memory usage for the project since the last update.
Number of Publisher Connections	(Count) Number of connections through which the publishers are running.
Number of Subscriber Connections	(Count) Number of connections through which the subscribers are running.

Table 18. Stream Statistics

KPI	Description
Total System CPU Usage	(Percent) The total CPU usage for the stream since the last update.
Total System CPU Usage	(Percent) The total system CPU usage for the stream since the last update.
Total User CPU Usage	(Percent) The total user CPU usage for the stream since the last update.
Rows per Sec	(Count) The number of rows processed by the stream, in seconds, since the last update.
Stream Depth Queued	(Count) Number of rows waiting to be processed.
Stream Rows Stored	(Count) The current number of records in the stream's store.

Table 19. Adapters Statistics

KPI	Description
AdapterTotalRows	Total number of rows in the adapter.
AdapterGoodRows	Number of good rows in the adapter.
AdapterBadRows	Number of bad rows in the adapter.
AdapterLatency	Time taken, in microseconds, for data to be processed.

Table 20. Publishers Statistics

KPI	Description
Total System CPU Usage	(Percent) Total CPU usage by the publisher's gateway thread.
Total System CPU Usage	(Percent) Total system CPU usage by the publisher's gateway thread.
Total User CPU Usage	(Percent) Total user CPU usage by the publisher's gateway thread.
Number of Rows Sent per Second	(Count) The number of data rows the client sent, per second, since the last update.

Table 21. Subscribers Statistics

KPI	Description
Total System CPU Usage	(Percent) Total CPU usage by the subscriber's gateway thread.
Total System CPU Usage	(Percent) Total system CPU usage by the subscriber's gateway thread.
Total User CPU Usage	(Percent) Total user CPU usage by the subscriber's gateway thread.
Number of Rows Received per Second	(Count) The number of data rows the client received, per second, since the last update.
Stream Rows Stored	(Count) The current number of records in the stream's store.
Stream Depth Queued	(Count) Number of rows waiting to be processed.

Table 22. Connection Statistics

KPI	Description
Number of Rows Sent per Second	(Count) The gateway client's performance, in data rows per second, sent by the client since the last update.
Number of Rows Received per Second	(Count) The gateway client's performance, in data rows per second, received since the last update.

See also

- *About Statistics* on page 142
- *Event Stream Processor Data Collections* on page 143
- *Graphing Performance Counters: the Statistics Chart* on page 166
- *Viewing All Statistics for a Node* on page 171

Changing the Screen Refresh Interval

Use the **Settings** panel to modify the screen refresh interval, in seconds.

1. In the Perspective Resources window, select a node, click the arrow, and select either:
 - **Monitor Node** to change the screen refresh interval for the ESP Node Monitor view.
 - Or **Monitor Cluster** to change the screen refresh interval for the ESP Cluster Monitor view.
2. In the left pane of either view, select **Settings**.
3. Change the value in the **Screen refresh interval (Seconds)** field. The number must be an integer, floating points are not allowed.

The minimum value is 5, maximum value is 999999, and default value is 30 seconds.

Setting the value too low can have a negative performance impact, and setting the value too high means you do not get system performance updates in a timely manner.

4. Click **Apply**.
The screen refresh interval is updated.

See also

- *Setting Up Statistics Collection* on page 141
- *Nodes* on page 169
- *Clusters* on page 181

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 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 23. 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.

See also

- *Assigning a Role to a Login or a Group* on page 115
- *Configuring the E-mail Server* on page 111
- *Alerts* on page 210
- *Testing an Alert-Triggered Script* on page 213

Event Stream Processor Alerts

Alerts that you can use for the Event Stream Processor Server.

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

Note: To configure alerts for an ESP cluster, register each node for which you want to set alerts and add the node to a perspective. Then set up alerts for each node using the Resource Properties view (select a server in the Perspective Resources view and select **Resource > Properties > Alerts**

Table 24. ESP Node Alerts

Alert	Description	Alert Type
ESP Node Availability	Indicates the availability of an ESP node.	State
Total CPU Availability	The sum of total CPU usage, since the last update, of all projects which are running on the node. This alert is the same as the Total CPU Availability alert for ESP Overview.	Threshold

Table 25. ESP Overview Alerts

Alert	Description	Alert Type
Total CPU Availability	The sum of total CPU usage, since the last update, of all projects which are running on the node. This alert is the same as the Total CPU Availability alert for the ESP node.	Threshold
Total System CPU Availability	The sum of total system (kernel on Windows) CPU usage, since the last update, of all projects which are running on the node.	Threshold
Total User CPU Availability	The sum of total user CPU usage, since the last update, of all projects which are running on the node.	Threshold
Physical Memory Availability	The sum of physical memory usage, since the last update, of all projects which are running on the node.	Threshold
Virtual Memory Availability	The sum of virtual memory usage, since the last update, of all projects which are running on the node.	Threshold

Alert	Description	Alert Type
Total Threads	The sum of the thread number, since the last update, of all projects which are running on the node.	Threshold
Publisher Connections	The total number of publishers running on the node.	Threshold
Rows Received per Second	The total number of rows input into the projects running on the node, per second.	Threshold
Rows Sent per Second	Total number of rows output from the projects running on the node, per second.	Threshold
Subscriber Connections	The total number of subscribers running on the node.	Threshold
Total Connections	The total number of connections on the node, including publishers and subscribers.	Threshold

Table 26. Project Alerts

Alert	Description	Alert Type
Total CPU Availability	The total CPU usage of the project since the last update.	Threshold
Total System CPU Availability	The total system CPU usage for the project since the last update.	Threshold
Total User CPU Availability	The total user CPU usage for the project since the last update.	Threshold
Physical Memory Availability	The physical memory usage for the project since the last update.	Threshold
Virtual Memory Availability	The virtual memory usage for the project since the last update.	Threshold
Total Threads	The thread number for the project since the last update.	Threshold
Publisher Rows Transferred	Number of rows input into the project per second.	Threshold

Alert	Description	Alert Type
Subscriber Rows Transferred	Number of rows output from the project per second.	Threshold
Publisher Connections	Number of connections through which publishers are running.	Threshold
Subscriber Connections	Number of connections through which subscribers are running.	Threshold
Total Connections	The total number of connections on the project, including publishers and subscribers.	Threshold

Table 27. Stream Alerts

Alert	Description	Alert Type
Total CPU Availability	The total CPU usage of the stream since the last update.	Threshold
Total System CPU Availability	The total system CPU usage for the stream since the last update.	Threshold
Total User CPU Availability	The total user CPU usage for the stream since the last update.	Threshold
Rows per Sec	Number of rows processed by the stream, in a second, since the last update.	Threshold
Stream Depth Queued	Number of rows waiting to be processed.	Threshold
Stream Rows Stored	The current number of records in the stream's store.	Threshold

Table 28. Gateway Connection Alerts

Alert	Description	Alert Type
Rows Sent per Sec	The gateway client connection performance. This is measured by the amount of data rows sent by the client, per second, since the last update.	Threshold

Alert	Description	Alert Type
Rows Received per Sec	The gateway client's performance. This is measured by the amount of data rows received by the client, per second, since the last update.	Threshold

Table 29. Publisher Alerts

Alert	Description	Alert Type
Total CPU Availability	The total CPU usage by the publisher's gateway thread.	Threshold
Total System CPU Availability	The total system CPU usage by the publisher's gateway thread.	Threshold
Total User CPU Availability	The total user CPU usage by the publisher's gateway thread.	Threshold
Rows Sent per Sec	The number of data rows the client sent, per second, since the last update.	Threshold

Table 30. Subscriber Alerts

Alert	Description	Alert Type
Total CPU Availability	The total CPU usage by the subscriber's gateway thread.	Threshold
Total System CPU Availability	The total system CPU usage by the subscriber's gateway thread.	Threshold
Total User CPU Availability	The total user CPU usage by the subscriber's gateway thread.	Threshold
Stream Depth Queued	Number of rows waiting to be processed.	Threshold
Stream Rows Stored	The current number of records in the stream's store.	Threshold
Rows Received per Sec	The number of data rows the client received, per second, since the last update.	Threshold

Table 31. Adapter Alerts

Alert	Description	Alert Type
Adapter Bad Rows	Number of bad rows in the adapter.	Threshold
Adapter Latency	Time it takes for the adapter to process data.	Threshold

See also

- *Alert Types and Severities for Event Stream Processor* on page 158
- *Alert-Triggered Scripts* on page 159
- *Alert-Triggered Script Examples* on page 160
- *Substitution Parameters for Scripts* on page 161
- *Alerts* on page 210
- *Viewing Overview Statistics and Alerts for a Node* on page 169
- *Viewing Overview Statistics and Alerts for a Cluster* on page 182

Alert Types and Severities for Event Stream Processor

Learn about the properties that define and control alerts.

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

Table 32. 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 33. Alert Severities

Severity	Description
Normal	No alert is issued.

Severity	Description
Warning	A problem has given cause for concern. An alert is issued; you can subscribe to alerts that fire at the Warning level.
Critical	A serious problem exists. An alert is issued; you can subscribe to alerts that fire at the Critical level.

See also

- *Event Stream Processor Alerts* on page 153
- *Alert-Triggered Scripts* on page 159
- *Alert-Triggered Script Examples* on page 160
- *Substitution Parameters for Scripts* on page 161
- *Alerts* on page 210
- *Viewing Overview Statistics and Alerts for a Node* on page 169
- *Viewing Overview Statistics and Alerts for a Cluster* on page 182

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\\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

- *Event Stream Processor Alerts* on page 153
- *Alert Types and Severities for Event Stream Processor* on page 158
- *Alert-Triggered Script Examples* on page 160
- *Substitution Parameters for Scripts* on page 161
- *Testing an Alert-Triggered Script* on page 213
- *Alerts* on page 210

Alert-Triggered Script Examples

Sample scripts for Windows and UNIX.

Example 1: An Alert-Triggered Windows Script

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 output 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:
```

Example 2: An Alert-Triggered UNIX Script

This is a UNIX script. Like the Windows script above, it outputs the parameter values you pass to it to a text file.

```
#!/bin/sh
outfile=/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

- *Event Stream Processor Alerts* on page 153
- *Alert Types and Severities for Event Stream Processor* on page 158
- *Alert-Triggered Scripts* on page 159
- *Substitution Parameters for Scripts* on page 161

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 34. 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.

Parameter	Description
%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 35. 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

Parameter	Description
%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

- *Event Stream Processor Alerts* on page 153
- *Alert Types and Severities for Event Stream Processor* on page 158
- *Alert-Triggered Scripts* on page 159
- *Alert-Triggered Script Examples* on page 160
- *Testing an Alert-Triggered Script* on page 213
- *Modifying an Alert* on page 212

Configure Sybase Control Center

Manage and Monitor Event Stream Processor

Administer the Sybase Event Stream Processor environment.

Displaying Resource Availability: the Heat Chart

Use the heat chart to view the status and availability of servers 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.

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.

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**.

Graphing Performance Counters: the Statistics Chart

To show performance trends, generate a graph for any set of performance counters over a specified period of time.

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 > 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. Increments depend 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

- *Configuring Repository Purging* on page 244
- *Viewing All Statistics for a Node* on page 171
- *Key Performance Indicators for Event Stream Processor* on page 147

Managing Workspaces

Add or remove workspaces in Event Stream Processor from a manager node in the Administration Console. From a manager node, you can view all existing workspaces in the cluster and create new or remove existing workspaces as necessary. One workspace can have several projects on it, but you need to create a workspace before adding projects to it or adding the project to a cluster.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espAdminRole role to be able to perform this task.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the right pane of the Administration Console, select a manager node, click the arrow, and select **Manage Workspaces...**
3. A new window displays and lists all the workspaces on the cluster to which the manager node belongs. You can:

- Click **Add workspace** to add a new workspace to the cluster. A new window displays. Specify the name for the new workspace you are adding.
- Select a workspace and click **Remove workspace** to remove that workspace from the cluster. This stops and removes all running projects from the workspace and then deletes the workspace.
- Click **Cancel** to exit the window.

Updating Access Control

If you are using the `policy.xml` file to manage access control, you can edit the `policy.xml` file online and reload the policy from the Administration Console without restarting the cluster. You can only do this from a manager node.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Register and authenticate the SCC agent for the node.
- Log in using the `espAdminRole` role to be able to perform this task.
- Enable the `policy.xml` file in the cluster configuration file.

Task

For details on configuring access control using the `policy.xml` file, see the *SAP Sybase Event Stream Processor Configuration and Administration Guide*.

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. Edit the `policy.xml` in the `$ESP_HOME\security` directory.
3. Select a node, click the arrow, and select **Reload Policy** to reload the `policy.xml` file. A new window displays indicating whether the request to reload the policy was issued. If it was issued successfully, the `policy.xml` file is reloaded without restarting the cluster.

See also

- *Configuring Policies for Monitoring and Administering Event Stream Processor* on page 124

Nodes

After you register and authenticate an ESP node, you can monitor its availability and performance using the Monitor Node view. Log in using the espMonitorRole role to monitor a node.

Statistics for the node are available only if the time-granularity option is enabled on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

See also

- *Configuring Event Stream Processor for Monitoring* on page 126
- *Changing the Screen Refresh Interval* on page 150

Viewing Overview Statistics and Alerts for a Node

View alerts for and monitor the performance of the selected ESP node by viewing overview statistics such as CPU, memory, and thread usage.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Overview**.

Name	Statistics
Server	<p>Node name - name of the ESP node.</p> <p>Host - host that server is currently running on.</p> <p>Node version - version of the ESP node.</p> <p>Platform - platform of the machine that the node is running on.</p> <p>Node type - whether node is a manager or controller.</p> <p>State - valid values are running or unknown.</p>
Alerts	<p>This is the header table containing all alerts that have fired for the selected ESP node after you open the ESP Node Monitor view. If you log off without closing this view, alerts that have fired since you logged back on display.</p> <p>Time - when the alert is triggered.</p> <p>Alert Name - name of the alert. This is based on the KPI.</p> <p>Resource - the resource for which the alert is triggered.</p> <p>Severity - alert severity rating. Possible severity ratings are Normal, Warning, or Critical, and are based on ranges of values you specified when setting the alert threshold.</p> <p>Value - the KPI value. The alert is triggered when the KPI value falls within the range of values you specified when setting the alert threshold.</p> <p>Threshold - the range of values you assigned to alert severity ratings. For example, if the low value for the Normal rating is 0 and the high value 100, the threshold is the range of 0 to 100.</p>

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The data on the graph starts displaying from the time you open the ESP Node Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.
Memory Usage History	Line graph displaying total memory usage over time, in kilobytes (KB). The data on the graph starts displaying from the time you open the ESP Node Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.
Thread Usage History	Line graph displaying number of threads used over time. The data on the graph starts displaying from the time you open the ESP Node Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.

See also

- *Viewing Overview Statistics and Alerts for a Cluster* on page 182
- *Alerts* on page 210
- *Event Stream Processor Alerts* on page 153
- *Alert Types and Severities for Event Stream Processor* on page 158

Viewing All Statistics for a Node

Monitor performance of a selected ESP node and the projects, streams, connections, adapters, publishers, and subscribers on the node. View statistics such as CPU usage, number of bytes and rows received and sent, and number of threads in use.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **All Statistics**.

You can view all statistics available for a node and the projects, streams, connections, adapters, publishers, and subscribers on the node.

See also

- *Graphing Performance Counters: the Statistics Chart* on page 166
- *Key Performance Indicators for Event Stream Processor* on page 147

Viewing Statistics for Projects on a Node

Monitor the performance of projects within the selected ESP node by viewing statistics such as CPU, memory, and thread usage, and rows received and sent.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Projects**.

Table Name	Statistics
Projects	This is the header table containing all projects within the selected ESP node. Workspace Name - name of workspace to which the project belongs. Project Name Status - project status. Valid values are running, stopped, or unknown.

These statistics display for the project you select in the header table:

Tab	Statistics
System	<p>The data on the graph starts displaying from the time you open the Node Monitor - Projects view.</p> <p>CPU History - line graph displaying the percentage of total CPU usage over time.</p> <p>Memory Usage History - line graph displaying total memory usage over time, in KB.</p> <p>Thread Usage History - line graph displaying number of threads used over time.</p>
Network	<p>Rows Transferred History - line graph displaying rows received and rows sent per second over time.</p> <p>The data on the graph starts displaying from the time you open the Node Monitor - Projects view.</p> <p>Connections History - line graph displaying total number of publishers, subscribers, and connections over time.</p>

3. In the right pane of the ESP Node Monitor view, select a project, click the arrow, and select **Project Properties**.

A new window displays and contains the following statistics for the project you selected:

- Workspace (name of workspace to which the project belongs)
- Project
- Status (project status; valid values are running, stopped, or unknown)
- Command Host (physical host on which the project is running)
- Command Port (command port on which the project command control gateway is listening)
- Gateway Host
- Gateway Port
- SQL Port (port assigned to the project instance for serving SQL query requests)
- SSL Enabled
- Big Endian
- Address Size (the size, in bytes, of a memory address on the deployed architecture)
- Data Size (the size, in bytes, of the date datatype on the deployed architecture)
- Money Precision
- WS Enabled (whether the ESP project is enabled for Web service access)
- Timer Interval (the value in the ESP project "timer-granularity" option)
- Active-Active (whether the project is running in active-active or high availability mode)

See also

- *Viewing Statistics for Projects on a Cluster* on page 186
- *Viewing Projects* on page 194
- *Starting a Project* on page 195
- *Stopping a Project* on page 196

Viewing Statistics for Streams on a Node

Monitor the performance of streams on a selected ESP node by viewing statistics such as CPU history and number of rows processed.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Streams**.

Table Name	Stream Details
Streams	<p>This is the header table containing all streams within the selected ESP node.</p> <p>Workspace - name of workspace on which the stream is running.</p> <p>Project - name of project to which the stream is attached.</p> <p>Stream - name of stream attached to the project running on the node.</p> <p>Queue Depth - the number of rows waiting to be processed.</p> <p>Rows In Store - the current number of records in the stream's store.</p>

These statistics display for the stream you select in the header table:

Statistics	Description
Rows Transferred History	A line graph displaying the number of rows processed per second.

Statistics	Description
CPU History	A line graph displaying the percentage of total CPU usage over time. The data on the graph starts displaying from the time you open the ESP Node Monitor view.

See also

- *Viewing Statistics for Streams on a Cluster* on page 188

Viewing Statistics for Connections on a Node

Monitor the performance of publishers and subscribers on a selected ESP node by viewing connection statistics such as number of rows received and sent.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Connections**.

These statistics display:

Statistics	Description
Connections	Line graph displaying total number of connections to the node.
Publishers and Subscribers	Line graph displaying number of publishers and number of subscribers of the node.
Rows Transferred History	Line graph displaying rows received and rows sent per second. The data on the graph starts displaying from the time you open the Node Monitor - Connections view.

See also

- *Viewing Statistics for Connections on a Cluster* on page 189

Viewing Statistics for Adapters on a Node

Monitor performance of adapters on a selected ESP node by viewing statistics such as numbers of rows processed.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Adapters**.

Table Name	Adapter Details
Adapters	<p>This is the header table containing all adapters within the selected ESP node.</p> <p>Workspace - name of workspace on which the project is running.</p> <p>Project - name of project to which the adapter is attached.</p> <p>Name - name of the adapter.</p> <p>Status - adapter status. Valid values are initial, stopped, ready, continuous, idle, done, or unknown.</p> <p>Stream - the stream in which the adapter operates.</p> <p>Type - the unique adapter ID assigned to each adapter. For example, sybase_ase_out or sybase_iq_out. See the Adapters Guide for more information.</p> <p>In/Out - displays whether the adapter is an input or output adapter.</p> <p>Group - the adapter group to which the adapter belongs.</p>

These statistics display for the adapter you select in the header table:

Table 36. Common Statistics

Name	Value
AdapterTotalRows	Total number of rows in the adapter.

Name	Value
AdapterGoodRows	Number of good rows in the adapter.
AdapterBadRows	Number of bad rows in the adapter.
AdapterLatency	Time taken for data to be processed.

Table 37. Extended Statistics

Name	Value
<Custom Statistic Name>	If the adapter you select has custom statistics to report, these display here. Note that not all adapters have custom statistics.

See also

- *Viewing Adapters* on page 198
- *Starting an Adapter* on page 199
- *Stopping an Adapter* on page 200
- *Viewing File Activity for the Sybase IQ Output Adapter* on page 201
- *Viewing Statistics for Adapters on a Cluster* on page 190

Viewing Statistics for Publishers on a Node

Monitor the performance of publishers on a selected ESP node by viewing statistics such as CPU history and number of rows sent.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Publishers**.

Table Name	Publisher Details
Publishers	<p>This is the header table containing all publishers within the selected ESP node.</p> <p>Workspace - name of workspace to which the project belongs.</p> <p>Project - name of the project to which the publisher is connected.</p> <p>IP Address - IP address of the publisher.</p> <p>Port - port number on the publisher's machine used to send subscribed data to Event Stream Processor.</p>

These statistics display for the publisher you select from the header table:

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The user CPU usage and system CPU usage are also shown on this line graph. The data on the graph starts displaying from the time you open the Node Monitor - Publishers view.
Rows Transferred History	Line graph displaying rows sent per second. The data on the graph starts displaying from the time you open the Node Monitor - Publishers view.

See also

- *Viewing Statistics for Publishers on a Cluster* on page 191

Viewing Statistics for Subscribers on a Node

Monitor the performance of subscribers on a selected ESP node by viewing statistics such as CPU history and number of rows received. SCC also uses subscribers to gather information used for monitoring. These SCC subscribers also display within this view.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Subscribers**.

Table Name	Subscriber Details
Subscribers	<p>This is the header table containing all subscribers within the selected ESP node.</p> <p>Workspace - name of the workspace to which the subscriber is connected.</p> <p>Project - name of the project to which the subscriber is connected.</p> <p>IP Address - IP address of the subscriber.</p> <p>Port - port number on the subscriber's machine used to receive published data from Event Stream Processor.</p> <p>Stream Name - name of the data stream to which the subscriber is subscribed.</p> <p>Queue Depth - number of rows waiting to be received by the subscriber.</p> <p>Rows in Store - number of rows in the subscriber's data store.</p>

These statistics display for the subscriber you select from the header table:

Statistics	Description
CPU History	<p>Line graph displaying the percentage of total CPU usage over time. The user CPU usage and system CPU usage are also shown on this line graph.</p> <p>The data on the graph starts displaying from the time you open the Node Monitor - Subscribers view.</p>
Rows Transferred History	<p>Line graph displaying rows received per second.</p> <p>The data on the graph starts displaying from the time you open the Node Monitor - Subscribers view.</p>

See also

- *Viewing Statistics for Subscribers on a Cluster* on page 192

Starting a Node

Start a node from the Administration Console.

Prerequisites

- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the right pane of the Administration Console, select an ESP node that is not currently running, click the arrow, and select **Start Node**.
3. A window appears and asks you to confirm whether you want to start the node. Click **Yes** to confirm.
A window appears and reports whether the node was successfully started.

See also

- *Configuring Event Stream Processor for Administration* on page 127
- *Stopping a Node* on page 180
- *Projects* on page 194
- *Adapters* on page 198
- *Starting a Project* on page 195
- *Starting an Adapter* on page 199

Stopping a Node

Stop a node from the Administration Console.

Prerequisites

- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

Stopping a node this way does not stop any projects running on the node unless the node is the only manager node.

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the right pane of the Administration Console, select an ESP node that is currently running, click the arrow, and select **Stop Node**.

3. A window appears and asks you to confirm whether you want to stop the node. Click **Yes** to confirm.
A window appears and reports whether the node was successfully stopped.

See also

- *Configuring Event Stream Processor for Administration* on page 127
- *Starting a Node* on page 179
- *Projects* on page 194
- *Adapters* on page 198
- *Stopping a Project* on page 196
- *Stopping an Adapter* on page 200

Viewing Schema for a Stream

View schema for a selected stream from the ESP Node Monitor view.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Node**.
2. In the left pane of the ESP Node Monitor view, select **Streams**.
3. In the right pane of the ESP Node Monitor view, select a stream, click the arrow, and select **Show Schema**.

A window appears displaying the schema for the selected stream (for example, column name, type, and primary key).

Clusters

After you register and authenticate an ESP manager node, you can monitor the availability and performance of the cluster to which it belongs using the Monitor Cluster view. This functionality is not available for controller nodes.

Log in with the espMonitorRole role to monitor a cluster. Statistics for the node and cluster are available only if the time-granularity option is enabled on the node in Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for more information on the project configuration file.

See also

- *Configuring Event Stream Processor for Monitoring* on page 126
- *Changing the Screen Refresh Interval* on page 150

Viewing Overview Statistics and Alerts for a Cluster

View alerts for and monitor the performance of the selected ESP cluster by viewing overview statistics such as CPU, memory, and thread usage history.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the resource, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Overview**.

Name	Statistics
Alerts	<p>This is the header table containing all alerts that have fired for the nodes in the cluster after you open the ESP Cluster Monitor view. If you log off without closing this view, alerts that have fired since you logged back on display.</p> <p>Time - when the alert is triggered.</p> <p>Alert Name - name of the alert. This is based on the KPI.</p> <p>Resource - the resource for which the alert is triggered.</p> <p>Severity - alert severity rating. Possible severity ratings are Normal, Warning, or Critical, and are based on ranges of values you specified when setting the alert threshold.</p> <p>Value - the KPI value. The alert is triggered when the KPI value falls within the range of values you specified when setting the alert threshold.</p> <p>Threshold - the range of values you assigned to alert severity ratings. For example, if the low value for the Normal rating is 0 and the high value 100, the threshold is the range of 0 to 100.</p>

These statistics display:

Statistics	Description
CPU History	Line graph displaying the average CPU user usage, CPU system usage, and total CPU usage across the cluster. The data on the graph starts displaying from the time you open the ESP Cluster Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.
Memory Usage History	Line graph displaying average total memory usage across the cluster, in kilobytes (KB). The data on the graph starts displaying from the time you open the ESP Cluster Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.
Thread Usage History	Line graph displaying the average number of threads used across the cluster. The data on the graph starts displaying from the time you open the ESP Cluster Monitor view. If you log off without closing this view, statistics from the time that you logged back on display.

See also

- *Viewing Overview Statistics and Alerts for a Node* on page 169
- *Alerts* on page 210
- *Event Stream Processor Alerts* on page 153
- *Alert Types and Severities for Event Stream Processor* on page 158

Viewing Topology for a Cluster

Display the topology view of a cluster which represents the entire cluster and consists of nodes and database servers.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

Topology is a graphical representation of how the ESP nodes and Sybase IQ, Adaptive Server Enterprise, and HANA database servers are connected to each other. It is a network diagram

Manage and Monitor Event Stream Processor

that provides a visual map and the status of Event Stream Processor nodes, output adapters, and the databases to which they connect.

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.

2. In the left pane of the ESP Cluster Monitor view, select **Topology**.

Nodes appear as icons and indicate different node roles. Valid roles include:

- Manager
- Controller
- Manager and Controller

Database servers appear as icons. You can connect ASE, Sybase IQ, and HANA database servers to any node in the cluster using an output adapter. You can also right-click the Sybase IQ and ASE servers to register unregistered servers.

Valid database servers include:

- ASE
- Sybase IQ
- HANA

3. To view node status, hover the mouse pointer over a selected node.

Valid statuses include:

- Running
- Unknown

4. To view if a node is authenticated, right-click on a selected node. If the **Monitor Node** command appears, the node is authenticated.

5. A line connecting controller nodes to Sybase IQ, ASE, and HANA database servers with one or more adapters in loaded projects. To view adapter status, hover the mouse pointer over a selected line.

Valid statuses include:

- Running
- Stopped
- Mixed

Viewing Nodes Within a Cluster

Monitor selected ESP nodes in the cluster by viewing information about their roles, statuses, and authentication.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Nodes**.

Table Name	Details
Nodes	<p>Node - name of node on the cluster.</p> <p>Host - host name where the server is running.</p> <p>Port - port number where the server is running.</p> <p>Role - role of the node within SCC. Valid roles include:</p> <ul style="list-style-type: none"> • Controller • Manager • Manager and Controller <p>Status - current status of the node. Valid states include:</p> <ul style="list-style-type: none"> • Running • Unknown <p>Authenticated - "Yes" if node is authenticated, "No" if node is not authenticated.</p>
Node Details	<p>Name - name of node.</p> <p>Version - version of node.</p> <p>Build information - details of build, including time and date.</p> <p>Platform - operating system in which SCC is running.</p>

3. (Optional) In the ESP Cluster Monitor view, select a node, click the arrow, and the following options appear:
 - **Monitor Node**
 - **Register Unregistered Nodes**
 - **Properties**

Registering Unregistered Nodes

If you have already registered a node, you can register other associated nodes from the Monitor Cluster view.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

Manage and Monitor Event Stream Processor

- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Nodes**.
3. In the ESP Cluster Monitor view, select a node, click the arrow, and select **Register Unregistered Nodes**.
4. A window appears stating that all unregistered nodes are now registered. Click **OK** to confirm.
The registered nodes display in the **Nodes** pane.

See also

- *Registering an ESP Node* on page 128
- *Authenticating an ESP Node* on page 131
- *Registering and Authenticating a Sybase Control Center Agent* on page 136

Viewing Statistics for Projects on a Cluster

Monitor the performance of projects within the selected ESP cluster by viewing statistics such as CPU, memory, and thread usage, and rows received and sent.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Projects**.

Table Name	Statistics
Projects	<p>This is the header table containing all projects within the selected ESP cluster.</p> <p>Node - name of the ESP node on which the workspace is running.</p> <p>Workspace Name - name of workspace to which the project belongs.</p> <p>Project Name</p> <p>Status - project status. Valid values are running, stopped, or unknown.</p>

These statistics display for the project you select in the header table:

Tab	Statistics
System	<p>CPU History - line graph displaying the total CPU usage over time for the project. The user CPU usage and system CPU usage are also shown on this line graph.</p> <p>Memory Usage History - line graph displaying the total memory usage over time for the project.</p> <p>Thread Usage History - line graph displaying the total thread usage over time for the project.</p> <p>The data on the graph starts displaying from the time you open the Cluster Monitor - Projects view.</p>
Network	<p>Rows Transferred History - line graph displaying rows received and rows sent per second over time.</p> <p>Connections History - line graph displaying total number of publishers, subscribers, and connections over time.</p> <p>The data on the graph starts displaying from the time you open the Cluster Monitor - Projects view.</p>

3. In the right pane of the ESP Cluster Monitor view, select a project, click the arrow, and select **Project Properties**.
 A new window displays and contains the following statistics for the project you selected:
 - Workspace (name of workspace to which the project belongs)
 - Project
 - Status (project status; valid values are running, stopped, or unknown)
 - Command Host (physical host on which the project is running)
 - Command Port (command port on which the project command control gateway is listening)
 - Gateway Host
 - Gateway Port
 - SQL Port (port assigned to the project instance for serving SQL query requests)
 - SSL Enabled
 - Big Endian
 - Address Size (the size, in bytes, of a memory address on the deployed architecture)
 - Data Size (the size, in bytes, of the date datatype on the deployed architecture)
 - Money Precision
 - WS Enabled (whether the ESP project is enabled for Web service access)
 - Timer Interval (the value in the ESP project "timer-granularity" option)

Manage and Monitor Event Stream Processor

- Active-Active (whether the project is running in active-active or high availability mode)

See also

- *Viewing Statistics for Projects on a Node* on page 172
- *Viewing Projects* on page 194
- *Starting a Project* on page 195
- *Stopping a Project* on page 196

Viewing Statistics for Streams on a Cluster

Monitor the performance of streams on an ESP cluster by viewing statistics such as CPU history and number of rows processed.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Streams**.

Table Name	Stream Details
Streams	<p>This is the header table containing all streams within the selected ESP cluster.</p> <p>Node - name of the ESP node on which the workspace is running.</p> <p>Workspace - name of workspace on which the stream is running.</p> <p>Project - name of project to which the stream is attached.</p> <p>Stream - name of stream attached to the project running on the node.</p> <p>Queue Depth - the number of rows waiting to be processed.</p> <p>Rows In Store - the current number of records in the stream's store.</p>

These statistics display for the stream you select in the header table:

Statistics	Description
Rows Transferred History	A line graph displaying the number of rows processed per second by the stream over time. The data on the graph starts displaying from the time you open the Cluster Monitor - Streams view.
CPU History	A line graph displaying the total CPU usage over time for the stream. The user CPU usage and system CPU usage are also shown on this line graph. The data on the graph starts displaying from the time you open the Cluster Monitor - Streams view.

See also

- *Viewing Statistics for Streams on a Node* on page 174

Viewing Statistics for Connections on a Cluster

Monitor the performance of publishers and subscribers on an ESP cluster by viewing connection statistics such as number of rows received and sent.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Connections**.

These statistics display:

Statistics	Description
Connections	Line graph displaying total number of connections across the cluster over time. The data on the graph starts displaying from the time you open the Cluster Monitor - Connections view.
Publishers and Subscribers	Line graph displaying total number of publishers and subscribers across the cluster over time. The data on the graph starts displaying from the time you open the Cluster Monitor - Connections view.

Statistics	Description
Rows Transferred History	Line graph displaying total number of rows received and sent per second across the cluster over time. The data on the graph starts displaying from the time you open the Cluster Monitor - Connections view.

See also

- *Viewing Statistics for Connections on a Node* on page 175

Viewing Statistics for Adapters on a Cluster

Monitor performance of adapters on an ESP cluster by viewing statistics such as number of rows processed.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Adapters**.

Name	Adapter Details
Adapters	<p>This is the header table containing all adapters within the selected ESP cluster.</p> <p>Node - name of the ESP node on which the workspace is running.</p> <p>Workspace - name of workspace on which the project is running.</p> <p>Project - name of project to which the adapter is attached.</p> <p>Name - name of the adapter.</p> <p>Status - adapter status. Valid values are initial, stopped, ready, continuous, idle, done, or unknown.</p> <p>Stream - the stream in which the adapter operates.</p> <p>Type - the unique adapter ID assigned to each adapter. For example, sybase_ase_out or sybase_iq_out. See the Adapters Guide for more information.</p> <p>In/Out - displays whether the adapter is an input or output adapter.</p> <p>Group - the adapter group to which the adapter belongs.</p>

These statistics display for the adapter you select in the header table:

Table 38. Common Statistics

Name	Value
AdapterTotalRows	Total number of rows in the adapter.
AdapterGoodRows	Number of good rows in the adapter.
AdapterBadRows	Number of bad rows in the adapter.
AdapterLatency	Time taken for data to be processed.

Table 39. Extended Statistics

Name	Value
<Custom Statistic Name>	If the adapter you select has custom statistics to report, these display here. Note that not all adapters have custom statistics.

See also

- *Viewing Statistics for Adapters on a Node* on page 176
- *Viewing Adapters* on page 198
- *Starting an Adapter* on page 199
- *Stopping an Adapter* on page 200
- *Viewing File Activity for the Sybase IQ Output Adapter* on page 201

Viewing Statistics for Publishers on a Cluster

Monitor the performance of publishers on an ESP cluster by viewing statistics such as CPU history and number of rows sent.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Publishers**.

Table Name	Publisher Details
Publishers	<p>This is the header table containing all publishers within the ESP cluster that the selected node belongs to.</p> <p>Node - name of the ESP node on which the workspace is running.</p> <p>Workspace - name of workspace to which the project belongs.</p> <p>Project - name of the project to which the publisher is connected.</p> <p>IP Address - IP address of the publisher.</p> <p>Port - port number on the publisher's machine used to send subscribed data to Event Stream Processor.</p>

These statistics display for the publisher you select from the header table:

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The user CPU usage and system CPU usage are also shown on this line graph. The data on the graph starts displaying from the time you open the Cluster Monitor - Publishers view.
Rows Transferred History	Line graph displaying rows sent per second. The data on the graph starts displaying from the time you open the Cluster Monitor - Publishers view.

See also

- *Viewing Statistics for Publishers on a Node* on page 177

Viewing Statistics for Subscribers on a Cluster

Monitor the performance of subscribers on an ESP cluster by viewing statistics such as CPU history and number of rows received. SCC also uses subscribers to gather information used for monitoring. These SCC subscribers also display within this view.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Subscribers**.

Table Name	Subscriber Details
Subscribers	<p>This is the header table containing all subscribers within the ESP cluster that the selected node belongs to.</p> <p>Node - name of the ESP node on which the workspace is running.</p> <p>Workspace - name of the workspace to which the subscriber is connected.</p> <p>Project - name of the project to which the subscriber is connected.</p> <p>IP Address - IP address of the subscriber.</p> <p>Port - port number on the subscriber's machine used to receive published data from Event Stream Processor.</p> <p>Stream Name - name of the data stream to which the subscriber is subscribed.</p> <p>Queue Depth - number of rows waiting to be received by the subscriber.</p> <p>Rows in Store - number of rows in the subscriber's data store.</p>

These statistics display for the subscriber you select from the header table:

Statistics	Description
CPU History	Line graph displaying the percentage of total CPU usage over time. The user CPU usage and system CPU usage are also shown on this line graph. The data on the graph starts displaying from the time you open the Cluster Monitor - Subscribers view.
Rows Transferred History	Line graph displaying rows received per second. The data on the graph starts displaying from the time you open the Cluster Monitor - Subscribers view.

See also

- *Viewing Statistics for Subscribers on a Node* on page 178

Viewing Schema for a Stream

View schema for a selected stream from the ESP Cluster Monitor view.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Log in using the espMonitorRole role to be able to perform this task.

Task

1. In the Perspective Resources window, select the node, click the arrow, and select **Monitor Cluster**.
2. In the left pane of the ESP Cluster Monitor view, select **Streams**.
3. In the right pane of the ESP Cluster Monitor view, select a stream, click the arrow, and select **Show Schema**.
A window appears displaying the schema for the selected stream (for example, column name, type, and primary key).

Projects

You can view, start, stop, add, remove projects, and view project log files from the Administration Console. Log in with the espAdminRole to be able to perform these actions.

See also

- *Configuring Event Stream Processor for Administration* on page 127
- *Starting a Node* on page 179
- *Stopping a Node* on page 180

Viewing Projects

Use the Administration Console to display a list of the projects associated with all registered ESP nodes or only a certain node.

Prerequisites

- Log in using the espAdminRole role to be able to perform this task.
- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

- Register and authenticate the SCC agent for the node.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Projects**. All the projects across the registered ESP nodes display on the right pane of the Administration Console. This info displays for each project:

Table Name	Statistics
Projects	<p>Node - name of the ESP node on which the workspace is running.</p> <p>Workspace Name - name of workspace to which the project belongs.</p> <p>Node Name - as entered when the resource was registered.</p> <p>Status - project status. Valid values are running, stopped, or unknown.</p>

See also

- *Viewing Statistics for Projects on a Node* on page 172
- *Viewing Statistics for Projects on a Cluster* on page 186
- *Starting a Project* on page 195
- *Stopping a Project* on page 196

Starting a Project

Start a project from the Administration Console.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

The steps below enable you to start a project which belongs to a specific node. To view and manage all projects belonging to a cluster rather than only a node, use the **Manage Projects...**

Manage and Monitor Event Stream Processor

command from a manager node in the Administration Console. For detailed steps, see *Managing Projects*.

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Projects**.
3. In the right pane of the Administration Console, select a project that is not currently running, click the arrow, and select **Start Project...**
4. A message displays asking you to confirm whether you want to start the project. Click **Yes** to confirm. The project status appears as "running".

See also

- *Viewing Statistics for Projects on a Node* on page 172
- *Viewing Statistics for Projects on a Cluster* on page 186
- *Viewing Projects* on page 194
- *Stopping a Project* on page 196
- *Starting an Adapter* on page 199
- *Starting a Node* on page 179

Stopping a Project

Stop a project from the Administration Console.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

The steps below enable you to stop a project running on a specific node. To view and manage all projects belonging to a cluster rather than only a node, use the **Manage Projects...** command from a manager node in the Administration Console. For detailed steps, see *Managing Projects*.

1. You can view all registered nodes or only a certain node from the Administration Console:

- To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Projects**.
 3. In the right pane of the Administration Console, select a project that is currently running, click the arrow, and select **Stop Project...**
 4. A message displays asking you to confirm whether you want to stop the project. Click **Yes** to confirm. The project status is marked by a red color.

See also

- *Viewing Statistics for Projects on a Node* on page 172
- *Viewing Statistics for Projects on a Cluster* on page 186
- *Viewing Projects* on page 194
- *Starting a Project* on page 195
- *Stopping a Node* on page 180
- *Stopping an Adapter* on page 200

Managing Projects

From a manager node, you can start, stop, add, or remove any project within the cluster to which the manager node belongs, even if the project is not running on the selected manager node. You can also start or stop a project from the Project subsection of the Administration Console. The difference with managing projects from the manager node is that the manager node allows you to see and manage all projects within the cluster rather than only the projects running on the selected node.

Prerequisites

- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.

Manage and Monitor Event Stream Processor

2. In the right pane of the Administration Console, select a manager node, click the arrow, and select **Manage Projects...**
3. A new window displays and contains detailed information about the projects running on the cluster to which the manager node belongs. You can:
 - Click **Add Project** to add a new project to the cluster. A new window displays. Specify a workspace, project name, project .ccx file name, and project .ccr file name for the new project. Click **Add** to add the new project to the cluster.
 - Select a project that is currently running and click **Remove Project** to remove that project from the cluster.
 - Select a project that is not currently running and click **Start Project**.
 - Select a project that is currently running and click **Stop Project**.
 - Click **Cancel** to exit the window.

Adapters

You can view, start and stop adapters, and view file activity for the Sybase IQ Output adapter from the Administration Console. Log in with the espAdminRole to be able to perform these actions.

See also

- *Configuring Event Stream Processor for Administration* on page 127
- *Starting a Node* on page 179
- *Stopping a Node* on page 180

Viewing Adapters

Use the Administration Console to display a list of the adapters associated with all registered ESP nodes or only a certain node.

Prerequisites

- Log in using the espAdminRole role to be able to perform this task.
- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.

- To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Adapters**.

All the adapters across the registered ESP nodes display on the right pane of the Administration Console. This info displays for each adapter:

 - **Name** - name of the adapter.
 - **Workspace Name** - name of workspace on which the project is running.
 - **Project Name** - name of project which is attached to the adapter.
 - **Node Name** - name of node on which the project is running.
 - **Type** - the unique adapter ID assigned to each adapter. For example, sybase_ase_out or sybase_iq_out. See the Adapters Guide for more information.
 - **In/Out** - whether the adapter is an input or output adapter.
 - **Status** - adapter status. Valid values are initial, stopped, ready, continuous, idle, done, or unknown.

See also

- *Viewing Statistics for Adapters on a Node* on page 176
- *Starting an Adapter* on page 199
- *Stopping an Adapter* on page 200
- *Viewing File Activity for the Sybase IQ Output Adapter* on page 201
- *Viewing Statistics for Adapters on a Cluster* on page 190

Starting an Adapter

Start an adapter from the Administration Console.

Prerequisites

- Log in using the espAdminRole role to be able to perform this task.
- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.

Manage and Monitor Event Stream Processor

2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Adapters**.
3. In the right pane of the Administration Console, select an adapter that is not currently running, click the arrow, and select **Start Adapter...**
4. A window appears and asks you to confirm whether you want to start the adapter. Click **Yes** to confirm.
A window appears and reports whether the adapter was successfully started.

See also

- *Viewing Statistics for Adapters on a Node* on page 176
- *Viewing Adapters* on page 198
- *Stopping an Adapter* on page 200
- *Viewing File Activity for the Sybase IQ Output Adapter* on page 201
- *Viewing Statistics for Adapters on a Cluster* on page 190
- *Starting a Project* on page 195
- *Starting a Node* on page 179

Stopping an Adapter

Stop an adapter from the Administration Console.

Prerequisites

- Log in using the espAdminRole role to be able to perform this task.
- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Adapters**.
3. In the right pane of the Administration Console, select an adapter that is currently running, click the arrow, and select **Stop Adapter...**
4. A window appears and asks you to confirm whether you want to stop the adapter. Click **Yes** to confirm.

A window appears and reports whether the adapter was successfully stopped.

See also

- *Viewing Statistics for Adapters on a Node* on page 176
- *Viewing Adapters* on page 198
- *Starting an Adapter* on page 199
- *Viewing File Activity for the Sybase IQ Output Adapter* on page 201
- *Viewing Statistics for Adapters on a Cluster* on page 190
- *Stopping a Node* on page 180
- *Stopping a Project* on page 196

Viewing File Activity for the Sybase IQ Output Adapter

View the file activity report for the Sybase IQ Output Adapter from the Administration Console to see each file processed by the adapter and its current state.

Prerequisites

- (Optional) Register the Sybase IQ server, to which the Sybase IQ Output adapter publishes data, with the same instance of SCC that you registered the node on.
- Create a database user and table in the database into which the Sybase IQ Output adapter is loading data. Run the `$ESP_HOME/adapters/iqoutput/enableFileActivity.sql` script on your Sybase IQ database as a user with permissions to create a user, and create a table for that user. For full information on the table and user you need to create, see *Enabling File Activity Monitoring for the Sybase IQ Adapter* in the *Event Stream Processor Adapters Guide*.
- Log in using the `espAdminRole` role to be able to perform this task.
- Enable the time-granularity option on the node in the Event Stream Processor project configuration (.ccr) file. See the *Event Stream Processor Administrators Guide* for information on the project configuration file.

Task

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Adapters**.
3. In the right pane of the Administration Console, select the Sybase IQ Output adapter, click the arrow, and select **File Activity**.

Manage and Monitor Event Stream Processor

4. A new window appears and displays the file activity report which contains details about each of the files processed by the adapter, including their current state.

By default, it displays records an hour ahead from the local time. If you want to display records from different periods of time, go to the Perspective Resources view menu, and select **View > Filter**.

If the Sybase IQ database cannot find the Sybase IQ database username and password from the SCC server, another window appears and asks for this information.

See also

- *Viewing Statistics for Adapters on a Node* on page 176
- *Viewing Adapters* on page 198
- *Starting an Adapter* on page 199
- *Stopping an Adapter* on page 200
- *Viewing Statistics for Adapters on a Cluster* on page 190

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 because it displays only selected resources.
- 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. To explore the hierarchy of object types, select **Navigation > Browse** in the left pane. Expand an object type by clicking its arrow icon.
3. Select an object type (any server type, for example) in the hierarchy. In the right pane, the Administration Console displays a list of resources of that type.

Note: Message rows in the right pane are placeholders for:

- Failed requests – to retry, select the message row and click the drop-down arrow that appears to the right. Select **Retry**.
- Slow-responding requests – SCC replaces these rows with real data as soon as it arrives.
- Large result sets – to display, select the message row and click the drop-down arrow that appears to the right. Select **Expand**. The results might take a minute to appear.

Hover the mouse over a message row to see a tooltip with more information.

4. (Optional) To create an object of the type now selected, click **Folder > Create** or **Folder > New**.
5. (Optional) To refresh the view, select **Folder > Refresh**.
6. 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.
7. 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.

See also

- *Searching and Filtering Resources* on page 203

Searching and Filtering Resources

Use the Administration Console's search and filter tools to quickly find the resources or objects within resources that interest you.

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 because it displays only selected resources.
 - 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. (Optional) You can use the Administration Console's tools to control which resources it displays:

Manage and Monitor Event Stream Processor

- a) In the left pane, click **Resource Selection**.
 - b) SCC refreshes the list of resources in the right pane with each selection you make in this pane. If you are making multiple selection changes, unselect **Automatically refresh details** to turn the refresh feature off.
 - c) Select or unselect resources to include or eliminate them from Administration Console displays.
3. To find a resource without navigating the hierarchy:
- a) In the left pane, select **Navigation > Search**.
 - b) (Required) On the Search tab, select the resource type and object type of the resource you want to find.
 - c) Enter a search string.
The search string can be the full or partial name of a resource.
 - d) (Optional) Select **Exact match** to find only the resource whose name is identical to the search string.
 - e) Click **Search**.
Results appear in the right pane.

Note: Message rows in the right pane are placeholders for:

- Failed requests – to retry, select the message row and click the drop-down arrow that appears to the right. Select **Retry**.
 - Slow-responding requests – SCC replaces these rows with real data as soon as it arrives.
 - Large result sets – to display, select the message row and click the drop-down arrow that appears to the right. Select **Expand**. The results might take a minute to appear.
Hover the mouse over a message row to see a tooltip with more information.
- f) To further narrow your search, enter a filter string in the field at the top of any column of search results. For example, in a search for databases, enter `wilma` above the Device column to display only results associated with the device `wilma`.

See also

- *Browsing and Managing Resources* on page 202
- *Configuring Retrieval Thresholds for the Administration Console* on page 113

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 141

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 205
- *Resuming and Suspending a Data Collection Job* on page 206
- *Adding a New Schedule to a Job* on page 207
- *Viewing or Deleting a Schedule* on page 208
- *Modifying the Data Collection Interval for a Job* on page 208
- *Resuming and Suspending the Scheduler* on page 209
- *Viewing the Job Execution History* on page 209

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 205
- *Resuming and Suspending a Data Collection Job* on page 206
- *Adding a New Schedule to a Job* on page 207
- *Viewing or Deleting a Schedule* on page 208
- *Modifying the Data Collection Interval for a Job* on page 208
- *Resuming and Suspending the Scheduler* on page 209
- *Viewing the Job Execution History* on page 209

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 205
- *Deleting a Data Collection Job* on page 205
- *Adding a New Schedule to a Job* on page 207
- *Viewing or Deleting a Schedule* on page 208
- *Modifying the Data Collection Interval for a Job* on page 208
- *Resuming and Suspending the Scheduler* on page 209
- *Viewing the Job Execution History* on page 209

Adding a New Schedule to a Job

Use the Properties view for a resource to add schedules to 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.
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**. If you choose **Later**, specify the start date and time.
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.

You cannot change the duration of a schedule (the once/repetitively/until setting) after you create it. To change the schedule duration, delete and recreate the schedule.

8. Click **Finish** to save the schedule.
9. Click **OK**.

See also

- *Executing and Stopping a Data Collection Job* on page 205
- *Deleting a Data Collection Job* on page 205

- *Resuming and Suspending a Data Collection Job* on page 206
- *Viewing or Deleting a Schedule* on page 208
- *Modifying the Data Collection Interval for a Job* on page 208
- *Resuming and Suspending the Scheduler* on page 209
- *Viewing the Job Execution History* on page 209

Viewing or Deleting a Schedule

Display schedule details or remove a schedule from 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. To display the schedules for a collection job, expand the job by clicking the arrow to the left of the job's name.
If there is no arrow to the left of the job's name, this job has no schedules.
4. Select a schedule.
The name, description, start and end dates, and repeat interval appear on the Schedule tab.
5. (Optional) To remove the selected schedule, click **Delete**.
6. Click **OK**.

See also

- *Executing and Stopping a Data Collection Job* on page 205
- *Deleting a Data Collection Job* on page 205
- *Resuming and Suspending a Data Collection Job* on page 206
- *Adding a New Schedule to a Job* on page 207
- *Modifying the Data Collection Interval for a Job* on page 208
- *Resuming and Suspending the Scheduler* on page 209
- *Viewing the Job Execution History* on page 209
- *Setting Up Statistics Collection* on page 141

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 205
- *Deleting a Data Collection Job* on page 205
- *Resuming and Suspending a Data Collection Job* on page 206
- *Adding a New Schedule to a Job* on page 207
- *Viewing or Deleting a Schedule* on page 208
- *Resuming and Suspending the Scheduler* on page 209
- *Viewing the Job Execution History* on page 209

Resuming and Suspending the Scheduler

Use the scheduler settings to resume or suspend all scheduled jobs.

Prerequisites

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

Task

1. From the main menu bar, select **Application > Administration**.
2. In the Sybase Control Center Properties dialog, select **Scheduler**.
3. Do one of the following:
 - To resume the scheduler, click **Resume**.
 - To suspend the scheduler, click **Suspend**.
4. Click **OK**.

See also

- *Executing and Stopping a Data Collection Job* on page 205
- *Deleting a Data Collection Job* on page 205
- *Resuming and Suspending a Data Collection Job* on page 206
- *Adding a New Schedule to a Job* on page 207
- *Viewing or Deleting a Schedule* on page 208
- *Modifying the Data Collection Interval for a Job* on page 208
- *Viewing the Job Execution History* on page 209

Viewing the Job Execution History

Use the Properties view to display a data collection job's execution history.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties**.
2. Select **Collection Jobs**.

3. Select a job.
4. Click the **History** tab.

See also

- *Executing and Stopping a Data Collection Job* on page 205
- *Deleting a Data Collection Job* on page 205
- *Resuming and Suspending a Data Collection Job* on page 206
- *Adding a New Schedule to a Job* on page 207
- *Viewing or Deleting a Schedule* on page 208
- *Modifying the Data Collection Interval for a Job* on page 208
- *Resuming and Suspending the Scheduler* on page 209

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 159
- *Creating an Alert* on page 151
- *Assigning a Role to a Login or a Group* on page 115
- *Configuring the E-mail Server* on page 111
- *Event Stream Processor Alerts* on page 153
- *Alert Types and Severities for Event Stream Processor* on page 158
- *Viewing Overview Statistics and Alerts for a Node* on page 169
- *Viewing Overview Statistics and Alerts for a Cluster* on page 182

Types, Severities, and States

Learn about the properties that define and control alerts.

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

Table 40. 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 41. 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 212
- *Modifying an Alert* on page 212
- *Testing an Alert-Triggered Script* on page 213

Manage and Monitor Event Stream Processor

- *Deleting an Alert* on page 214
- *Alert Subscriptions* on page 214
- *Alert Notifications* on page 217
- *Creating an Alert* on page 151

Viewing Alerts

Display alert notifications and alerts that have been configured for a given resource.

- To display generated alerts (notifications):
 - a) Select **View > Open > Alert Monitor** from the application menu bar.
For a given alert, the Alert Monitor displays only the most recent unresolved notifications at each severity level. That is, if an alert fires five times at the warning level, only the notification of the fifth firing is listed—even if the previous four alerts remain unresolved.
 - b) To display information about a generated alert, select the alert in the Alert Monitor and click **Properties**.
- To display configured alerts:
 - a) In the Perspective Resources view, select a resource and select **Resource > Properties**.
 - b) Click **Alerts** to view configured alerts for the selected resource.
(This is a different route to the information displayed in the second step, above.)

See also

- *Types, Severities, and States* on page 211
- *Modifying an Alert* on page 212
- *Testing an Alert-Triggered Script* on page 213
- *Deleting an Alert* on page 214
- *Alert Subscriptions* on page 214
- *Alert Notifications* on page 217
- *Creating an Alert* on page 151

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

- *Types, Severities, and States* on page 211
- *Viewing Alerts* on page 212
- *Testing an Alert-Triggered Script* on page 213
- *Deleting an Alert* on page 214
- *Alert Subscriptions* on page 214
- *Alert Notifications* on page 217
- *Creating an Alert* on page 151

Testing an Alert-Triggered Script

Execute a script to make sure it works properly.

Prerequisites

Configure an alert with a script.

Task

1. In the Perspective Resources view, select a resource and select **Resource > Properties**.
2. Select **Alerts**.
3. Select the alert to test.
4. On the Script tab, click **Modify**.
5. If the script requires parameter values, click **Select Parameters** to enter them in the **Execution Parameters** box.

You can include a number of predefined substitution parameters, which are replaced by values from the alert. The parameter values are passed on the command line to the script. For the test execution, use values that test all the parameters used by the script. See the substitution parameters topic (linked below) for more information.

Note: When you test a script, Sybase Control Center supplies test values for the **%Severity %** and **%Source_Application%** parameters (“Testing” and “TestScriptExecution,” respectively). Any test values you supply for these parameters are discarded. This prevents

the test results from being confused with real script results after testing and in the SCC repository.

6. Click **Test** to perform a test execution of your script.

If your script takes parameters, the test may fail if parameter values are missing or incorrect.

See also

- *Types, Severities, and States* on page 211
- *Viewing Alerts* on page 212
- *Modifying an Alert* on page 212
- *Deleting an Alert* on page 214
- *Alert Subscriptions* on page 214
- *Alert Notifications* on page 217
- *Alert-Triggered Scripts* on page 159
- *Substitution Parameters for Scripts* on page 161
- *Creating an Alert* on page 151

Deleting an Alert

Use the Properties view of your resource to delete an alert.

1. In the Perspective Resources view, select a resource and select **Resource > Properties**.
2. Select **Alerts**.
3. Select an alert and click **Drop**.
4. Click **Yes** to confirm the deletion.

See also

- *Types, Severities, and States* on page 211
- *Viewing Alerts* on page 212
- *Modifying an Alert* on page 212
- *Testing an Alert-Triggered Script* on page 213
- *Alert Subscriptions* on page 214
- *Alert Notifications* on page 217
- *Creating an Alert* on page 151

Alert Subscriptions

When an alert subscription is configured, the alert notifies the specified user or group of users by e-mail message when the alert fires.

You can configure an alert subscription to send e-mail notifications when the alert reaches a severity of warning, a severity of critical, or both.

You can also configure an alert subscription to escalate after a period of time that you specify. If the alert is not resolved within the escalation period, Sybase Control Center e-mails an escalation message to the user or group whose address you provide for escalations, as well as to the primary subscriber. The escalation message is identical to the primary notification message. Sybase recommends that if you configure alert subscriptions to escalate, you do so only for the most urgent alerts, those with a severity of critical.

See also

- *Types, Severities, and States* on page 211
- *Viewing Alerts* on page 212
- *Modifying an Alert* on page 212
- *Testing an Alert-Triggered Script* on page 213
- *Deleting an Alert* on page 214
- *Alert Notifications* on page 217
- *Creating an Alert* on page 151

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 42. 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 216
- *Enabling and Disabling Alert Subscription* on page 216

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 215
- *Enabling and Disabling Alert Subscription* on page 216

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 215
- *Unsubscribing from an Alert* on page 216

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 211
- *Viewing Alerts* on page 212
- *Modifying an Alert* on page 212
- *Testing an Alert-Triggered Script* on page 213
- *Deleting an Alert* on page 214
- *Alert Subscriptions* on page 214
- *Creating an Alert* on page 151

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.

Manage and Monitor Event Stream Processor

4. Click the **History** tab.
5. (Optional) Click the **Resolutions** tab.

See also

- *Resolving Alerts* on page 218

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 217

Log Files for Event Stream Processor

Event Stream Processor has four log files which record messages about system and component events.

Event Stream Processor maintains these logs:

- The Sybase Control Center for Event Stream Processor (ESPMAP) log - captures messages about startup errors, tracks verification processes, and can help you diagnose connectivity issues.
- The Sybase Control Center agent for Event Stream Processor (ESPAP) log - captures messages about SCC agent activities and can help you diagnose issues.

- The node log - captures messages about the cluster manager and controller activities and can help you diagnose issues.
- The project log - captures messages about the project activities and can help you diagnose issues.

See also

- *Logging* on page 245
- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Modifying the Event Stream Processor Log Configuration* on page 246

Viewing the SCC Agent for Event Stream Processor Log File

View the log file for the SCC agent for Event Stream Processor from the Administration Console.

Prerequisites

- Register and authenticate an ESP node.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

The SCC agent log records messages about SCC agent activities and can help you diagnose issues.

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select **ESP Nodes**.
3. In the right pane of the Administration Console, select a node, click the arrow, and select **View Agent Log**.

A new window appears and displays the log file entries:

Column Name	Description
Date	The date and time at which the event occurred and the message was logged.

Column Name	Description
Severity	Severity levels span from zero to five and correspond to: <ul style="list-style-type: none"> • FATAL - (0) Processing cannot continue. • ERROR - (1) Processing can continue but a serious issue is encountered. May require corrective action on your part. • WARNING - (2) Bringing your attention to an unusual situation or event. For example, you may see a warning when you have only 20 percent of memory left. • INFO - (3) Information about system activity. No action is required. • DEBUG - (4) Provides details about the code to assist you in resolving issues. This may have a negative impact on performance. • TRACE - (5) Similar to DEBUG but with more details.
Message	Displays the log entries. Each log entry displays descriptive text according to the severity level to which it corresponds.

4. (Optional) Select **Settings** from the top right corner of the window to select what lines you want to see from the log file, and select **Apply Settings**.

You can choose to retrieve all the lines from the file, the last n lines, lines from the past n days, or lines from certain dates.

Viewing the Node Log File

View the log file for an ESP node from the Administration Console.

Prerequisites

- Register and authenticate an ESP node.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

The node log file captures messages generated about the cluster manager and controller activities and can help you diagnose issues.

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, select **ESP Nodes**.

- In the right pane of the Administration Console, select a node, click the arrow, and select **View Node Log**.

A new window appears and displays the log file entries:

Column Name	Description
Date	The date and time at which the event occurred and the message was logged.
Severity	Severity levels span from zero to five and correspond to: <ul style="list-style-type: none"> FATAL - (0) Processing cannot continue. ERROR - (1) Processing can continue but a serious issue is encountered. May require corrective action on your part. WARNING - (2) Bringing your attention to an unusual situation or event. For example, you may see a warning when you have only 20 percent of memory left. INFO - (3) Information about system activity. No action is required. DEBUG - (4) Provides details about the code to assist you in resolving issues. This may have a negative impact on performance. TRACE - (5) Similar to DEBUG but with more details.
Message	Displays the log entries. Each log entry displays descriptive text according to the severity level to which it corresponds.

- (Optional) Select **Settings** from the top right corner of the window to select what lines you want to see from the log file, and select **Apply Settings**.

You can choose to retrieve all the lines from the file, the last n lines, lines from the past n days, or lines from certain dates.

See also

- Registering and Authenticating a Sybase Control Center Agent* on page 136

Viewing the Project Log File

View the log file for a project from the Administration Console.

Prerequisites

- Register and authenticate an ESP node.
- Register and authenticate the SCC agent for the node.
- Log in using the espAdminRole role to be able to perform this task.

Task

The project log file records messages about project activity and can help you diagnose issues.

Manage and Monitor Event Stream Processor

1. You can view all registered nodes or only a certain node from the Administration Console:
 - To view all registered nodes, in the application menu, select **View > Open > Administration Console**.
 - To view only a certain node, in the Perspective Resources view, select the node and select **Resource > Administration Console**.
2. In the left pane of the Administration Console, expand **ESP Nodes** and select **Projects**.
3. In the right pane of the Administration Console, select a project, click the arrow, and select **View Project Log**.

A new window appears and displays the log file entries:

Column Name	Description
Date	The date and time at which the event occurred and the message was logged.
Severity	Severity levels span from zero to seven, and display the scale of importance. Zero is the highest severity, and seven is the lowest severity. For more information on severity levels, see the <i>Logging</i> topic in the <i>ESP Administrators Guide</i> .
Message	Displays the log entries. Each log entry displays descriptive text according to the severity level to which it corresponds.

4. (Optional) Select **Settings** from the top right corner of the window to select what lines you want to see from the log file, and select **Apply Settings**.

You can choose to retrieve all the lines from the file, the last n lines, lines from the past n days, or lines from certain dates.

See also

- *Registering and Authenticating a Sybase Control Center Agent* on page 136

Manage Sybase Control Center

All Sybase products using Sybase Control Center share high-level management features, including the Administration Console, data collection jobs, alerts, logs, resources, perspectives, and views.

Use these management features to perform high-level management tasks. These tasks apply to any Sybase product you manage with Sybase Control Center.

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

- *Importing Resources for Batch Registration* on page 129

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 224
- *Removing a Resource from a Perspective* on page 224
- *Modifying a Resource's Name and Connection Properties* on page 225
- *Searching for Resources in the Resource Explorer* on page 226
- *Importing Resources for Batch Registration* on page 129

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 223
- *Removing a Resource from a Perspective* on page 224
- *Modifying a Resource's Name and Connection Properties* on page 225
- *Searching for Resources in the Resource Explorer* on page 226
- *Importing Resources for Batch Registration* on page 129

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. Before removing a resource, make sure it is not in use by an open view.

- Close any views that display the resource.
- If you prefer not to close the Administration Console, unselect the resource:
 - a) In the left pane of the Administration Console, click **Resource Selection**.
 - b) Locate the resource in the list and click the box to unselect it.
- 2. If the Perspective Resources view is not open, click the **Show/Hide Perspective Resources View** icon in the perspective toolbar.
- 3. In the Perspective Resources view, select the resources to remove. Use **Shift-click** or **Control-click** to select multiple resources.
- 4. Select **Resource > Remove**.
- 5. Click **Yes** to confirm the removal.

See also

- *Unregistering a Resource* on page 223
- *Adding a Resource to a Perspective* on page 224
- *Modifying a Resource's Name and Connection Properties* on page 225
- *Searching for Resources in the Resource Explorer* on page 226
- *Adding a Resource to a Perspective* on page 139

Modifying a Resource's Name and Connection Properties

Change the properties of a resource registered with Sybase Control Center.

1. In the Perspective Resources view, select a resource and select **Resource > Properties**.
2. (Optional) On the General Properties page, modify the name or description of the resource.
Enter the actual name of the managed 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.
3. (Optional) On the Connection Information page, modify:
 - the host name
 - the port number
 - other options for the managed resource
4. Click **OK** (to apply the changes and close the properties dialog) or **Apply** (to apply the changes and leave the dialog open).

See also

- *Unregistering a Resource* on page 223
- *Adding a Resource to a Perspective* on page 224
- *Removing a Resource from a Perspective* on page 224
- *Searching for Resources in the Resource Explorer* on page 226

- *Registering an ESP Node* on page 128
- *Importing Resources for Batch Registration* on page 129

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 223
- *Adding a Resource to a Perspective* on page 224
- *Removing a Resource from a Perspective* on page 224
- *Modifying a Resource's Name and Connection Properties* on page 225

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.

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 227
- *Renaming a Perspective* on page 227

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 227
- *Renaming a Perspective* on page 227

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 227
- *Removing a Perspective* on page 227

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.

Note: Sybase Control Center views are not related to database views; they serve a completely different purpose.

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: <ul style="list-style-type: none"> • In the Perspective Resources view, select a resource, click the drop-down arrow to the right of the resource name, and select the view to open. • In the application menu bar, select View > Open and choose a view.
Close a view	Select the view to close. In the application menu bar, select View > Close . You can also click the X in the view's upper right corner.
Maximize a view	Click the box in the view's upper right corner. The view enlarges to fill the entire perspective window. Click the box again to return the view to its former size.
Minimize a view	Click the _ in the view's upper right corner. The view shrinks to a small tab at the bottom of the perspective window.

Task	Action
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 229

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 228

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 reenables, 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 231
- *Refreshing or Converting an Instance* on page 232
- *Removing an Instance* on page 233
- *Shared-Disk Mode* on page 234
- *sccinstance Command* on page 235

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 230
- *Refreshing or Converting an Instance* on page 232
- *Removing an Instance* on page 233
- *Shared-Disk Mode* on page 234
- *sccinstance Command* on page 235

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 230
- *Deploying an Instance from a Shared Disk Installation* on page 231
- *Removing an Instance* on page 233
- *Shared-Disk Mode* on page 234
- *sccinstance Command* on page 235

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 230
- *Deploying an Instance from a Shared Disk Installation* on page 231
- *Refreshing or Converting an Instance* on page 232
- *Shared-Disk Mode* on page 234
- *sccinstance Command* on page 235

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, SAP Sybase ESP clusters, or SAP 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 230
- *Deploying an Instance from a Shared Disk Installation* on page 231
- *Refreshing or Converting an Instance* on page 232
- *Removing an Instance* on page 233
- *sccinstance Command* on page 235

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	ScsSADataserver 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
tds	Tabular Data Stream™ port (used to communi- cate 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 230
- *Deploying an Instance from a Shared Disk Installation* on page 231

- *Refreshing or Converting an Instance* on page 232
- *Removing an Instance* on page 233
- *Shared-Disk Mode* on page 234

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-3_2/backup`.
- Each incremental backup is stored in a file in `SCC-3_2/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 241
- *Forcing an Immediate Backup* on page 242
- *Restoring the Repository from Backups* on page 243
- *Configuring Repository Purging* on page 244

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**.

Manage Sybase Control Center

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 240
- *Forcing an Immediate Backup* on page 242
- *Restoring the Repository from Backups* on page 243
- *Configuring Repository Purging* on page 244

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 240
- *Modifying the Backup Schedule* on page 241
- *Restoring the Repository from Backups* on page 243
- *Configuring Repository Purging* on page 244

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-3_2/backup/
<generated_directory_name> to SCC-3_2/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 /opt/sybase/SCC-3_2/backup/repository.270110161105/  
scc_repository.db  
/opt/sybase/SCC-3_2/services/Repository
```

3. If you have no incremental backups to load,
 - a) Also copy the log file from SCC-3_2/backup/
<generated_directory_name> to SCC-3_2/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 /opt/sybase/SCC-3_2/backup/repository.270110161105/  
scc_repository.log  
/opt/sybase/SCC-3_2/services/Repository
```

- b) Skip to step 5 on page 244.
4. (Optional) To load incremental backups, 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 sybase\SCC-3_2\services\Repository  
  
..\..\bin\sa\bin_<platform>\dbsrv11.exe scc_repository -ad  
sybase\SCC-3_2\backup\incremental
```

UNIX:

Manage Sybase Control Center

```
cd /opt/sybase/SCC-3_2/services/Repository
../../bin/sa/bin_<platform>/dbsrv11 scc_repository -ad
/opt/sybase/SCC-3_2/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, Sybase Control Center starts normally (that is, no further recovery occurs). If you copied a full backup to the `Repository` directory, the database recovers the repository from the full backup.

Example: Loading incremental backups into the repository database

These commands start SQL Anywhere® on a 32-bit Windows machine:

```
% cd C:\sybase\SCC-3_2\services\Repository
% ../../bin\sa\bin_windows32\dbsrv11.exe scc_repository -ad
C:\sybase\SCC-3_2\backup\incremental
```

These commands start SQL Anywhere on a 64-bit machine running Solaris:

```
$ cd /opt/sybase/SCC-3_2/services/Repository
$ ../../bin/sa/bin_sunsparc64/dbsrv11 scc_repository -ad
/opt/sybase/SCC-3_2/backup/incremental
```

See also

- *Scheduling Backups of the Repository* on page 240
- *Modifying the Backup Schedule* on page 241
- *Forcing an Immediate Backup* on page 242
- *Configuring Repository Purging* on page 244

Configuring Repository Purging

Change repository purging options.

Prerequisites

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

Task

As you decide how to purge your repository, consider that:

- Purging keeps the repository from absorbing too much disk space.
- By default, purging is enabled. It occurs once a day and purges data older than one day.
- Statistics and alert history can help you detect trends in server performance and user behavior. The Sybase Control Center statistics chart can graph performance data over a

period of a year or more if the data is available. If you have enough disk space, consider saving data for a longer period of time or disabling the purging of statistics or alert history.

- Changing the purge frequency and other options might affect Sybase Control Center performance.

Note: If you configure purging as part of a scheduled backup of the repository, disable automatic purging on the Size Management tab.

1. From the main menu bar, select **Application > Administration**.
2. Select **Repository**.
3. Click the **Size Management** tab.
4. To turn automatic purging on or off, click **Automatically purge the repository periodically**.

Turn this option off if purging is configured as part of your scheduled full or incremental backups.

5. Click purge options to turn them on or off:
 - **Purge statistics**
 - **Purge alert history**
6. In **Purge data older than**, enter the number of days after which to purge repository data.
7. Click **Apply**, then **OK**.

See also

- *Scheduling Backups of the Repository* on page 240
- *Modifying the Backup Schedule* on page 241
- *Forcing an Immediate Backup* on page 242
- *Restoring the Repository from Backups* on page 243

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:

Manage Sybase Control Center

- 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`.

See also

- *Log Files for Event Stream Processor* on page 218

Viewing the Sybase Control Center for Event Stream Processor Log

View event logs for Sybase Control Center for Sybase Event Stream Processor.

Sybase Control Center for Sybase Event Stream Processor uses Log4J for message logging. The Sybase Control Center for Sybase Event Stream Processor log files are located at:

- Windows – `%SYBASE%\SCC-3_2\plugins\ESPMAP\log\espmap.log`
- UNIX – `$SYBASE/SCC-3_2/plugins/ESPMAP/log/espmap.log`

1. Display the log file using a log viewer or another method of your choice.
2. Look for entries of interest such as login attempts or the failure of a service to start.

See also

- *Modifying the Event Stream Processor Log Configuration* on page 246
- *Viewing Sybase Control Center Server Logs* on page 247
- *Viewing the Sybase Control Center Client Log* on page 248
- *Changing the Logging Level* on page 248
- *Logging or Message Levels* on page 249
- *Changing Logging Configuration* on page 250
- *Troubleshooting Tips* on page 259
- *Log Files for Event Stream Processor* on page 218

Modifying the Event Stream Processor Log Configuration

Change the log level or logging configuration settings for Sybase Control Center for Event Stream Processor.

1. Navigate to `%SYBASE%\SCC-3_2\plugins\ESPMAP\agent-plugin.xml`.

2. Open the `agent-plugin.xml` file, and modify the settings as needed.
3. Save and close the `agent-plugin.xml` file.
4. Restart the SCC server.

Messages related to SCC for Event Stream Processor are recorded on the console and the `esmap.log` file. The `esmap.log` file is located in `%SYBASE%\SCC-3_2\plugins\ESPMAP\log`.

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Viewing Sybase Control Center Server Logs* on page 247
- *Viewing the Sybase Control Center Client Log* on page 248
- *Changing the Logging Level* on page 248
- *Logging or Message Levels* on page 249
- *Changing Logging Configuration* on page 250
- *Log Files for Event Stream Processor* on page 218

Viewing Sybase Control Center Server Logs

View event logs for the Sybase Control Center server.

Sybase Control Center logs events to several places:

- The console from which Sybase Control Center is launched.
- The Sybase Control Center agent log: `<SCC-install-directory>/log/agent.log`
- The repository log: `<SCC-install-directory>/log/repository.log`
- The component log for each installed Sybase Control Center product module. The path to the component log takes this form: `<SCC-install-directory>/plugins/<component>/log/<component>.log`

1. Display one of the log files using a log viewer or a method of your choice.
2. Look for entries of interest such as login attempts or the failure of a service to start.

On the console and in the Sybase Control Center agent log file, some components prepend the component name to log entries.

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Modifying the Event Stream Processor Log Configuration* on page 246
- *Viewing the Sybase Control Center Client Log* on page 248
- *Changing the Logging Level* on page 248
- *Logging or Message Levels* on page 249

- *Changing Logging Configuration* on page 250

Viewing the Sybase Control Center Client Log

Display the event log for the current session of your Sybase Control Center browser client.

In the perspective tab window (the main window), do either of the following to display the client log:

- Enter **Ctrl+Alt+L**.
- Select **View > Open > Log Window**.

Note: The client log reader displays the 100 most recent log messages for the current login session.

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Modifying the Event Stream Processor Log Configuration* on page 246
- *Viewing Sybase Control Center Server Logs* on page 247
- *Changing the Logging Level* on page 248
- *Logging or Message Levels* on page 249
- *Changing Logging Configuration* on page 250

Changing the Logging Level

Adjust the logging level that determines which events Sybase Control Center records in the server logs. This task requires you to restart Sybase Control Center.

If you are having a problem with Sybase Control Center, you might be able to discover the cause of the problem by changing the server logging level so that more events are recorded.

1. Shut down Sybase Control Center.
2. Restart Sybase Control Center using the -m option to change the logging level. In <SCC-installation-dir>/bin, enter:

```
scc -m <logging-level>
```

The logging levels are OFF (logs nothing), FATAL (logs only the most severe events), ERROR, WARN, INFO, DEBUG, and ALL (logs everything).

3. Examine the server log for clues about what might be causing the problem.
4. When you have resolved the problem, set the logging level back to WARN, the default. Your log may become unmanageably large if you leave it at the DEBUG or ALL level.

Example

These commands, which must be executed in the installation directory, start Sybase Control Center with the logging level set to debug:

```
Windows: bin\scc -m DEBUG
UNIX: bin/scc -m DEBUG
```

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Modifying the Event Stream Processor Log Configuration* on page 246
- *Viewing Sybase Control Center Server Logs* on page 247
- *Viewing the Sybase Control Center Client Log* on page 248
- *Logging or Message Levels* on page 249
- *Changing Logging Configuration* on page 250
- *Starting and Stopping Sybase Control Center in Windows* on page 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78

Logging or Message Levels

Describes values you can use to control the types of events that are logged by Sybase Control Center.

These are the logging levels, from highest to lowest. The higher the level, the more serious an event must be to be logged. Each level includes all the levels above it—for example, if you set the logging level to WARN, you log events for the WARN, ERROR, and FATAL levels.

OFF	Nothing is logged. This is the highest level.
FATAL	Logs only very severe error events that lead the server to abort. This is the highest level at which events are logged.
ERROR	Logs error events that might allow the server to continue running.
WARN	Logs potentially harmful situations. WARN is the default logging level during normal operation (that is, after system initialization).
INFO	Logs informational messages that track the progress of the server in a coarse-grained fashion. INFO is the default logging level during the system initialization process.
DEBUG	Logs a larger set of events that provides a finer-grained picture of how the server is operating. This level is recommended for troubleshooting.
ALL	Logs all loggable events. This is the lowest level.

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Modifying the Event Stream Processor Log Configuration* on page 246
- *Viewing Sybase Control Center Server Logs* on page 247
- *Viewing the Sybase Control Center Client Log* on page 248
- *Changing the Logging Level* on page 248

- *Changing Logging Configuration* on page 250
- *scc Command* on page 86

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	<ul style="list-style-type: none">• Agent log – <code>log4j.appender.agent.File</code>• Repository log – <code>log4j.appender.repository.File</code>• Collection statistics log – <code>log4j.appender.collection-stats.File</code>• Alert server log – <code>log4j.appender.alert.File</code>• Gateway log – <code>log4j.appender.gateway.File</code>
Change the maximum size that a log file can reach before Sybase Control Center creates a new file	<ul style="list-style-type: none">• Agent log – <code>log4j.appender.agent.MaxFileSize</code>• Repository log – <code>log4j.appender.repository.MaxFileSize</code>• Collection statistics log – <code>log4j.appender.collection-stats.MaxFileSize</code>• Alert server log – <code>log4j.appender.alert.MaxFileSize</code>• Gateway log – <code>log4j.appender.gateway.MaxFileSize</code>

To	Modify
Change the number of log files Sybase Control Center backs up before deleting the oldest file	<ul style="list-style-type: none"> • Agent log – log4j.appender.agent.MaxBackupIndex • Repository log – log4j.appender.repository.MaxBackupIndex • Collection statistics log – log4j.appender.collection-stats.MaxBackupIndex • Alert server log – log4j.appender.alert.MaxBackupIndex • Gateway log – log4j.appender.gateway.MaxBackupIndex

5. Save and exit the file.
6. Start Sybase Control Center to make the logging changes take effect.

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246
- *Modifying the Event Stream Processor Log Configuration* on page 246
- *Viewing Sybase Control Center Server Logs* on page 247
- *Viewing the Sybase Control Center Client Log* on page 248
- *Changing the Logging Level* on page 248
- *Logging or Message Levels* on page 249
- *Starting and Stopping Sybase Control Center in Windows* on page 75
- *Starting and Stopping Sybase Control Center in UNIX* on page 78

Sybase Control Center Console

The console is a command-line interface for displaying details about the status of the Sybase Control Center server and its subsystems.

When you use the **scc** command to start Sybase Control Center, it displays start-up messages and then displays the console prompt.

Note: The console prompt does not appear if you start Sybase Control Center as a service, if you direct the output of **scc** to a file, or if you start Sybase Control Center in the background.

See also

- *Launching Sybase Control Center* on page 74

Console Commands

Use the Sybase Control Center console to get status information on Sybase Control Center and its ports, plug-ins, and services.

help Command

Display syntax information for one or more Sybase Control Center console commands.

Syntax

```
help [command_name]
```

Parameters

- **command_name** – optional. status, info, or shutdown. If you omit *command_name*, **help** returns information on all the console commands.

Examples

- **Example 1** – returns information on the **status** command:

```
help status
```

Permissions

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

See also

- *info Command* on page 252
- *shutdown command* on page 253
- *status Command* on page 254

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 252
- *shutdown command* on page 253
- *status Command* on page 254

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 252
- *info Command* on page 252
- *status Command* on page 254

status Command

Display the status of the Sybase Control Center agent, plug-in, or service components of Sybase Control Center.

Syntax

```
status [-a | --agent]
[-h | --help]
[-p | --plugin [plugin-name]]
[-s | --service [service-name]]
```

Parameters

- **-a | --agent** – display the status of the Sybase Control Center agent component.
- **-h | --help** – display information about the **info** command.
- **-p | --plugin [plugin-name]** – display the status of the specified Sybase Control Center plug-in (for example, ASEMap, the Adaptive Server[®] management module). Omit the plugin-name argument to return a list of plug-ins.
- **-s | --service [service-name]** – display the status of the specified Sybase Control Center service (for example, the Alert service or the Messaging service). Omit the service-name argument to return a list of services.

Examples

- **Example 1** – displays status information on the Repository service:

```
status --service Repository
```

Permissions

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

See also

- *help Command* on page 252
- *info Command* on page 252
- *shutdown command* on page 253

Troubleshoot Sybase Control Center for Event Stream Processor

Solve problems that occur in Sybase Control Center for Sybase Event Stream Processor.

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 114
- *Setting Up Security* on page 91

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 109
- *Configuring Memory Usage* on page 82

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 141

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 244
- *Graphing Performance Counters: the Statistics Chart* on page 166

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 208

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 82

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/binSend the heap dump files to Sybase technical support for analysis.

See also

- *Configuring Memory Usage* on page 82

Statistics Do Not Display

Problem: Some statistics do not display in Sybase Control Center for Event Stream Processor.

Solution: Check that you started the projects in Event Stream Processor with the `-t` option. Set this option to a non-zero value in the project deployment configuration `.ccr` file. For example, `<Option name="time-granularity" value=5"/>`. Refer to the *Event Stream Processor Administrators Guide* for more information.

Troubleshooting Tips

Tips for troubleshooting issues with errors, and resetting your Sybase Control Center configuration.

To obtain error information about any issues you encounter in Sybase Control Center for Event Stream Processor, refer to the `SCC-3_2/log` directory.

If you need to reset your Sybase Control Center for Event Stream Processor configuration, stop SCC, refer to the `scc_repository.db.orig` file in the `SCC-3_2\services\Repository` directory, and:

- Delete the `scc_repository.log` file
- Copy the `scc_repository.db.orig` file to the `scc_repository.db` file
- Register and authenticate your SCC resources again

If you are unable to start Sybase Control Center, and you see this error:

```
Plugin 'com.sybase.ua.services.plugin.PluginRegisterService' registration failed: Failed to register plugin com.sybase.ua.plugins.espmap_3.2.7. Failed to start database engine java.lang.IllegalStateException: ASAEngine: The "scc_repository" dataserver Engine could not be started.
```

The problem is that the Sybase Control Center "scc_repository" database is out of sync with the `scc_repository.log` log file, and so SCC cannot start. To resolve this issue, follow the instructions above for reverting to a clean SCC repository.

See also

- *Viewing the Sybase Control Center for Event Stream Processor Log* on page 246

Glossary: Sybase Control Center for Event Stream Processor

Glossary of Sybase Control Center terms related to Event Stream Processor.

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.

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.

database – a collection of tables that are related by primary and foreign keys. The tables hold the information in the database. The tables and keys together define the structure of the database.

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.

instance – an SCC agent or server run from a shared disk installation. See also shared-disk mode.

job – a task performed by the scheduler in Sybase Control Center.

key performance indicator (KPI) – a single metric used to evaluate the status or performance of a monitored resource. A KPI value can be a state (such as running, error, or stopped) or a numerical value. KPIs are grouped into collections (and also, for some product modules, into key performance areas, or KPAs). KPI values are collected by scheduled collection jobs and appear on monitoring screens and in the statistics and heat charts. Examples of KPIs are resource state and CPU usage.

key performance area (KPA) – a group of related key performance indicators.

managed resource – see resource.

message row – a row that appears in the right pane of the Administration Console in place of a slow-responding request, a failed request, or a large result set. Rows with slow-responding requests are populated as soon as the data arrives. You can retry failed requests or expand large result sets—select the row and click the drop-down arrow to see options.

node – a topology object representing a server or other entity type, displayed in the form of an icon.

perspective – a named tab in Sybase Control Center that displays information related to a collection of managed resources (such as servers) and a set of views associated with those resources. The views in a perspective are chosen by users of the perspective. You can create as many perspectives as you need, and customize them to monitor and manage your resources. Perspectives allow you to group resources in ways that make sense in your environment—for example by location, department, or project.

product module – a plug-in component of Sybase Control Center that manages and monitors a particular Sybase product. SCC product modules are available for Adaptive Server, Data Assurance (a Replication Server option), replication (Replication Server, Replication Agent, and Mirror Replication Agent), Sybase Event Stream Processor, and Sybase IQ.

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, Data Assurance Server, Replication Server, Replication Agent, Mirror Replication Agent, Sybase Event Stream Processor, 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 remote command and control agent for Sybase Control Center that runs on a managed server. The SCC agent is installed automatically as part of the Sybase server.

schedule – the definition of a task (such as the collection of a set of statistics) and the time interval at which Sybase Control Center executes the task.

screen refresh interval – the period in seconds between refreshes of screens in the monitor views (ESP Node Monitor and ESP Cluster Monitor). Refreshing a screen redraws it with the most recent available data. Set the screen refresh interval on the Settings screen of either monitor view. See also collection repeat interval.

shared-disk mode – a feature that enables multiple instances of Sybase Control Center to execute from a single installation on a shared disk. Instances can be SCC servers, agents, or a mixture of the two.

singleton installation – a Sybase Control Center installation that runs a single SCC agent or server. Contrast with instance; see also shared-disk mode.

store – a store is one or more dbspaces that store persistent or temporary data for a special purpose. See catalog store, main store, or temporary store.

trend period – See chart trend period.

view (SCC) – a window in a Sybase Control Center perspective that displays information about one or more managed resources. Some views also let you interact with managed resources or with SCC itself. For example, the Perspective Resources view lists all the resources managed by the current perspective. Other views allow you to configure alerts, view the topology of a replication environment, and graph performance statistics.

Index

- Xmx maximum memory option 41, 84
- XX:MaxPermSize permanent memory option 41, 84

A

- access control
 - reloading policy file 168
- accessibility 11
- adapter statistics
 - ESP cluster 190
 - ESP node 176
- adapters
 - starting 199
 - stopping 200
 - Sybase IQ Output adapter 201
 - view file activity 201
- adapters, ESP 198
 - monitor 198
 - view 198
- adding a project 197
- adding a workspace, ESP 167
- administration
 - enabling for ESP 127
- Administration Console
 - adding a project 197
 - column filtering 6
 - display tools and options 4
 - displaying only selected resources 202, 203
 - ESP adapters 198
 - ESP projects 194
 - message rows in result sets 203
 - proxy rows 203
 - removing a project 197
 - setting data retrieval thresholds 113
 - starting a node 179
 - starting a project 195, 197
 - starting an adapter 199
 - stopping a node 180
 - stopping a project 196, 197
 - stopping an adapter 200
 - using 202
- Adobe Flex 11
- alert notifications 261
- alert severities
 - Event Stream Processor 158
- alert storm 261
- alert storm suppression 261
- alert subscriptions
 - disabling 216
 - enabling 216
- alert type 261
- alert types
 - Event Stream Processor 158
- alert-triggered scripts 159, 213
 - examples 160
 - substitution parameters 161
- alerts 261
 - about 210
 - configured, deleting 214
 - configured, modifying 212
 - configured, viewing 212
 - configuring duplicate alerts 152
 - configuring e-mail server 62, 111
 - configuring escalations 152
 - configuring storm suppression 152
 - configuring subscriptions 152
 - configuring to execute scripts 152
 - creating 151
 - displaying history 217
 - displaying resolutions 217
 - effects of repository purging on history 244
 - escalations 214
 - Event Stream Processor 153
 - log 159
 - modifying subscriptions 215
 - monitoring 212
 - not being generated 256
 - notifications, about 217
 - notifications, viewing 212
 - resolving 218
 - script examples 160
 - scripts executed by 159
 - setting triggering states and thresholds 151
 - subscribing to 215
 - subscriptions 214
 - substitution parameters for scripts 161
 - testing 213
 - triggering scripts, about 159
 - types, states, and severities 211
 - unsubscribing from 216

Index

- ALL logging level 249
- all statistics
 - ESP node 171
- authenticate 261
- authenticating
 - an ESP node 131
 - SCC agents 136
 - SCC with a managed resource 18, 140
- authentication
 - about 43, 92
 - configuring for LDAP 47, 95
 - configuring for UNIX 45, 94
 - configuring for Windows 44, 93
 - reloading policy file 168
- authentication type, ESP 132
- authentication type, ESP cluster
 - Kerberos 134
 - LDAP 133
 - Native OS 132
 - RSA 134
- authentication, ESP
 - clearing authentication parameters 136
- authorization 63, 114
- availability 261

B

- background, running SCC or SCC agent in 35, 78
- backups
 - about 240
 - changing the schedule 241
 - forcing 242
 - restoring from 243
 - scheduling 240
 - suspending and resuming 241
- badges, status 3

C

- changing the screen refresh interval 150
- clearing authentication parameters
 - ESP node 136
- client log, viewing 248
- Cluster policies to enable administering of ESP 124
- clusters
 - monitor 181
- collections 261
- column filtering in SCC 6

- columns
 - sorting by 4
- configuring
 - SCC agent connection data 136
- connection statistics
 - ESP cluster 189
 - ESP node 175
- console
 - about 251
 - commands 251
- conventions, style and syntax 9
- csi_config.xml file 47, 95

D

- data collection jobs
 - adding 141
 - adding schedules 207
 - creating 141
 - deleting 205
 - displaying history 209
 - executing 205
 - not saving data 141
 - removing schedules 208
 - resuming 206
 - stopping 205
 - suspending 206
 - viewing schedules 208
- data collection schedules
 - adding 141
 - modifying 208
- data collections
 - Event Stream Processor 143
 - troubleshooting timeouts 257
- data retrieval thresholds
 - setting for Administration Console 113
- databases 261
- DEBUG logging level 249
- defined 261–263
- display options in Sybase Control Center 4
- drivers
 - ODBC, registering 14, 75

E

- e-mail server, configuring for alerts 62, 111
- environment variables
 - SCC_MEM_MAX 39–41, 82, 84, 85
 - SCC_MEM_PERM 39–41, 82, 84, 85

- ERROR logging level 249
 - errors
 - OutOfMemory 258
 - REJECT_DUPLICATE_RESOURCE_AND_COLLECTION policy 257
 - timeouts for data collections 257
 - ESP cluster
 - adapter statistics 190
 - connection statistics 189
 - node statistics 184
 - overview statistics 22, 182
 - project statistics 186
 - publisher statistics 191
 - stream statistics 188
 - subscribers statistics 192
 - topology statistics 183
 - update authentication type 132
 - viewing stream schema 194
 - ESP node
 - adapter statistics 176
 - all statistics 171
 - authenticating 131
 - connection statistics 175
 - overview statistics 20, 169
 - project statistics 172
 - publisher statistics 177
 - registering 17, 128
 - stream statistics 174
 - subscriber statistics 178
 - viewing stream schema 181
 - ESP nodes
 - monitor 169
 - registering unregistered nodes 185
 - espAdminRole 127
 - espMonitorRole 16, 126
 - evaluation
 - quick start instructions 13
 - Event Stream Processor
 - adapters 198
 - adding a project 197
 - alert severities 158
 - alert types 158
 - alerts 153
 - configuring for administration 127
 - configuring for monitoring 16, 126
 - configuring log settings 246
 - data collections 143
 - espAdminRole 124, 127
 - espMonitorRole 16, 124, 126
 - log file 246
 - mapping groups to roles 16, 126, 127
 - projects 194
 - removing a project 197
 - starting a project 195, 197
 - starting an adapter 199
 - starting the ESP Server 138
 - stopping a project 196, 197
 - stopping an adapter 200
 - Sybase IQ Output adapte, file activity 201
 - troubleshooting 255
 - troubleshooting tips 259
 - view adapters 198
 - view file activity, Sybase IQ Output adapter 201
 - view projects 194
 - events 261
 - expiration dates for login accounts 120
- ## F
- F11 (browser full screen mode toggle) 6
 - F5 (browser refresh)
 - logging out of Sybase Control Center 256
 - FATAL logging level 249
 - Flash Player 15
 - foreground, running SCC or SCC agent in 35, 78
 - full backups 240
 - full screen mode 6
- ## G
- getting started after installing 15
 - glossaries
 - SCC for Event Stream Processor terms 261
 - graphing statistics 166
 - grid format, using 4
 - groups 66, 120
 - adding login accounts 65, 117
 - assigning monitoring and administration roles 64, 115
 - creating 65, 116
 - in LDAP, mapping to SCC roles 57, 106
 - in OS, mapping to SCC roles 57, 106
 - remove login 117
 - removing 116
 - removing roles 115
 - SCC Administrator 57, 106
 - sybase 57, 106

Index

H

- heat chart 261
 - customizing columns 19, 165
 - display tools and options 4
 - displaying 19, 165
 - filtering resources displayed 19, 165
 - icons 3
 - launch icon 2
- help command (console) 252
- historical performance monitoring 166
- history displays for alerts 217

I

- icons
 - for server status 3
 - in SCC toolbar 2
 - minimize/maximize sections of a view 6
- incremental backups 240
- info command (console) 252
- INFO logging level 249
- instances 261
 - about 27, 69, 234
 - converting 232
 - deploying 25, 67, 231
 - deploying and managing 28, 70, 235
 - file locations 26, 68, 232
 - refreshing 232
 - removing 233
- interfaces files, importing resources from 129

J

- Java system properties
 - displaying information about 252
- jobs 262
 - modifying collection intervals 208
 - resuming 209
 - suspending 209
- jvmopt memory options for Windows services 39, 41, 82, 84

K

- key performance areas 262
- key performance indicators 262
 - getting values from heat chart icons 19, 165

- Key Performance Indicators
 - Event Stream Processor 147
- keyboard shortcuts 7
- keyboard shortcuts for Adobe Flex 11
- KPAs 262
- KPIs 262
 - Event Stream Processor 147

L

- layout for Sybase Control Center views 229
- LDAP
 - configuration properties 48, 97
 - configuring authentication 47, 95
 - configuring to authenticate SCC logins 43, 92
- log files 218
 - node 220
 - project 221
 - SCC agent 219
- log4j.properties file 250
- logging in to Sybase Control Center 42, 90
 - troubleshooting 255
- logging in to Sybase Control Center - first user 15
- logging levels 249
- logging out of Sybase Control Center 91
 - unintentionally, using F5 browser refresh 256
- login accounts
 - assigning monitoring and administration roles 64, 115
 - authenticating 18, 140
 - creating automatically (UNIX) 45, 94
 - creating automatically (Windows) 44, 93
 - default 118
 - expiration date, imposing 120
 - granting privileges with roles and groups 57, 106
 - modifying 120
 - native SCC, adding 118
 - predefined 66, 120
 - removing 119
 - removing roles 115
 - suspending 120
- login accounts, default
 - about 15
- login modules 43, 92
- login session timeout 91
 - setting 63, 112
- logs
 - agent log, viewing 247
 - alert services 159

- alert services log, about 245
 - changing the logging level 248
 - client log, about 245
 - client log, viewing 248
 - component logs, about 245
 - configuring 250
 - Event Stream Processor logs, configuring 246
 - Event Stream Processor logs, viewing 246
 - repository log, about 245
 - repository log, viewing 247
 - SCC agent log, about 245
 - script execution log, about 245
 - server logs, about 245
 - server logs, viewing 247
- M**
- managed resources 223, 262
- managed servers
 - See managed resources
- management
 - enabling for ESP 127
- managing workspaces, ESP 167
- mapping groups to roles
 - for Event Stream Processor 16, 126, 127
- memory
 - configuring 39, 82
 - displaying information about 252
 - warnings at startup 258
- memory leak 258
- memory, insufficient 255
- message levels 249
- message rows 262
 - about 113
 - using 203
- minimize/maximize icon 6
- monitor
 - clusters 181
 - ESP nodes 169
- monitoring
 - enabling for ESP 16, 126
 - performance 166
- N**
- Node policies to enable monitoring of ESP 124
- node statistics
 - ESP cluster 184
- nodes
 - Administration Console 179, 180
 - log file 220
 - monitor 169
 - registering unregistered nodes 185
 - starting a node 179
 - stopping a node 180
- nodes, Event Stream Processor 262
- O**
- ODBC drivers
 - registering 14, 75
- online help
 - resetting 256
- operating system
 - configuring to authenticate SCC logins 43, 92
- OutOfMemory errors, SCC 258
- overview statistics
 - ESP cluster 22, 182
 - ESP node 20, 169
- P**
- parameters for scripts 161
- passencrypt utility 59, 108
- passwords
 - encrypting 59, 108
 - for repository database dba account, changing 86
 - for SCC default login account 15
- performance monitoring 166
- Perspective Heat Chart view 19, 165
- Perspective Resources view
 - about 223, 226
 - display tools and options 4
 - icons 3
 - show/hide icon 2
- perspectives 262
 - about 226
 - adding resources 139, 224
 - creating 139, 227
 - removing 227
 - removing a resource 224
 - renaming 227
- pluggable authentication modules for UNIX
 - authentication 45, 94
- policy.xml file 168
- policy.xml file, ESP 124
- port conflicts 255

Index

- ports
 - changing 86
 - configuring 60, 109
 - default 86
 - displaying information about 252
 - postinstallation tasks 15
 - product modules 262
 - displaying versions 8
 - production environment, setting up SCC in 24
 - Project policies to enable monitoring of ESP 124
 - project statistics
 - ESP cluster 186
 - ESP node 172
 - projects
 - adding 197
 - log file 221
 - managing 197
 - removing 197
 - starting 195, 197
 - statistics, cluster 186
 - statistics, ESP node 172
 - stopping 196, 197
 - projects, ESP 194
 - monitor 194
 - view 194
 - properties
 - for resources, changing 225
 - publisher statistics
 - ESP cluster 191
 - ESP node 177
- ## Q
- quick start instructions 13
- ## R
- registering
 - an ESP node 17, 128
 - SCC agents 136
 - registering ESP nodes
 - unregistered 185
 - registration
 - about 223
 - REJECT_DUPLICATE_RESOURCE_AND_COLLECTION policy errors 257
 - reloading policy file 168
 - removing a project 197
 - removing a workspace, ESP 167
 - repository 240, 262
 - backing up 242
 - changing backup schedule 241
 - changing database dba password 86
 - configuring purging 244
 - restoring from backup 243
 - scheduling backups 240
 - resource explorer
 - launch icon 2
 - Resource Explorer
 - about 223
 - display tools and options 4
 - searching in 226
 - resources 262
 - about 223
 - adding to a perspective 139, 224
 - authenticating 18, 140
 - browsing and managing 202
 - changing connection properties 225
 - changing name 225
 - displaying availability 19, 165
 - filtering by column in Admin Console 4
 - importing in batch 129
 - modifying data collection schedules 208
 - removing from a perspective 224
 - searching and filtering in Admin Console 203
 - searching for 226
 - selecting for display in Admin Console 203
 - unregistering 223
 - restarts
 - configuring in UNIX 35, 78
 - configuring in Windows 32, 75
 - role mapping
 - for Event Stream Processor 16, 126, 127
 - role-mapping.xml file 57, 106
 - roles
 - assigning to users and groups 64, 115
 - mapping SCC roles to LDAP or OS groups 57, 106
 - predefined 66, 120
 - product level 63, 114
 - removing 115
 - system level 63, 114
 - roles, ESP
 - espAdminRole 127
 - espMonitorRole 16, 126
 - row retrieval threshold
 - setting for Administration Console 113
 - RSSD user name, using to authenticate 18, 140

S

- SAP Sybase ESP
 - versions supported 1
- Save data collected from this job checkbox 141
- SCC 179, 180, 218–221
- SCC Administrator group 57, 106
- SCC agent 263
 - deploying and managing instances 28, 70, 235
 - deploying instances from a shared disk 25, 67, 231
 - shared-disk mode 26, 68, 230
 - starting in UNIX 35, 78
 - starting in UNIX as a service 35, 78
 - starting in Windows 32, 75
 - starting in Windows as a service 32, 75
 - stopping in UNIX 35, 78
 - stopping in Windows 32, 75
 - unauthenticating 136
- scc command 86
 - using to launch Sybase Control Center 14, 74
- SCC_MEM_MAX 39–41, 82, 84, 85, 258
- SCC_MEM_PERM 39–41, 82, 84, 85
- SCC-enabled login account 262
- scc.bat 14, 32, 75
- scc.sh 35, 78
- sccadmin account
 - about 15
- sccAdminRole 66, 120
- sccd shell script 35, 78
- sccinstance command 28, 70, 235
- sccUserRole 66, 120
- scheduler
 - resuming 209
 - suspending 209
- schedules 204, 263
 - adding to a job 207
 - creating for a data collection job 141
 - removing from a job 208
 - viewing 208
- schema, ESP
 - viewing stream schema 181, 194
- screen refresh interval
 - changing 150
- screen refresh interval, Event Stream Processor 263
- screens
 - maximizing 6
 - maximizing and minimizing sections of a view 6
- scripts
 - alert-triggered 213
 - alert-triggered, examples 160
 - substitution parameters 161
 - triggered by alerts 159
- security 43, 92
 - configuring 42, 91
- security providers
 - configuring 43, 92
- server log, configuring for Event Stream Processor 246
- server log, viewing for Event Stream Processor 246
- server logs, viewing 247
- servers
 - authenticating 18, 140
 - displaying availability 19, 165
 - importing in batch 129
 - modifying data collection schedules 208
 - searching for 226
 - unregistering 223
- services
 - enabling and disabling 86
 - listing 252
- services, UNIX
 - configuring SCC memory options for 41, 85
 - running SCC or SCC agent as 35, 78
- services, Windows
 - configuring Sybase Control Center memory options for 41, 84
 - running SCC or SCC agent as 32, 75
- settings
 - changing the screen refresh interval 150
- severities for alerts 211
- shared-disk mode 263
 - about 27, 69, 234
 - enabling and disabling 26, 68, 230
- shutdown command (console) 253
- singleton installation 263
- sorting by column 4
- sql.ini files, importing resources from 129
- start up
 - automatic, configuring in UNIX 35, 78
 - automatic, configuring in Windows 32, 75
- starting a project 195, 197
- starting an adapter 199
- starting Sybase Control Center 14, 74

Index

- statistics
 - about 142
 - availability 142
 - performance 142
 - troubleshooting display issues 258
 - statistics chart
 - displaying data for a longer period 256
 - effects of repository purging on 244
 - graphing performance counters 166
 - troubleshooting 256
 - statistics, ESP
 - adapters - cluster 190
 - adapters - node 176
 - all statistics - node 171
 - connections - cluster 189
 - connections - node 175
 - nodes - cluster 184
 - overview - cluster 22, 182
 - overview - node 20, 169
 - projects - node 172
 - projects -cluster 186
 - publishers - cluster 191
 - publishers - node 177
 - streams - cluster 188
 - streams - node 174
 - subscribers - cluster 192
 - subscribers - node 178
 - topology - cluster 183
 - status command (console) 254
 - status icons and badges for resources 3
 - stopping a project 196, 197
 - stopping an adapter 200
 - store 263
 - storm suppression for alerts 152
 - stream statistics
 - ESP cluster 188
 - ESP node 174
 - streams, ESP
 - viewing stream schema 181, 194
 - subscriber statistics
 - ESP cluster 192
 - ESP node 178
 - substitution parameters for scripts 161
 - Sybase Control Center
 - accessibility 11
 - connecting a browser to 15
 - console commands 251
 - deploying and managing instances 28, 70, 235
 - deploying instances from a shared disk 25, 67, 231
 - display tools and options 4
 - displaying component versions 8
 - failure to start 255
 - keyboard shortcuts 7
 - log files 247
 - logging in 42, 90
 - logging out 91
 - logging out unintentionally with F5 256
 - shared-disk mode 26, 68, 230
 - starting 14, 74
 - starting in UNIX 35, 78
 - starting in UNIX as a service 35, 78
 - starting in Windows 32, 75
 - starting in Windows as a service 32, 75
 - stopping in UNIX 35, 78
 - stopping in Windows 32, 75
 - Sybase Control Center agent 1
 - configuring 136
 - Event Stream Processor 137
 - Sybase Control Center for SAP Sybase ESP 1
 - sybase group 57, 106
 - system properties
 - displaying information about 252
 - system-wide features
 - configuring 42, 91
- ## T
- terms
 - SCC for Event Stream Processor 261
 - text conventions 9
 - thresholds for data retrieval
 - setting for Administration Console 113
 - timeout
 - errors on data collections 257
 - setting for login sessions 63, 112
 - toolbar icons 2
 - topology statistics
 - ESP cluster 183
 - trouble displaying statistics 258
 - troubleshooting
 - Event Stream Processor 255
 - statistics not displaying 258
 - tips for Event Stream Processor 259
 - troubleshooting tips
 - Event Stream Processor 259
 - types of alerts 211

U

UNIX

- configuring authentication 45, 94
- running SCC or SCC agent in the background 35, 78
- running SCC or SCC agent in the foreground 35, 78
- starting, stopping SCC or SCC agent 35, 78
- unregistered nodes, ESP
 - registering 185
- updating the authentication type, ESP
 - Kerberos 134
 - LDAP 133
 - Native OS 132
 - RSA 134
- user accounts
 - default 118
 - native SCC, adding 118
 - native SCC, not using 43, 92
- user information
 - modifying 120
- user interface, about 2

V

- versions of SCC components
 - displaying 8
- view
 - ESP adapters 198
 - ESP projects 194
 - file activity, Sybase IQ Output adapter 201
- view layouts, Sybase Control Center
 - cascade 229

- close all 229
- horizontal tiling 229
- minimize all 229
- restore all 229
- vertical tiling 229

View menu 6

views

- icons for managing 2
- maximizing and minimizing sections 6
- views, SCC 263
- views, Sybase Control Center
 - about 228
 - bringing to front of perspective 228
 - closing 228
 - maximizing 228
 - minimizing 228
 - opening 228
 - restoring 228

W

WARN logging level 249

Windows

- configuring authentication 44, 93
- starting, stopping Sybase Control Center or SCC agent 32, 75

workspaces, ESP

- adding 167
- managing 167
- removing 167

