



Sybase Control Center for Replication

3.2.6

DOCUMENT ID: DC01677-01-0326-01

LAST REVISED: March 2012

Copyright © 2012 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 Replication	1
New Features in Sybase Control Center for Replication	1
Replication Environment	2
Compatibility Matrix	3
User Interface Overview	4
Toolbar Icons	4
Status Icons	5
Common Display Options	6
Keyboard Shortcuts	9
Style and Syntax Conventions	10
Accessibility Features	12
Sybase Control Center Accessibility Information	12
Get Started	13
Quick Start for an Evaluation	13
Get Started in a Production Environment	19
Deploying an Instance from a Shared Disk Installation	60
Enabling and Disabling Shared-Disk Mode	61
Shared-Disk Mode	62
sccinstance Command	63
Launching Sybase Control Center	67
Registering the ODBC Driver in Windows	68
Starting and Stopping Sybase Control Center in Windows	68
Starting and Stopping Sybase Control Center in UNIX	71
Configuring Memory Usage	75
scc Command	78
Logging in to Sybase Control Center	82
Logging out of Sybase Control Center	83

Setting Up Security	83
Security	84
Configuring Authentication for Windows	85
Configuring a Pluggable Authentication Module (PAM) for UNIX	86
Configuring an LDAP Authentication Module	88
Mapping Sybase Control Center Roles to LDAP or OS Groups	96
Encrypting a Password	98
Configuring Ports	98
Configuring the E-mail Server	101
Configuring the Automatic Logout Timer	102
User Authorization	103
Assigning a Role to a Login or a Group	103
Removing a Role from a Login or a Group	104
Adding a Group	105
Removing a Group	105
Adding a Login Account to a Group	106
Removing a Login Account from a Group	106
Adding a Login Account to the System	107
Removing a Login Account from the System	108
Modifying a User Profile	108
Logins, Roles, and Groups	109
Configure	111
Registering a Replication Server, Replication Agent, or Mirror Replication Agent	112
Importing Resources for Batch Registration	113
Creating a Perspective	115
Adding a Resource to a Perspective	115
Role Assignment in Sybase Control Center for Replication	115
Authenticating a Login Account for a Managed Resource	116
Roles or Permissions for Setting Up Replication Monitoring Accounts	116

Encrypted Authentication for Replication	117
Setting Up the Latency Monitoring Process	117
Latency Monitoring Process	119
Setting Up Statistics Collection	119
About Statistics	121
Replication Data Collections	122
Key Performance Indicators and Key Performance Areas for Replication	123
Setting Replication Parameters	144
Configuring Replication Server and Replication Agent Parameters	144
Configuring Connection Parameters for Replication Server	145
Configuring Logical Connection Parameters for Replication Server	146
Configuring Route Parameters for Replication Server	146
Configuring Adaptive Server RepAgent Thread Parameters	147
Configuring Replication Server to Gather Zero Values for Counters	147
Creating an Alert	148
Replication Alerts	152
Alert Types, Severities, and States for Replication	157
Alert-Triggered Scripts	162
Substitution Parameters for Scripts	163
Optional Configuration Steps	164
Manage and Monitor	167
Heat Chart	167
Displaying Resource Availability	167
Historical Performance Monitoring	168
Graphing Performance Counters	169
Manage Sybase Control Center	170
Job Scheduling	170

Alerts	175
Resources	183
Perspectives	186
Views	187
Instances	189
Repository	199
Logging	205
Sybase Control Center Console	209
Manage and Monitor the Replication Environment	213
Availability Monitoring	213
Replication Latency	230
Troubleshoot Sybase Control Center for Replication	233
Collection Job for Adaptive Server Fails	233
Data on Dashboards or Charts Is Missing	233
No Data Is Returned for Performance Counters	
Collection Job	233
Same Data Is Displayed for Two or More Adaptive	
Server Collection Job Executions	234
Servers Are Missing in the Topology View	234
Servers, Routes, or Connections Display an Unknown	
Status	234
Some Connector Lines Do Not Display a Status Icon	
.....	234
Topology Display Is Incorrect When Two Replication	
Servers With the Same Name Are Registered	235
Topology View Displays Blank Canvas	235
Topology View Takes a Long Time to Display	235
Unable to Retrieve Topology Relationship Data or	
Server Objects	235
Problems with Basic Sybase Control Center	
Functionality	236
Cannot Log In	236
Sybase Control Center Fails to Start	236
Browser Refresh (F5) Causes Logout	237
Alerts Are Not Generated	237

Performance Statistics Do Not Cover Enough Time	237
Resetting the Online Help	237
Data Collections Fail to Complete	238
Memory Warnings at Startup	238
SCC Out of Memory Errors	239
Glossary: Sybase Control Center for Replication	241
Index	245

Contents

About Sybase Control Center for Replication

Sybase® Control Center for Replication is a Web-based solution that replaces Replication Monitoring Services (RMS), a middle-tier server, in monitoring the status and availability of servers in a replication environment.

Sybase Control Center for Replication allows you to monitor and manage large, complex, and geographically dispersed replication environments. It lets you search, sort, and filter servers and component objects to support a larger environment than what the current Replication Manager and Replication Monitoring Services can handle.

Sybase Control Center for Replication provides status information at a glance, using server monitors and a heat chart for displaying the availability or status of a specific server. The server monitors display high-level information, such as server version and platform. The server monitors also display critical performance counters to aid you in troubleshooting replication performance.

To help you control the flow of data and configure replication parameters to improve server performance, Sybase Control Center for Replication provides a quick administration tool that you can easily access through every replication monitor.

In addition to the monitors, Sybase Control Center for Replication provides a topology view that graphically displays the servers, the connections between servers, data flow in the environment, and a replication path's sources and targets. Graphs and charts are also available for monitoring performance counters.

New Features in Sybase Control Center for Replication

A brief description for each new and enhanced feature, and links to associated topics.

Table 1. New and enhanced Sybase Control Center for Replication features

Feature	Topics
Multi-Path Replication™ – monitor multiple Rep-Agent thread connections from the primary database, multiple dedicated routes between two Replication Servers, and multiple replicate connections to the same replicate database.	<i>Replication Paths</i> on page 227 <i>Replication Monitors</i> on page 216
Note: Multipath replication is supported in Replication Server 15.7 and later.	

Feature	Topics
Common Security Infrastructure – updates to the CSI improve SCC security and provide consistency across Sybase products.	—

Replication Environment

A replication environment is a set of servers—including Replication Servers, data servers, and Replication Agents—that participate in replication.

A replication environment does not need to contain all the server types in a replication system domain. It can be a subset of the replication domain identified by the ID Server. In the replication system, one Replication Server® serves as the ID Server. The ID Server assigns unique ID numbers to every Replication Server and database in the system. The ID Server also maintains version information for the system.

These are the servers that you can monitor in your Sybase Control Center replication environment:

- **Data Server** – a database within a data server may act as the source (primary database) of transaction data or may be the destination (replicate or standby database) of transaction data.
- **Replication Agent™ for Adaptive Server® Enterprise (ASE)** – is embedded within an Adaptive Server Enterprise and is called RepAgent Thread. The RepAgent Thread reads the transaction data from the primary database transaction log and sends it to Replication Server for distribution to the replicate databases.
- **Replication Server** – receives the replicated transaction data from a replication agent or another Replication Server, converts it into SQL, and applies the SQL to the replicate or standby databases. Replication Server delivers operational data across complex and broadly distributed data-infrastructures in near real-time. The primary and replicate databases can be Sybase or non-Sybase data servers.
- **Replication Agent** – reads the primary database transactions from database transaction logs, then sends those transactions to Replication Server for distribution to the replicate databases. Replication Agents allow non-Sybase data servers to act as the primary data servers in a replication system based on Sybase replication technology.
- **Mirror Replication Agent** – a component in a Mirror Activator disaster recovery environment and, in conjunction with a storage replication system, reads the primary database transactions from mirror log devices, then sends those transactions to Replication Server for distribution to the standby databases. The Mirror Activator provides an integrated disaster recovery solution with:
 - Standby databases protected from disk corruption (by logical, not literal, replication).
 - Synchronous replication, with zero data loss and guaranteed data integrity (transaction consistency).

- Complete coverage for databases, as well as non-database systems.
- **Sybase IQ** – a highly optimized analytics server that works with diverse data, including unstructured data and different data sources. It serves as a replicate database in the replication environment.

The replication environment also includes:

- **Components** – are objects in a server in a replication environment. Examples of components in a Replication Server are connections, routes, and queues.
- **Replication path** – is the set of servers through which transactions pass when moving from the primary to the replicate database.

See also

- *Data Flow Control* on page 217
- *Replication Servers* on page 218
- *Replication Agents* on page 225
- *Replication Paths* on page 227
- *Compatibility Matrix* on page 3

Compatibility Matrix

Sybase Control Center supports Replication Server and Replication Agent in these Replication bundles: Real-time Loading Edition (RTLE), Replication Server Heterogeneous Edition (RSHE), and Replication Server Options (RSO).

Sybase Control Center supports these server versions:

Server	Version
Replication Server	12.6 or later
Replication Agent	15.0 or later
Mirror Replication Agent	15.0 or later
Adaptive Server Enterprise	15.0.2 or later for Sybase Control Center for Adaptive Server, which offers full monitoring capabilities. 15.0.3 or later for Adaptive Server Cluster Edition monitoring. 12.5.4 or later for Sybase Control Center for Replication, which monitors only the RepAgent threads in the registered Adaptive Server.
Sybase IQ	15.1 ESD #2.1 or later for Sybase Control Center monitoring. 15.3 or later for Sybase Control Center administration.

About Sybase Control Center for Replication

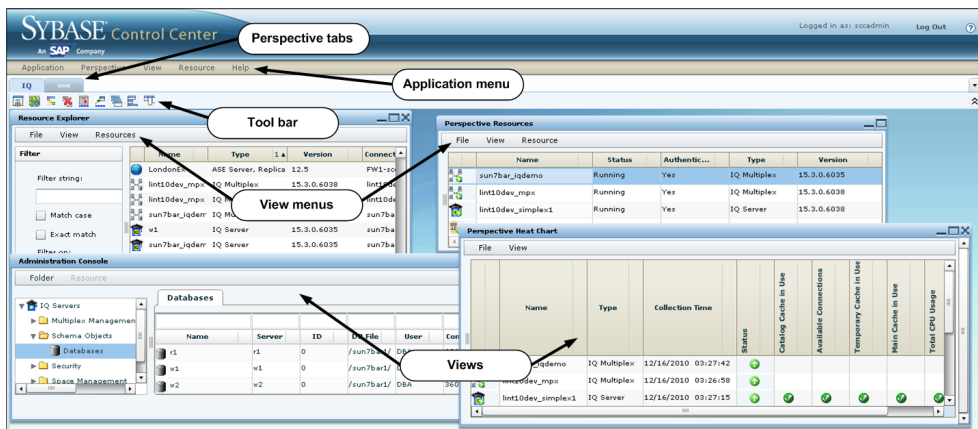
See also

- *Replication Environment* on page 2
- *Data Flow Control* on page 217
- *Replication Servers* on page 218
- *Replication Agents* on page 225
- *Replication Paths* on page 227
- *Resources* on page 183
- *Registering a Replication Server, Replication Agent, or Mirror Replication Agent* on page 112

User Interface Overview

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

Figure 1: Sybase Control Center User Interface

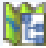









Toolbar Icons

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

Table 2. Toolbar icons

Icon	Name	Description
	Show/Hide Perspective Resources View	Displays or minimizes the Perspective Resources view, which lists registered resources in this perspective.

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

Status Icons


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

Resource Status Icons in the Perspective Resources View and Heat Chart






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

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

Table 3. Resource Status Icons

Icon	Status	Description
	Running	Resource is up and running

About Sybase Control Center for Replication




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

KPI Status Icons in the Heat Chart

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

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

Table 4. KPI Status Icons

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

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











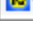
The Perspective Resources view, Resource Explorer, Administration Console, Alert Monitor, heat chart, and other views in Sybase Control Center—including those in product modules—use a tabular grid format to display information about managed resources. You can use options provided by the grid format to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

Table 5. Column Sorting Options

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

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

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

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


Filter by Column

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

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


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

See also

- *Resources* on page 183

- *Registering a Replication Server, Replication Agent, or Mirror Replication Agent* on page 112

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.

Key Sequence	Action
End	Highlight the last item in the active view (or the active section of a view), for example the last row in a table.
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.

Style and Syntax Conventions

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

Table 6. Style Conventions

Key	Definition
<code>monospaced(fixed-width)</code>	<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).

Key	Definition
<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 7. Syntax Conventions

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

Accessibility Features

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

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

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

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

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

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

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

Sybase Control Center Accessibility Information

Sybase Control Center uses the Adobe Flex application.

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

Note: To use Sybase Control Center 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. *Registering a Replication Server, Replication Agent, or Mirror Replication Agent*

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

5. *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.

6. *Displaying Resource Availability*

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

7. *Displaying Replication Server Status*

View the most recent Replication Server statistics. Sybase Control Center automatically updates the page when new status and statistical values are collected.

8. *Displaying Replication Agent Status*

Review high-level server availability by viewing the Replication Agent or Mirror Replication Agent status.

See also

- *Get Started in a Production Environment* on page 19

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.

Registering a Replication Server, Replication Agent, or Mirror Replication Agent

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

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

Table 8. New Resource Type Details

Field	Description
Resource Name	(Required) Name of the resource to register. Enter the actual name of the server, using uppercase and lowercase letters. If the name registered in Sybase Control Center does not exactly match the server name, some monitoring functions, including the topology view, do not work.
Resource Type	Select a resource type: <ul style="list-style-type: none"> • ASE Server, Replication Only – monitor only the RepAgent threads for an Adaptive Server that is older than version 15.0.2.0. Choose this type for an Adaptive Server that is part of a replication environment. • Replication Agent – monitor Replication Agent. • Replication Server – monitor Replication Server.
Description	A brief description to help you identify the resource.

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

Table 9. New resource connection details

Field	Description
Server Host Name/Host Name	Local host name
Port Number	Local host port number
RSSD Host Name	(Required) Name of the RSSD host
RSSD Port Number	(Required) RSSD host port number

Note: **RSSD Host Name** and **RSSD Port Number** are required only if you are registering a Replication Server.

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

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

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

6. Click **Next**.
7. (Optional) Click **Add this resource to the current perspective**. You must add a resource to a perspective (not necessarily the current perspective) before you can manage or monitor it.
8. (Optional) Click **Open the resource explorer to view this new resource**. (This option is not present when the Resource Explorer is open.)
9. Click **Finish**.

Authenticating a Login Account for a Managed Resource

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

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

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

1. Connect a browser to Sybase Control Center and log in.
2. If the Perspective Resources view is not open, click the **Show/Hide Perspective Resources View** icon in the toolbar.
3. In the Perspective Resources view, select your resource and select **Resource > Authenticate** from the view menu.
4. Select **Use my current SCC login** or **Specify different credentials**.
5. If you chose **Specify different credentials**, enter the login and password for Sybase Control Center to use to connect to your resource.
6. If the selected resource is a Replication Server, also enter the RSSD user name and password.
7. Click **OK** to save and exit the dialog.

Displaying Resource Availability

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

Get Started

1. From the application menu bar, select **View > Open > Heat Chart**.
2. (Optional) To display information about the status represented by an icon in the chart, hover the mouse over the icon.
 - Status column – icon tooltips describe the status of the resource (Running or Stopped, for example).
 - All columns to the right of Status – icon tooltips give the value of the KPI listed at the top of the column.
3. (Optional) To display tools for filtering (narrowing the list of resources in the heat chart) or changing the columns, select **View > Filter** from the Perspective Heat Chart menu bar. The Filter and Column tools appear in the left pane.
4. (Optional) To use filtering, select **View > Filter** from the view's menu bar and enter a search term in the **Filter string** field.

The search term can be any string that appears in the tabular portion of the heat chart, such as the name, or part of the name, of a server or a resource type (ASE Server, for example).
5. (Optional) Select a filtering setting:
 - **Match case** – search for resources whose displayed data includes the search term, including uppercase and lowercase letters; or
 - **Exact match** – search for resources whose displayed data includes an item identical to the search term.
6. (Optional) Select a column from the **Filter on** list to restrict your search to that column.
7. (Optional) Click **Columns** to customize your heat chart.
8. (Optional) Unselect any column that should not appear in your heat chart.
9. (Optional) Click the sorting arrow in the column headers to sort the column values in either ascending or descending order.
10. (Optional) Click the resource's row and pull down the menu to the right of the resource name to view options for the selected resource.
11. (Optional) To resize the Filter and Columns tools pane, move your mouse over the border between the tools pane and the resource table. When the mouse cursor changes to a resize icon, click and drag the border to the left or the right.
12. (Optional) To hide the Filter and Columns tools, unselect **View > Filter**.

Displaying Replication Server Status

View the most recent Replication Server statistics. Sybase Control Center automatically updates the page when new status and statistical values are collected.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have permission to perform this task.

Task

Note: You can also view the server status through the status badges on the server icons in the Perspective Resources, Perspective Heat Chart, Resource Explorer, and Replication Topology views.

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor** .
The Replication Server monitor view appears, with the Overview as the default display.
2. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

Displaying Replication Agent Status

Review high-level server availability by viewing the Replication Agent or Mirror Replication Agent status.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Agent and click **Resource > Monitor** .
The monitor view appears displaying high-level information, which includes the status of the selected Replication Agent.
2. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

Get Started in a Production Environment

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

Prerequisites

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

1. *Deploying an Instance from a Shared Disk Installation*
(Optional) Create a Sybase Control Center server or agent from an installation on a shared disk.
2. *Starting and Stopping Sybase Control Center in Windows*

Get Started

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

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

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

4. *Configuring Memory Usage*

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

5. *Logging in to Sybase Control Center*

Enter the Sybase Control Center Web console.

6. *Setting Up Security*

Configure login authentication and map roles.

7. *Configuring the E-mail Server*

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

8. *Configuring the Automatic Logout Timer*

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

9. *User Authorization*

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

10. *Configure*

Configure login accounts, statistics collection, and other replication monitoring options.

Deploying an Instance from a Shared Disk Installation

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

Prerequisites

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

Task

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

Note: You can create an instance on one host and run it on another host, but doing so interferes with the predeployment checks run by **sccinstance**. Such a deployment might

generate errors (port conflicts, for example). If you are confident that the errors are caused by problems that will not be present on the host where you plan to run the instance, use the **-force** option to create the instance.

2. Change to `SCC-3_2/bin`.
3. Create the instance as an SCC agent if you plan to run a managed server on this host. Create the instance as an SCC server if you plan to manage other Sybase servers from this host.

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

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

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

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

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

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

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

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

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

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

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


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

Next

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

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

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

See also

- *Starting and Stopping Sybase Control Center in Windows* on page 27

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 or Sybase IQ multiplexes.

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

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

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

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

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

sccinstance Command

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

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

Syntax

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

Parameters

- **-agent** – use with **-create** or **-refresh** to create or refresh an SCC agent. In a **-create** or **-refresh** command, **-agent** is the default, so you can omit it.
- **-create** – deploy a new instance. Use alone or with **-agent** to create an agent instance, or with **-server** to create a server instance.
- **-d | debug** – display debugging messages with the output of this command.
- **-disable** – turn off shared-disk mode for this installation. Generates an error if any instance is running.

Get Started

- **-enable** – turn on shared-disk mode for this installation. Shared-disk mode is required if you intend to run more than one server or agent from a single installation of SCC.
- **-f | -force** – execute **sccinstance** even if there are potential conflicts, such as port clashes or a running SCC process. Sybase does not recommend using **-force** to remove or refresh a running instance in a Windows environment.
- **-h | --help** – display help and usage information for the **sccinstance** command.
- **-host *host-name*** – specify the host for this instance. Use with **-create**; required only when the instance name does not match the name of the host on which this instance will run. (The instance name defaults to the name of the current host unless you use **-instance** to specify another name.)
- **-instance [*instance-name*]** – specify an instance. Use with **-create**, **-remove**, or **-refresh**, or use alone to display the instance’s status. You can omit **-instance** when you are addressing the only SCC instance or the only instance of the specified type (server or agent) on the current host.

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

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

Port information:

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

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	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

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

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

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

Examples

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

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

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

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

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

or

```
sccinstance -create -service
```

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

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

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

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

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

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

```
sccinstance -refresh -server
```

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

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

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

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

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

```
sccinstance
```

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

```
sccinstance -list
```

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

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

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

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

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

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

Permissions

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

Starting and Stopping Sybase Control Center in Windows

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

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

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

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

Note: To start an SCC agent or server as a service:

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

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

Here are the steps for each starting and stopping option:

- **Start Sybase Control Center or the SCC agent:**

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

For a singleton installation:

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

For an instance:

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

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

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

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

For a singleton installation:

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

For an instance:

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

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

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

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

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

For a singleton installation:

```
net start "sybase control center 3.2.6"
```

```
The Sybase Control Center 3.2.6 service is starting.....
The Sybase Control Center 3.2.6 service was started
successfully.
```

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

```
net start "sybase control center 3.2.6 (Boston-1)"
```

```
The Sybase Control Center 3.2.6 (Boston-1) service is
starting.....
The Sybase Control Center 3.2.6 (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 3.2.6"
```

```
The Sybase Control Center 3.2.6 service is stopping.....
The Sybase Control Center 3.2.6 service was stopped
successfully.
```

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

```
net stop "sybase control center 3.2.6 (Boston-1)"
```

```
The Sybase Control Center 3.2.6 (Boston-1) service is
stopping.....
```

```
The Sybase Control Center 3.2.6 (Boston-1) service was
stopped successfully.
```

See also

- *Deploying an Instance from a Shared Disk Installation* on page 20

Starting and Stopping Sybase Control Center in UNIX

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

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

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

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

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

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

Here are the steps for each starting and stopping option:

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

Bourne shell:

```
. SYBASE.sh
```

C shell:

```
source SYBASE.csh
```

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

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

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

For a singleton installation:

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

For an instance:

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

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

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

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

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

Bourne shell (sh) or Bash

For a singleton installation:

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

For an instance:

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

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

C shell

For a singleton installation:

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

For an instance:

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

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

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

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

```
shutdown
```

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

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

For a singleton installation:

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

For an instance:

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

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

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

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

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

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

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

b) Open `sccd` and make these changes:

- Change the line that sets the SYBASE variable to the location of your Sybase installation (that is, the parent of `SCC-3_2`, the Sybase Control Center installation directory). By default, this directory is called `Sybase`.
- If you are not using shared-disk mode, or you are using shared-disk mode to run a single instance whose name is the same as the host name, skip to step *5.c* on page 33 or step *5.d* on page 33.
- 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
<p>Maximum memory</p> <ul style="list-style-type: none"> • <code>jvmopt=-Xmx</code> – if you are running SCC 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 • SCC fails to start and may display an error: <code>Could not create the Java Virtual machine.</code> • An <code>OutOfMemory</code> error says SCC is out of heap space • A warning message about system memory appears during the start process • The machine where SCC is installed has less than 4GB of memory. (Starting SCC 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 SCC 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 <code>OutOfMemory</code> error says SCC 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 <code>OutOfMemory</code> 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 SCC service. When you use this method, your changes to the memory options persist—Sybase Control Center uses them every time it starts as a service.

See also

- *Logging in to Sybase Control Center* on page 36

Changing a Memory Option on the Command Line

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

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

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

Windows example:

```
> set SCC_MEM_MAX=512
```

UNIX example:

```
bash$ export SCC_MEM_MAX=512
```

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

See also

- *Changing a Memory Option for an SCC Windows Service* on page 35
- *Changing a Memory Option for an SCC UNIX Service* on page 36

Changing a Memory Option for an SCC Windows Service

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

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

1. If Sybase Control Center is running, shut it down.
2. Open the SCC properties file:


```
<SCC-install-directory>\SCC-3_2\bin\scc.properties
```
3. Add (or modify, if it already exists) a **jvmopt** line specifying the memory size in Java format. Use `m` for megabytes or `g` for gigabytes.

Get Started

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 35
- *Changing a Memory Option for an SCC UNIX Service* on page 36

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 35
- *Changing a Memory Option for an SCC Windows Service* on page 35

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@sybase`, not `fred@sybase.com`.

See also

- *Configuring Memory Usage* on page 34

Setting Up Security

Configure login authentication and map roles.

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

Note: These security topics are intended for use in a production environment. If you are evaluating or testing SCC, see *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*

Get Started

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 54

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 <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. 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 (scAdminRole).
 - 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.
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="1snjikfwregfqr43hu5io..." />
```

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 49

LDAP Configuration Properties

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

Property	Default Value	Description
ServerType	None	<p>Optional. The type of LDAP server you are connecting to:</p> <ul style="list-style-type: none"> • <code>sunone5</code> -- SunOne 5.x OR iPlanet 5.x • <code>msad2k</code> -- Microsoft ActiveDirectory, Windows 2000 • <code>nsds4</code> -- Netscape Directory Server 4.x • <code>openldap</code> -- 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"> • <code>RoleFilter</code> • <code>UserRoleMembership</code> • <code>RoleMemberAttributes</code> • <code>AuthenticationFilter</code> • <code>DigestMD5Authentication</code> • <code>UseUserAccountControl</code>
ProviderURL	<code>ldap://localhost:389</code>	<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 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.
SecurityProtocol	None	<p>The protocol to be used when connecting to the LDAP server.</p> <p>To use an encrypted protocol, use "ssl" instead "ldaps" in the url.</p> <hr/> <p>Note: ActiveDirectory requires the SSL protocol when setting the value for the password attribute. This occurs when creating a user or updating the password of an existing user.</p> <hr/>
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"> • simple — For clear-text password authentication. • DIGEST-MD5 — For more secure hashed password authentication. This method requires that the server use plain text password storage and only works with JRE 1.4 or later.

Property	Default Value	Description
AuthenticationFilter	<p>For most LDAP servers: (& (uid={uid}) (object- class=person))</p> <p>or</p> <p>For Active Directory email lookups: (&(user- Principal- Name={uid}) (object- class=user)) [ActiveDirec- tory]</p> <p>For Active Directory Windows username lookups: (& (sAMAccount- Name={uid}) (object- class=user))</p>	<p>The filter to use when looking up the user.</p> <p>When performing a username based lookup, this filter is used to determine the LDAP entry that matches the supplied username.</p> <p>The string "{uid}" in the filter is replaced with the supplied username.</p>
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>
AuthenticationSearchBase	none	<p>The search base used to authenticate users. If this value is not specified, the LDAP DefaultSearch-Base is used.</p>

Property	Default Value	Description
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 also use this DN to create the users in the LDAP server. When the self-registration feature is used, this user may also need the requisite permissions to create a user record. This behavior can occur if you do not set <code>useUserCredentialsToBind</code> to <code>true</code>. In this case, the LDAP attributer uses this DN to update the user attributes.</p>
BindPassword	none	<p>BindPassword is the password for BindDN, which is used to authenticate any user. BindDN and BindPassword are used to separate the LDAP connection into units.</p> <p>The <code>AuthenticationMethod</code> property determines the bind method used for this initial connection.</p> <p>Sybase recommends encrypting passwords and provides a password encryption utility for the purpose. If you encrypt BindPassword, include <code>encrypted=true</code> in the line that sets the option. For example:</p> <pre><options name="BindPassword" encrypted="true" value="lsnjikf-wregfqr43hu5io..." /></pre> <p>If you do not encrypt BindPassword, the option might look like this:</p> <pre><options name="BindPassword" value="s3cr3T" /></pre>
RoleSearchBase	none	<p>The search base used to retrieve lists of roles. If this value is not specified, the LDAP Default-SearchBase is used.</p>

Property	Default Value	Description
RoleFilter	<p>For SunONE/iPlanet: (& (object-class=ldapsu-bentry) (objectclass=nsroledefinition))</p> <p>For Netscape Directory Server: (object-class=groupof-names) (object-class=groupofunique-names))</p> <p>For ActiveDirectory: (object-class=groupof-names) (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 or this property is not initialized, no roles are available.</p>
RoleMemberAttributes	<p>For Netscape Directory Server: member,unique-member</p>	<p>The role's member attributes defines a comma-delimited list of attributes that roles may have that define a list of DN's of people who are in the role.</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 only has a default value when the Netscape server type is chosen.</p>
RoleNameAttribute	cn	<p>The attribute for retrieved roles that is the common name of the role. If this value is "dn" it is interpreted specially as the entire dn of the role as the role name.</p>
RoleScope	onelevel	<p>The role 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
UserRoleMembershipAttributes	For iPlanet/SunONE: nsRoleDN For ActiveDirectory: memberOf For all others: none	The user's role membership attributes property is used to define an attribute that a user has that contains the DN's of all of the roles as user is a member of. These comma-delimited values are then cross-referenced with the roles retrieved in the role search base and search filter to come up with a list of user's roles.
UserFreeformRoleMembershipAttributes	None	The "freeform" role membership attribute list. Users who have attributes in this comma-delimited list are automatically granted access to roles whose names are equal to the attribute value. For example, if the value of this property is "department" and user's LDAP record has the following values for the department attribute, { "sales", "consulting" }, then the user will be granted roles whose names are "sales" and "consulting".
Referral	ignore	The behavior when a referral is encountered. The valid values are those dictated by LdapContext, for example, "follow", "ignore", "throw".
DigestMD5AuthenticationFormat	DN For OpenLDAP: User-name	The DIGEST-MD5 bind authentication identity format.
UseUserAccountControlAttribute	For most LDAP servers: false For ActiveDirectory: true	The UserAccountControl attribute to be used for detecting disabled user accounts, account expirations, password expirations and so on. ActiveDirectory also uses this attribute to store the above information.
controlFlag	optional	Indicates whether authentication with this login module is sufficient to allow the user to log in, or whether the user must also be authenticated with another login module. Rarely set to anything other than "sufficient" for any login module. Note: controlFlag is a generic login module option rather than an LDAP configuration property.

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

Login Module	OS Group	Sybase Control Center Roles
	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 41

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
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

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

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

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

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

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

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

Examples

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

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

Set only the database messaging service to port 3639.

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

Set the HTTP port to 9292.

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

Set the Jini RMI daemon to port 9696.

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

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

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

Set the RMI port to 9991.

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

Set the Tabular Data Stream port to 9997.

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

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

Configuring the E-mail Server

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

Prerequisites

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

Task

1. From the 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 37

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. The roles can be assigned directly to login accounts or to groups; a login account inherits the roles of any group to which it belongs. Component product modules assign some roles automatically.

Sybase Control Center classifies roles as follows:

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

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

See also

- *Configure* on page 58

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 56
- *Adding a Login Account to a Group* on page 57
- *Logins, Roles, and Groups* on page 57

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 56
- *Adding a Login Account to a Group* on page 57
- *Logins, Roles, and Groups* on page 57

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 56
- *Adding a Group* on page 56
- *Logins, Roles, and Groups* on page 57

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 10. 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 11. 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 56
- *Adding a Group* on page 56
- *Adding a Login Account to a Group* on page 57

Configure

Configure login accounts, statistics collection, and other replication monitoring options.

1. *Registering a Replication Server, Replication Agent, or Mirror Replication Agent*
Register a resource (for example, a server that can be monitored) to make Sybase Control Center aware of it and its connection information.
2. *Importing Resources for Batch Registration*

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

3. *Creating a Perspective*

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

4. *Adding a Resource to a Perspective*

Add one or more resources to the current perspective.

5. *Role Assignment in Sybase Control Center for Replication*

In Sybase Control Center for Replication version 3.2.3, you no longer need to grant special roles for administrative or monitoring privileges on the Replication Server.

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. *Setting Up the Latency Monitoring Process*

Start the latency monitoring process for replication paths in the replication environment.

8. *Setting Up Statistics Collection*

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

9. *Setting Replication Parameters*

Configure replication parameters to improve server performance. Sybase Control Center for Replication allows you to configure the parameters for Replication Server, Replication Agent, connection and logical connection, route, and Adaptive Server RepAgent thread.

10. *Creating an Alert*

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

11. *Optional Configuration Steps*

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

See also

- *User Authorization* on page 55
- *Logins, Roles, and Groups* on page 109
- *Setting Up Security* on page 83
- *Assigning a Role to a Login or a Group* on page 103

Deploying an Instance from a Shared Disk Installation

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

Prerequisites

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

Task

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

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

2. Change to `SCC-3_2/bin`.
3. Create the instance as an SCC agent if you plan to run a managed server on this host. Create the instance as an SCC server if you plan to manage other Sybase servers from this host.

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

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

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

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

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

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

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

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

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

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

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

```
sccinstance -list
```

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

Next

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

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

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

See also

- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *Instances* on page 189

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 62
- *sccinstance Command* on page 63
- *Instances* on page 189

Shared-Disk Mode

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

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

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

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

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

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

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

See also

- *Enabling and Disabling Shared-Disk Mode* on page 61
- *sccinstance Command* on page 63
- *Instances* on page 189

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

Get Started

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
```

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

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

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

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

Permissions

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

See also

- *Enabling and Disabling Shared-Disk Mode* on page 61
- *Shared-Disk Mode* on page 62

Launching Sybase Control Center

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

Prerequisites

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

Task

1. Start Sybase Control Center.

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

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

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

See also

- *Sybase Control Center Console* on page 209
- *Logging in to Sybase Control Center* on page 82

Registering the ODBC Driver in Windows

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

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

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

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

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

See also

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

Starting and Stopping Sybase Control Center in Windows

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

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

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

- Use the **scc.bat** command to start Sybase Control Center or the SCC agent manually. The command gives you access to the Sybase Control Center console, which you can use to shut down and to display information about services, ports, system properties, and

environment variables. You can also use **scc.bat** to change the logging level for troubleshooting purposes. Using **scc.bat** prevents you from taking advantage of the automatic start and restart features available to services.

- Use the Services list under the Windows Control Panel to start, stop, and configure the Sybase Control Center service for an SCC server or agent.
- Use the **net start** and **net stop** commands. This is another way to run Sybase Control Center or the SCC agent as a service.

Note: To start an SCC agent or server as a service:

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

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

Here are the steps for each starting and stopping option:

- **Start Sybase Control Center or the SCC agent:**

- (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 or the SCC agent:**

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

For a singleton installation:

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

For an instance:

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

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

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

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

For a singleton installation:

```
net start "sybase control center 3.2.6"
```

```
The Sybase Control Center 3.2.6 service is starting.....  
The Sybase Control Center 3.2.6 service was started  
successfully.
```

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

```
net start "sybase control center 3.2.6 (Boston-1)"
```

```
The Sybase Control Center 3.2.6 (Boston-1) service is  
starting.....  
The Sybase Control Center 3.2.6 (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 3.2.6"
```

```
The Sybase Control Center 3.2.6 service is stopping.....
```

```
The Sybase Control Center 3.2.6 service was stopped
successfully.
```

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

```
net stop "sybase control center 3.2.6 (Boston-1)"
```

```
The Sybase Control Center 3.2.6 (Boston-1) service is
stopping.....
The Sybase Control Center 3.2.6 (Boston-1) service was
stopped successfully.
```

See also

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

Starting and Stopping Sybase Control Center in UNIX

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

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

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

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

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

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

Here are the steps for each starting and stopping option:

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

Get Started

- a) Change to the Sybase directory (the parent of the Sybase Control Center installation directory).
- b) Execute one of the following to set environment variables.

Bourne shell:

```
. SYBASE.sh
```

C shell:

```
source SYBASE.csh
```

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

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

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

For a singleton installation:

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

For an instance:

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

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

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

You can use **nohup**, **&**, and **>** to run Sybase Control Center or the SCC agent in the background, redirect output and system error to a file, and suppress the SCC console.

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

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

Bourne shell (sh) or Bash

For a singleton installation:

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

For an instance:

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

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

C shell

For a singleton installation:

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


For an instance:

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

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

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

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

```
shutdown
```

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

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

For a singleton installation:

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

For an instance:

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

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

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

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

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

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

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

b) Open `sccd` and make these changes:

- Change the line that sets the SYBASE variable to the location of your Sybase installation (that is, the parent of `SCC-3_2`, the Sybase Control Center installation directory). By default, this directory is called `Sybase`.
- If you are not using shared-disk mode, or you are using shared-disk mode to run a single instance whose name is the same as the host name, skip to step *5.c* on page 74 or step *5.d* on page 74.

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.

See also

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

Configuring Memory Usage

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

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

Modify this value	When	Guidelines
Maximum memory <ul style="list-style-type: none"> • <code>jvmopt=-Xmx</code> – if you are running SCC 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 • SCC fails to start and may display an error: Could not create the Java Virtual machine. • An OutOfMemory error says SCC is out of heap space • A warning message about system memory appears during the start process • The machine where SCC is installed has less than 4GB of memory. (Starting SCC 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 SCC 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 SCC 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 SCC service. When you use this method, your changes to the memory options persist—Sybase Control Center uses them every time it starts as a service.

See also

- *Registering the ODBC Driver in Windows* on page 68
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *scc Command* on page 78

Changing a Memory Option on the Command Line

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

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

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

Windows example:

```
> set SCC_MEM_MAX=512
```

UNIX example:

```
bash$ export SCC_MEM_MAX=512
```

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

See also

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

Changing a Memory Option for an SCC Windows Service

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

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

1. If Sybase Control Center is running, shut it down.
2. Open the SCC properties file:

```
<SCC-install-directory>\SCC-3_2\bin\scc.properties
```
3. Add (or modify, if it already exists) a **jvmopt** line specifying the memory size in Java format. Use `m` for megabytes or `g` for gigabytes.

For example:

```
jvmopt=-Xmx512m
```

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

See also

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

Changing a Memory Option for an SCC UNIX Service

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

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

1. If Sybase Control Center is running, shut it down.
2. Open the `sccd` file: `/etc/init.d/sccd`

Get Started

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

For example:

```
SCC_MEM_MAX=512
```

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

See also

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

scc Command

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

Syntax

```
scc[.bat | .sh] [-a | --address RMI-service-address]  
[-b | --bitwidth]  
[--dbpassword]  
[-disable | --disable service-name,service-name...]  
[-enable | --enable service-name,service-name...]  
[-h | --help]  
[-I | --info [information-category]]  
[-instance [instance-name]]  
[-m | --message message-level]  
[-password | --password password]  
  
[-p | --port {port-name=port-number |  
          service-name:property-name=port-number}]  
[{-start | --start} | {-stop | --stop}]  
[-status | --status]  
[-user | --user login-name]  
[-v | -version | --version]
```

Parameters

- **-a | --address *RMI-service-address*** – the address for the RMI service to use; must be an IP address on this machine or the name of this machine (which is the default).
- **-b | --bitwidth** – returns a string identifying the bit width (32 or 64) of the underlying platform; Sybase Control Center uses this option to determine which libraries to use for its internal database. If you use this option, the **scc** command does not start Sybase Control Center.
- **--dbpassword** – changes the password of the default dba account provided for the repository database. It prompts you for the new password, validates it, and starts the

Sybase Control Center server. This option does not work if you start Sybase Control Center in the background—the server fails to start if there is no console.

- **-disable** | **--disable** *service-name,service-name...* – disable the specified Sybase Control Center services. This option does not work while Sybase Control Center is running or as part of a command that starts SCC. To use it, shut down SCC, execute `scc --disable`, then restart. See under `--ports` for service names; separate each service from the next with a comma.
- **-enable** | **--enable** *service-name,service-name...* – enable the specified Sybase Control Center services. See under `--ports` for service names; separate each service from the next with a comma. When you use this option, **scc** does not start Sybase Control Center—use a separate command to start SCC.
- **-h** | **--help** – display help and usage information for the **scc** command. If you use this option, **scc** does not start Sybase Control Center.
- **-I** | **--info** [*information-category*] – display the specified categories of information about Sybase Control Center. Separate each category from the next with a comma. The information categories are:
 - `all` – returns all the information provided by the `sys`, `ports`, and `services` categories. Default option.
 - `sys` – returns general information about this instance of Sybase Control Center, including the version, the home (installation) directory, the host machine’s name and IP address, the RMI port number, the messaging level, and details about the platform and Java installation.
 - `ports` – lists all the ports on which the Sybase Control Center agent and its services listen, indicates whether each port is in use, and shows the service running on each port.
 - `services` – lists all the services known to the Sybase Control Center agent, indicates whether each service is enabled, and lists other services on which each service depends.
 - `sysprop` – lists all the Java system properties known the Java VM and their values.
 - `env` – lists the complete Java VM process environment.
- **-instance** [*instance-name*] – use with other options (**-start** and **-stop**, for example) to specify a Sybase Control Center instance in a shared disk deployment. If you do not enter a name for the instance, it defaults to the host name.
- **-m** | **--message** *message-level* – set the amount of detail recorded in system logs; also known as the logging level. Valid values are OFF, FATAL, ERROR, WARN, INFO, DEBUG, and ALL. WARN is the default.
- **-password** | **--password** – specify the password of the user account Sybase Control Center will use to stop servers or query them for status. Use this option with **--user**. When you enter a command with **--user** but without **--password**, the console prompts you to enter a password.
- **-p** | **--port** {*port-name=port-number* | *service-name:property-name=port-number*} – configure the specified service to run on the specified port. Changing ports is useful if you discover a port conflict between Sybase Control Center and other software on the same

Get Started

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 68
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71
- *Configuring Memory Usage* on page 75
- *Configuring Ports* on page 98
- *Logging or Message Levels* on page 207

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@sybase`, not `fred@sybase.com`.

See also

- *Launching Sybase Control Center* on page 67

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

- *Configuring the Automatic Logout Timer* on page 102

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 86
- *Configuring an LDAP Authentication Module* on page 88
- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 96

Configuring Authentication for Windows

Authentication through the Windows operating system is enabled by default. 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

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 96
- *Adding a Login Account to the System* on page 107

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. 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.

See also

- *Mapping Sybase Control Center Roles to LDAP or OS Groups* on page 96
- *Adding a Login Account to the System* on page 107

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.LDAPAuthenticator"
type="authenticator"/>
```

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 96

LDAP Configuration Properties

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

Property	Default Value	Description
ServerType	None	<p>Optional. The type of LDAP server you are connecting to:</p> <ul style="list-style-type: none"> • <code>sunone5</code> -- SunOne 5.x OR iPlanet 5.x • <code>msad2k</code> -- Microsoft ActiveDirectory, Windows 2000 • <code>nsds4</code> -- Netscape Directory Server 4.x • <code>openldap</code> -- 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"> • <code>RoleFilter</code> • <code>UserRoleMembership</code> • <code>RoleMemberAttributes</code> • <code>AuthenticationFilter</code> • <code>DigestMD5Authentication</code> • <code>UseUserAccountControl</code>
ProviderURL	<code>ldap://localhost:389</code>	<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> <p><code>ldap://<hostname>:<port></code></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 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.
SecurityProtocol	None	<p>The protocol to be used when connecting to the LDAP server.</p> <p>To use an encrypted protocol, use "ssl" instead "ldaps" in the url.</p> <hr/> <p>Note: ActiveDirectory requires the SSL protocol when setting the value for the password attribute. This occurs when creating a user or updating the password of an existing user.</p> <hr/>
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"> • simple — For clear-text password authentication. • DIGEST-MD5 — For more secure hashed password authentication. This method requires that the server use plain text password storage and only works with JRE 1.4 or later.

Property	Default Value	Description
AuthenticationFilter	<p>For most LDAP servers: (& (uid={uid}) (object- class=person))</p> <p>or</p> <p>For Active Directory email lookups: (&(user- Principal- Name={uid}) (object- class=user)) [ActiveDirec- tory]</p> <p>For Active Directory Windows username lookups: (& (sAMAccount- Name={uid}) (object- class=user))</p>	<p>The filter to use when looking up the user.</p> <p>When performing a username based lookup, this filter is used to determine the LDAP entry that matches the supplied username.</p> <p>The string "{uid}" in the filter is replaced with the supplied username.</p>
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>
AuthenticationSearchBase	none	<p>The search base used to authenticate users. If this value is not specified, the LDAP DefaultSearch-Base is used.</p>

Property	Default Value	Description
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 also use this DN to create the users in the LDAP server. When the self-registration feature is used, this user may also need the requisite permissions to create a user record. This behavior can occur if you do not set <code>useUserCredentialsToBind</code> to <code>true</code>. In this case, the LDAP attributer uses this DN to update the user attributes.</p>
BindPassword	none	<p>BindPassword is the password for BindDN, which is used to authenticate any user. BindDN and BindPassword are used to separate the LDAP connection into units.</p> <p>The <code>AuthenticationMethod</code> property determines the bind method used for this initial connection.</p> <p>Sybase recommends encrypting passwords and provides a password encryption utility for the purpose. If you encrypt BindPassword, include <code>encrypted=true</code> in the line that sets the option. For example:</p> <pre><options name="BindPassword" encrypted="true" value="lsnjikf-wregfqr43hu5io..." /></pre> <p>If you do not encrypt BindPassword, the option might look like this:</p> <pre><options name="BindPassword" value="s3cr3T" /></pre>
RoleSearchBase	none	<p>The search base used to retrieve lists of roles. If this value is not specified, the LDAP Default-SearchBase is used.</p>

Property	Default Value	Description
RoleFilter	<p>For SunONE/iPlanet: (<code>&i (object-class=ldapsu-bentry) (objectclass=nsroledefinition)</code>)</p> <p>For Netscape Directory Server: (<code>object-class=groupof-names) (object-class=groupofunique-names)</code>)</p> <p>For ActiveDirectory: (<code>object-class=groupof-names) (object-class=group)</code>)</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 or this property is not initialized, no roles are available.</p>
RoleMemberAttributes	<p>For Netscape Directory Server: <code>member,unique-member</code></p>	<p>The role's member attributes defines a comma-delimited list of attributes that roles may have that define a list of DN's of people who are in the role.</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 only has a default value when the Netscape server type is chosen.</p>
RoleNameAttribute	<p><code>cn</code></p>	<p>The attribute for retrieved roles that is the common name of the role. If this value is "dn" it is interpreted specially as the entire dn of the role as the role name.</p>
RoleScope	<p><code>onelevel</code></p>	<p>The role search scope. The supported values for this are:</p> <ul style="list-style-type: none"> • <code>onelevel</code> • <code>subtree</code> <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
UserRoleMembershipAttributes	For iPlanet/SunONE: nsRoleDN For ActiveDirectory: memberOf For all others: none	The user's role membership attributes property is used to define an attribute that a user has that contains the DN's of all of the roles as user is a member of. These comma-delimited values are then cross-referenced with the roles retrieved in the role search base and search filter to come up with a list of user's roles.
UserFreeformRoleMembershipAttributes	None	The "freeform" role membership attribute list. Users who have attributes in this comma-delimited list are automatically granted access to roles whose names are equal to the attribute value. For example, if the value of this property is "department" and user's LDAP record has the following values for the department attribute, { "sales", "consulting" }, then the user will be granted roles whose names are "sales" and "consulting".
Referral	ignore	The behavior when a referral is encountered. The valid values are those dictated by LdapContext, for example, "follow", "ignore", "throw".
DigestMD5AuthenticationFormat	DN For OpenLDAP: User-name	The DIGEST-MD5 bind authentication identity format.
UseUserAccountControlAttribute	For most LDAP servers: false For ActiveDirectory: true	The UserAccountControl attribute to be used for detecting disabled user accounts, account expirations, password expirations and so on. ActiveDirectory also uses this attribute to store the above information.
controlFlag	optional	Indicates whether authentication with this login module is sufficient to allow the user to log in, or whether the user must also be authenticated with another login module. Rarely set to anything other than "sufficient" for any login module. Note: controlFlag is a generic login module option rather than an LDAP configuration property.

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

Login Module	OS Group	Sybase Control Center Roles
	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.

Get Started

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 88

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
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

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

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

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

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

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

Examples

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

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

Set only the database messaging service to port 3639.

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

Set the HTTP port to 9292.

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

Set the Jini RMI daemon to port 9696.

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

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

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

Set the RMI port to 9991.

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

Set the Tabular Data Stream port to 9997.

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

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

See also

- *scc Command* on page 78

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**.

Get Started

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

- *Alerts* on page 175
- *Replication Data Collections* on page 122
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123
- *Launching Sybase Control Center* on page 67
- *Logging in to Sybase Control Center* on page 82

Configuring the Automatic Logout Timer

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

Prerequisites

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

Task

1. From the 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

- *Logging out of Sybase Control Center* on page 83

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 116

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

- *Alerts* on page 175
- *Replication Data Collections* on page 122
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123
- *Removing a Role from a Login or a Group* on page 104

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 103

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 105
- *Adding a Login Account to a Group* on page 106
- *Removing a Login Account from a Group* on page 106

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 105

Get Started

- *Adding a Login Account to a Group* on page 106
- *Removing a Login Account from a Group* on page 106

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 105
- *Removing a Group* on page 105
- *Removing a Login Account from a Group* on page 106

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 105
- *Removing a Group* on page 105
- *Adding a Login Account to a Group* on page 106

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

Get Started

- 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.

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.

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**.

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 12. 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 13. 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.

Get Started

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.

Configure

Configure login accounts, statistics collection, and other replication monitoring options.

1. *Registering a Replication Server, Replication Agent, or Mirror Replication Agent*
Register a resource (for example, a server that can be monitored) to make Sybase Control Center aware of it and its connection information.
2. *Importing Resources for Batch Registration*
(Optional) Import and register multiple servers from an interfaces or sql.ini file.
3. *Creating a Perspective*
Create a perspective in which you can add and manage resources.
4. *Adding a Resource to a Perspective*
Add one or more resources to the current perspective.
5. *Role Assignment in Sybase Control Center for Replication*
In Sybase Control Center for Replication version 3.2.3, you no longer need to grant special roles for administrative or monitoring privileges on the Replication Server.
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. *Setting Up the Latency Monitoring Process*
Start the latency monitoring process for replication paths in the replication environment.
8. *Setting Up Statistics Collection*
Use the Properties view of your managed resource to create a data collection job and add a schedule to the job.
9. *Setting Replication Parameters*
Configure replication parameters to improve server performance. Sybase Control Center for Replication allows you to configure the parameters for Replication Server, Replication Agent, connection and logical connection, route, and Adaptive Server RepAgent thread.
10. *Creating an Alert*
Use the Add Alert wizard to create an alert instance for your resource.
11. *Optional Configuration Steps*
Perform additional configuration, including user authorization, alerts, data collection scheduling, backups, and setting purging options for the repository.

See also

- *User Authorization* on page 55
- *Logins, Roles, and Groups* on page 109

Configure

- *Setting Up Security* on page 83
- *Assigning a Role to a Login or a Group* on page 103

Registering a Replication Server, Replication Agent, or Mirror Replication Agent

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

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

Table 14. New Resource Type Details

Field	Description
Resource Name	(Required) Name of the resource to register. Enter the actual name of the server, using uppercase and lowercase letters. If the name registered in Sybase Control Center does not exactly match the server name, some monitoring functions, including the topology view, do not work.
Resource Type	Select a resource type: <ul style="list-style-type: none">• ASE Server, Replication Only – monitor only the RepAgent threads for an Adaptive Server that is older than version 15.0.2.0. Choose this type for an Adaptive Server that is part of a replication environment.• Replication Agent – monitor Replication Agent.• Replication Server – monitor Replication Server.
Description	A brief description to help you identify the resource.

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

Table 15. New resource connection details

Field	Description
Server Host Name/Host Name	Local host name
Port Number	Local host port number
RSSD Host Name	(Required) Name of the RSSD host
RSSD Port Number	(Required) RSSD host port number

Note: **RSSD Host Name** and **RSSD Port Number** are required only if you are registering a Replication Server.

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

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

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

6. Click **Next**.
7. (Optional) Click **Add this resource to the current perspective**. You must add a resource to a perspective (not necessarily the current perspective) before you can manage or monitor it.
8. (Optional) Click **Open the resource explorer to view this new resource**. (This option is not present when the Resource Explorer is open.)
9. Click **Finish**.

See also

- *Resources* on page 183
- *Compatibility Matrix* on page 3
- *Common Display Options* on page 6
- *Unregistering a Resource* on page 184

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`

Configure

For more information on `interfaces` files, see the appendix on configuration files in *Configuration Guide Open Client and Open Server 15.0 for UNIX*.

For more information on `sql.ini` files, see the chapter on network communications using `sql.ini` in the Adaptive Server Enterprise 15.0 *Configuration Guide for Windows*.

Note: The Import Resources wizard imports servers in batches of a single type (Adaptive Server, Sybase IQ, or Replication Server, for example). If your `interfaces` or `sql.ini` file includes resources of more than one type, you must perform this procedure for each resource type.

1. In the application menu, select **View > Open > Resource Explorer**.
2. In the Resource Explorer, select **Resources > Import**.
The Import Resources wizard opens; **Interfaces file** is already selected.
3. Click **Next**.
The Directory Service Connection page appears.
4. Click **Browse** and navigate to the `interfaces` file you want to import from.
You cannot type in the **File name** field.
5. Click **Next**.
6. On the Import Resource Type page, select the type of server you want to import.
7. On the Resource Selection page, click to select the servers you want to import.
Select only servers of the type you chose on the Import Resource Type page. If you import servers with incorrect types, Sybase Control Center will not be able to monitor or manage them properly.
8. Resources of your chosen type may require connection parameters in addition to those present in the file—RSSD host name and port for Replication Server, for example, or character set and language for Adaptive Server. Enter any required connection parameters.
9. Click **Next**.
10. (Optional) Click **Add these resources to the current perspective**. You must add a resource to a perspective (not necessarily the current perspective) before you can manage or monitor it.
11. Click **Next**.
The Confirmation page displays a list of the resources you have selected.
12. Click **Finish** if you are ready to import, or click **Back** to return to the previous screens and change your selections.
When you click **Finish**, Sybase Control Center imports and registers the resources and displays a summary page.
13. Click **Close** to finish the wizard.

The newly imported resources appear in the Resource Explorer. If you elected to add them to the current perspective, the resources also appear in the Perspective Resources view.

See also

- *Resources* on page 183
- *Unregistering a Resource* on page 184

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**.

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

- *Removing a Resource from a Perspective* on page 185

Role Assignment in Sybase Control Center for Replication

In Sybase Control Center for Replication version 3.2.3, you no longer need to grant special roles for administrative or monitoring privileges on the Replication Server.

Sybase Control Center automatically assigns repMonitorRole and repAdminRole privileges to users who have sa permission, and repAdminRole privileges to users who have **create**

Configure

object permission on Replication Server. The same roles are automatically assigned to users when they successfully authenticate a Replication Agent resource.

Sybase Control Center automatically assigns `aseMonitorRole` privileges to users who have `mon_role`, and `aseAdministratorRole` privileges to users who have `sa_role` privileges on an Adaptive Server Replication Only resource. After successful resource authentication, these assigned roles let users perform monitoring and administrative tasks on the server.

Sybase Control Center validates roles every 30 minutes. The Monitor option is unavailable from the Perspective Resources window until the resource is authenticated.

Authenticating a Login Account for a Managed Resource

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

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

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

1. Connect a browser to Sybase Control Center and log in.
2. If the Perspective Resources view is not open, click the **Show/Hide Perspective Resources View** icon in the toolbar.
3. In the Perspective Resources view, select your resource and select **Resource > Authenticate** from the view menu.
4. Select **Use my current SCC login** or **Specify different credentials**.
5. If you chose **Specify different credentials**, enter the login and password for Sybase Control Center to use to connect to your resource.
6. If the selected resource is a Replication Server, also enter the RSSD user name and password.
7. Click **OK** to save and exit the dialog.

See also

- *Setting Up the Latency Monitoring Process* on page 117
- *User Authorization* on page 103

Roles or Permissions for Setting Up Replication Monitoring Accounts

Sybase Control Center uses the user name and password account that you provide when you authenticate a server to retrieve monitoring information.

The user account you have specified must have these roles or permissions:

- Replication Server – a **create object** or sa permission.
- Replication Server System Database (RSSD) database – read privilege on all tables and create and execute privileges on stored procedures in the RSSD database. It is recommended to use the account created by **rs_init** for the RSSD.
- Replication Agent or Mirror Replication Agent – an administrative account. There is only one administrative account in a Replication Agent.
- Adaptive Server for RepAgent threads – a system administrator role. Sybase Control Center uses the **dbcc monitor** command to retrieve statistics.

Note: sa or dbo permission are required for setting parameters and controlling data flow.

Encrypted Authentication for Replication

Sybase Control Center uses encrypted passwords to connect to a Replication Server, an Adaptive Server, an Adaptive Server that hosts a Replication Server System Database (RSSD), and a SQL Anywhere that hosts an embedded RSSD that are configured for network password encryption.

If the Replication Server, Adaptive Server, RSSD, or embedded RSSD is configured to require login encryption, Sybase Control Center establishes a connection using a password that is encrypted during network transmission.

Note: Password encryption is not supported in the Replication Agent for systems administrator password, and in a non-Sybase data server that retrieves latency and **rs_ticket** performance data.

For information on setting up encrypted password on Replication Server and Adaptive Server, see *Replication Server Configuration Guide > Password Encryption* and *Adaptive Server System Administration Guide: Volume 1 > Setting Configuration Parameters*.

Setting Up the Latency Monitoring Process

Start the latency monitoring process for replication paths in the replication environment.

Prerequisites

- Install and configure all the servers to be monitored. Verify that these servers have been registered and added to the Perspective Resources, and that you have the permission to perform this task.

Note: All resources in the path must be registered and authenticated for latency monitoring to take place. If ASE servers are not registered and authenticated, connection information cannot be obtained for latency monitoring.

- The primary and replicate Replication Servers must be at version 15.1 or later. For a non-Sybase primary database, the Replication Agent must also be version 15.1 or later.

Configure

- To use **rs_ticket** in a replication environment to measure the latency of non-Sybase replicate databases, you must install JDBC drivers to enable Sybase Control Center for Replication to connect to those non-Sybase replicate databases. See *Sybase Control Center Installation Guide > Installing the JDBC Drivers for Non-Sybase Replicate Databases*.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Paths**.
3. Select a path from the list.

Note: When selecting or highlighting a path, click outside the Path Name column. Clicking the path name opens the path dashboard.

4. Click **Quick Admin** to display the latency monitoring process administration view for the selected path.

Note: If the selected path's primary and replicate Replication Servers, or the Replication Agent for a non-Sybase primary database, are not at version 15.1 or later, the **Quick Admin** option is disabled.

All replicate databases to which the selected path replicates data are listed in the administration view. For Adaptive Server databases, all parameters are populated automatically. For non-Sybase databases, you need to provide the parameters manually in the administration view.

5. Specify the replicate database information, start and end dates, and the repeat interval time.
6. For non-Sybase databases, click **Test Connection** to verify that the connection information is valid.

Repeat this step for all the listed replicate databases.

7. Click **Start**.

You can click the bar graph icon in the upper-right corner of the Latency window in the Replication Server Monitor's Paths view or in the Replication Path Dashboard to see the detailed latency history.

See also

- *Authenticating a Login Account for a Managed Resource* on page 116
- *Setting Up Statistics Collection* on page 119
- *Stopping the Latency Monitoring Process* on page 231
- *Displaying Latency* on page 230
- *Latency Chart Details* on page 231

- *Replication Latency* on page 230
- *Data on Dashboards or Charts Is Missing* on page 233

Latency Monitoring Process

Latency monitoring means executing the **rs_ticket** command at the primary database, retrieving the latency data from the replicate database, and storing the results in the Sybase Control Center repository.

rs_ticket gathers time information at various points in the replication path as it flows from the primary database to the replicate database. You can then use this latency information to identify bottlenecks in the replication path. Latency information helps ensure that data is flowing efficiently between the primary and replicate databases.

The latency monitoring process collects data on a specified schedule for each replicate Replication Server and for the primary Replication Server for the path you configured. The primary collection executes the **rs_ticket** command in the primary database. Each replicate collection performs a collection in all its paths. When a path collection occurs, the results of the **rs_ticket** command are retrieved from the replicate database, stored in the repository, and displayed in the latency history chart in the Replication Server Monitor's Paths view or in the Replication Path Dashboard for server performance monitoring.

Each collection repeats until the end date or until you manually stop the latency monitoring process. Once the process is stopped, all collections are deleted from the scheduler.

See also

- *Displaying Replication Path Status* on page 229
- *Displaying Latency* on page 230
- *Stopping the Latency Monitoring Process* on page 231

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**.

Configure

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.
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

- *Setting Up the Latency Monitoring Process* on page 117
- *Setting Replication Parameters* on page 144
- *Job Scheduling* on page 170
- *Viewing or Deleting a Schedule* on page 173

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.

Configure

- 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

- *Replication Data Collections* on page 122
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123

Replication Data Collections

Lists and describes predefined data collections you can use to create jobs for Replication.

Replication Objects	Replication Data Collection	Description
Replication Server	Monitoring Statistics (default collection*) <hr/> Note: This collection must be scheduled to enable alerts.	Collects the availability metrics displayed in the Replication Server Monitor and heat chart.
	Performance Counters	A diagnostic tool that collects historical performance metrics that you can later graph in the statistics chart.
ASE Server, Replication Only	RepAgent Statistics (default collection*) <hr/> Note: This collection must be manually started if the Adaptive Server is registered as an ASE Server type. "collection_ase_rat" is the name of the RepAgent Statistics collection in an ASE Server registered Adaptive Server.	Collects availability metrics for RepAgent thread.
Replication Agent	RA Monitoring Statistics (default collection*)	Collects availability metrics displayed in the Replication Agent Monitor and heat chart.
	Performance Counters	Collects historical performance metrics that you can later graph in the statistics chart.

Replication Objects	Replication Data Collection	Description
*Default collections are automatically created upon authentication and are set to repeat indefinitely at a 60-second interval.		

See also

- *About Statistics* on page 121
- *Alerts* on page 175
- *Creating an Alert* on page 148
- *Assigning a Role to a Login or a Group* on page 103
- *Configuring the E-mail Server* on page 101
- *Job Scheduling* on page 170

Key Performance Indicators and Key Performance Areas for Replication

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

Each Replication KPA includes a subset of the KPIs.

Table 16. Key Performance Indicators

KPI	Description	Data Collection
Managed Object States		
Primary State	Managed object primary state.	Monitoring Statistics
Secondary State	Managed object secondary state.	
Replication Server Performance Statistics		
Insert commands	Insert command on active object.	Performance Counters
Update commands total	Update command on active object.	
Delete commands total	Delete command on active object.	
Writetext commands total	Writetext command on active object.	
Execute commands total	Execute command on active object.	

Configure

KPI	Description	Data Collection
AIO collect time total	Amount of time that the Asynchronous I/O Daemon spent in collecting async I/O requests in which Daemon is obliged to poll until notified of their completion.	
AIO after poll total	Number of pending asynchronous I/O requests just after the Async I/O Daemon finishes polling for their completion.	
AIO before poll total	Number of pending asynchronous I/O requests just before the Async I/O Daemon goes off to poll for their completion.	
AIO poll time total	Amount of time that the Asynchronous I/O Daemon spent in polling I/O completions. This includes the time taken to notify AIO clients of their I/O requests completion.	
AIO read retry total	The number of times an async read is made and the OpSys tells the client that it must try again.	
AIO write retry total	The number of times an async write is made and the OpSys tells the client that it must try again.	
Alarm milliseconds spent total	Number of milliseconds actually spent before the alarm was expired.	
Alarm cancel total	Number of alarms cancelled.	
Alarm examined total	Number of examined alarms to determine whether they are expired or not.	
Alarm expired cancelled total	Number of alarms that were found canceled upon each 'handle expiration' execution.	
Alarm expired total	Number of alarms that have expired upon each 'handle expiration' execution.	

KPI	Description	Data Collection
Alarm milliseconds expired interval total	Number of milliseconds until the next alarm expiration.	
Alarm expired time total	Amount of time elapsed while handling expired alarms.	
Alarm milliseconds requested total	Number of milliseconds originally requested for the alarm when it was set.	
Alarm set total	Number of alarms set.	
Thread start total	Total number of started threads.	
Thread stop total	Number of threads stopped.	
Begins total	Begin Tran statements sent to Replication Server System Database (RSSD).	
Deletes total	Delete statements sent to RSSD.	
Inserts total	Insert statements sent to RSSD.	
Selects total	Select statements sent to RSSD.	
Cache exceed total	Number of time System Table Services (STS) cached was exceeded.	
Updates total	Update statements sent to RSSD.	
Replication Server		
Replication Server State	The state of the Replication Server.	Monitoring Statistics
% Memory usage	The percent of the current memory used by the server.	
Maximum Latency	The highest latency of any replication path in the server.	
% Partition Usage	The percent full of all partitions in the Replication Server.	
Path		
Path Status	The state of the replication path between a primary and a replicate database.	Monitoring Statistics

Configure

KPI	Description	Data Collection	
Path Latency	The time between when a transaction is committed at the primary database and when it is committed at the replicate database.		
Change in Path Activity	The number of commands sent to the replicate database by the Data Server Interface (DSI) thread since the last monitoring cycle.		
Path Data Loss	Indicates possible data loss between the primary and replicate databases.		
rs_ticket Count	The ticket's unique sequence number.		
Primary Database Start Time	The time, in milliseconds, the rs_ticket was executed at the primary database.		
Executor Thread Latency	The number of milliseconds from the start, to the time the ticket passes through the executor thread.		
Distributor Thread Latency	The number of milliseconds from the primary Replication Server time, to the time the ticket passes through the distributor thread.		
Replication Server Interface Latency	The number of milliseconds from the distributor (DIST) thread time, to the time the ticket passes through the Replication Server Interface (RSI) thread.		
Data Server Interface Thread	The number of milliseconds from the RSI thread time, to the time the ticket passes through the DSI thread.		
Path RDB Latency	The number of milliseconds from the DSI thread time, to the time the ticket arrived at the replicate database.		
Connection			
Time closing an OB fadeout total	The amount of time spent closing a connection marked to be faded out.		Performance Counters
Time waiting on OB fadeout total	The amount of time spent waiting for a connection marked to be faded out.		

KPI	Description	Data Collection
Commands read total	Commands read from an inbound queue by a DIST thread.	
Resolved delete statements total	Delete commands encountered by a DIST thread and resolved by subscription resolution engine (SRE).	
Unresolved statements total	DIST commands with no subscription resolution that a DIST thread discarded. This implies either there is no subscription or the 'where' clause associated with the subscription does not result in row qualification.	
Resolved insert statements total	Insert commands that a DIST thread encountered, which SRE resolved.	
Resolved update statements total	Update commands that a DIST thread encountered, which SRE resolved.	
TD Begins total	Begin transaction commands that a DIST thread propagated.	
TD Commits or Rollbacks total	Commit or Rollback commands that a DIST thread processed.	
Transactions processed total	Transactions read by DIST thread from an inbound queue.	
All large threads used total	This counter is incremented each time a Large Parallel Transaction must wait because there are no available parallel DSI threads.	
All threads used total	This counter is incremented each time a Parallel Transaction must wait because there are no available parallel DSI threads.	
Commands applied successfully total	Commands that a DSI successfully applied to the target database.	

Configure

KPI	Description	Data Collection
Time dispatching large groups total	Time spent by the Data Server Interface, scheduler thread (DSI/S) dispatching a large transaction group to a DSI/E. This includes the time spent finding a large group to dispatch.	
Time dispatching regular groups total	Time that the DSI/S spent dispatching a regular transaction group to a DSI/E.	
Time finding regular groups total	Time that the DSI/S spent finding a group to dispatch.	
Time loading SQT cache total	Time that the DSI/S spent loading Stable Queue Transaction (SQT) cache.	
Time executing special group total	Time that the DSI/S spent determining if a transaction is special, and executing it if it is.	
DSI/E in sleep mode total	Number of DSI/E threads that the DSI/S put to sleep prior to loading SQT cache. These DSI/E threads have just completed their transaction.	
Time putting DSI/E in sleep mode total	Time that the DSI/S spent putting free DSI/E threads to sleep.	
Transaction groups read total	Transaction groups that the DSI read. If grouping is disabled, grouped and ungrouped transaction counts are the same.	
Ungrouped transactions read total	Ungrouped transactions that the DSI read. If grouping is disabled, grouped and ungrouped transaction counts are the same.	
Transaction groups committed total	Transactions that a DSI thread successfully committed.	

KPI	Description	Data Collection
Transaction groups sent total	Transaction groups that the DSI thread sent to the target database. A transaction group can contain at most dsi_max_xacts_in_group transactions. This counter is incremented each time a 'begin' for a grouped transaction is executed.	
Transaction groups success total	Transaction groups that a DSI thread successfully applied to a target database. This includes transactions that were successfully committed or rolled back according to their final disposition.	
Transaction groups committed total	Transactions in groups sent by a DSI thread that committed successfully.	
Transaction groups sent total	Transactions contained in transaction groups that a DSI thread sent.	
Closed groups maximum bytes total	Transaction groups that a DSI thread closed due to the next transaction that caused the dsi_xact_group_size to exceed.	
Closed groups next new origin total	Transaction groups that a DSI closed due to no open group from the origin of the next transaction. For example, there is a new origin in the next transaction, or the schedule forced a flush of the current group from the origin leaving no open group from that origin.	
Closed groups partitioning total	Transaction groups that a DSI thread closed because of a Transaction Partitioning rule.	
Closed groups maximum log transaction total	Transaction groups that a DSI thread closed due to the next transaction that caused the dsi_max_xacts_in_group to exceed.	

Configure

KPI	Description	Data Collection
Groups wait partitioning total	Transaction groups forced to wait for another group to complete (processed serially based on Transaction Partitioning rule).	
True return of rs_dsi_check_thread_lock total	Number of rs_dsi_check_thread_lock invocations returning true. The function determined the calling thread holds locks required by other threads. A roll-back and retry occurred.	
Delete statements read total	rs_delete commands that a DSI executor (DSIEXEC) thread processed.	
Batch size total	Size, in bytes, of commands batches, which a DSI submitted.	
Batch time total	Time, in milliseconds, to process command batches, which a DSI submitted.	
Bytes succeed total	Bytes that a DSI/E successfully applied to the target database.	
Commands read total	Commands read from an outbound queue by a DSIEXEC thread.	
Commands succeed total	Commands that a DSI/E successfully applied to the target database.	
Function string mapping time total	Time, in milliseconds, to perform function string mapping on commands.	
Input command count total	Number of input commands in command batches that a DSI submitted.	
Output command count total	Number of output commands in command batches that a DSI submitted.	
Result time total	Time, in milliseconds, to process the results of commands batches that a DSI submitted.	
Sequencing check on command batches time total	Time, in milliseconds, to check the sequencing on command batches, which required some kind of synchronization such as wait_for_commit .	

KPI	Description	Data Collection
Sequencing check on commits time total	Time, in milliseconds, to check the sequencing on commits.	
Transaction time total	Time, in milliseconds, for a DSI/E thread to process transactions. This includes function string mapping, sending and processing results. A transaction may span command batches.	
Errors deadlock total	Times that a DSI thread failed to apply a transaction due to deadlocks in the target database (ASE Error 1205).	
Errors out of lock total	Times that a DSI thread failed to apply a transaction due to no locks available in the target database (ASE Error 1204).	
Writetext total	rs_writetext commands that a DSIEXEC thread processed.	
Insert total	rs_insert commands that a DSIEXEC thread processed.	
MemUsed group total	Memory that a DSI/S thread consumed for transaction groups.	
Send RPC time total	Time, in milliseconds, spent in sending RPCs to the RDS.	
Send time total	Time, in milliseconds, spent in sending command buffers to the RDS.	
System transaction read total	Internal system transactions that a DSIEXEC thread processed.	
Transaction groups scheduled total	Transactions groups scheduled to a DSIEXEC thread.	
Ungrouped transactions scheduled total	Transactions in transaction groups scheduled to a DSIEXEC thread.	
Updates read total	rs_update commands that a DSIEXEC thread processed.	
Bytes received total	Bytes that a RepAgent thread received. This size includes the TDS header size when in 'passthru' mode.	

Configure

KPI	Description	Data Collection
Mini abort total	mini-abort commands (in Adaptive Server, SAVEXACT records) that a RepAgent thread processed. mini-abort instructs Replication Server to rollback commands to a specific OQIQ value.	
Commands received total	Commands that a RepAgent thread received.	
Write wait time total	The amount of time the RepAgent spent waiting for the Stable Queue Manager (SQM) Writer thread to drain the number of outstanding write requests to get the number of outstanding bytes to be written under the threshold.	
Yield time total	The amount of time the RepAgent spent yielding the processor while handling Log Transfer Language (LTL) commands each time the processor was yielded.	
Execution time total	The amount of time, in milliseconds, RepAgent user thread is scheduled by Open Client/Server (OCS).	
Receive time total	The amount of time, in milliseconds, spent receiving network packets or language commands.	
Update rs_locator total	Updates to RSSD. rs_locator where type = 'e' executed by a RepAgent thread.	
Replication Agent Thread State	The state of the Replication Agent thread in the Replication Server that is associated with the connection.	Monitoring Statistics
DSI State	The state of the DSI thread in the Replication Server that is associated with the connection.	
DSI Activity	The number of commands that the DSI thread sent to the replicate database.	

KPI	Description	Data Collection
DSI Activity per Second	The number of commands that the DSI thread sent to the replicate database per second.	
Transaction Ignored	The number of transactions that were ignored.	
Transaction Skipped	The number of transactions that were skipped.	
Transaction Succeeded	The number of transactions that were successfully committed to the replicate database.	
Transaction Failed	The number of transactions that were failed.	
Transaction Retried	The number of transactions that were retried.	
DSI SQT Cache Percent Used	The percent of DSI SQT cache size used.	
Transactions removed per second	The number of transactions removed from the DSI SQT cache per second.	
SQT cache low boundry	The smallest size to which DSI SQT cache is configured before transactions start being removed from the cache.	
Transactions read per second	The number of transactions that the DSI read per second.	
Commands per Transaction	The number of commands per transaction.	
Transactions per Group	The number of transactions per group.	
Commands per Batch	The number of commands per batch.	
Logical Connection		
Operation in Progress State	The state of the operation in progress.	Monitoring Statistics
Route		
Route State	The state of the route at both the origin and destination Replication Servers.	Monitoring Statistics

Configure

KPI	Description	Data Collection
Route Activity	The number of bytes sent to the destination Replication Server.	
Route Activity per Second	The number of bytes sent to the destination Replication Server per second.	
Alarm time spent total	Number of milliseconds actually spent before an alarm expired.	Performance Counters
Time spent cancelling alarms total	Amount of time elapsed while cancelling an alarm.	
Expired alarms to be cancelled total	Number of alarms that were found to have been expired while attempting to cancel them.	
Alarm required time total	Number of milliseconds originally requested for the alarm when it was set.	
Alarm set total	Amount of time elapsed while setting an alarm.	
Connection active time total	Amount of time elapsed while a thread makes a connection active.	
Connection charset time total	Amount of time elapsed while checking a connection's character set.	
Connection fadeout freetime total	Amount of time elapsed while this thread's free connections are marked for fadeout. Note that this thread is probably <i>*not*</i> the thread doing the work. Most likely it is the Alarm Handler doing the work.	
Connection fadeout time total	Amount of time elapsed while this thread's assigned connections are marked for fadeout. Note that this thread is probably <i>*not*</i> the thread doing the work. Most likely it is the Alarm Handler doing the work.	
Connection give-up total	Number of times the thread gave up a connection.	
Connection give-up time total	Amount of time elapsed while a thread gives up a connection.	

KPI	Description	Data Collection
Connection idle time total	Amount of time elapsed while a thread makes a connection idle.	
Connection open total	Number of times the thread made a connection request and received one that had to be opened.	
Connection request total	Number of times the thread made a connection request and received one that was already open.	
Connection sort order time total	Amount of time elapsed while checking a connection's sort order.	
Connection time total	Amount of time elapsed while a thread obtains a connection.	
Message queue get time total	Number of centiseconds spent getting a message from an internal message queue.	
Message queue put time total	Number of centiseconds spent putting a message on an internal message queue.	
Sleep time total	Amount of time elapsed after a thread puts itself to sleep.	
Yield time total	Amount of time elapsed after a thread yields the CPU.	
Blocks read total	Number of blocking (SQM_WAIT_C) reads that an RSI thread performed against SQM thread that manages a RSI queue.	
Bytes sent total	Bytes that an RSI sender thread delivered.	
Messages sent total	RSI messages that an RSI thread sent. These messages contain the distribute command.	
Packets sent total	Packets that an RSI sender thread sent.	
Send packet time total	Time, in milliseconds, spent in sending packets of data to the RRS.	

Configure

KPI	Description	Data Collection
Bytes received total	Bytes that an RSI user thread received. This size includes the TDS header size when in 'passthru' mode.	
Command length total	Length of an RSI command.	
Commands received total	Commands that an RSI user thread received. Includes RSI message, get truncation requests, etc.	
Execution time total	The amount of time, in milliseconds, RSI User thread is scheduled by OCS.	
Receive packet time total	The amount of time, in milliseconds, spent receiving network packets.	
GetTrunc send time total	The amount of time, in milliseconds, spent responding to 'get truncation' requests.	
SetTrunc send time total	The amount of time, in milliseconds, spent responding to 'set truncation' requests.	
Queue		
Queue Writes	The number of commands written to the queue in the Replication Server.	Monitoring Statistics
Queue Writes per Second	The number of commands written to the queue in the Replication Server per second.	
Queue Reads	The number of commands read from the queue in the Replication Server.	
Queue Reads per Second	The number of commands read from the queue in the Replication Server per second.	
Queue Bytes	The number of bytes written to the queue in the Replication Server.	
Queue Bytes per Second	The number of bytes written to the queue in the Replication Server per second.	

KPI	Description	Data Collection
Queue Size	The size of the queue in the Replication Server.	
Queue SQT Cache Percent Used	The percent of SQT cache size used for inbound queue.	
Duplicates Transaction during monitoring interval	The number of duplicate commands detected and ignored by the SQM.	
Transactions removed per second	The number of transactions removed from the SQT cache per second.	
SQT cache low boundry	The smallest size to which SQT cache would be configured before transactions start being removed from cache.	
Commands processed by DIST per second	The number of commands that the distributor thread processed for an inbound queue per second.	
Commands processed without rep def per second	The number of commands that did not have an associated replication definition, which the distributor thread processed for an inbound queue per second.	
Percent of blocks that were full when written to the queue	The percentage of blocks that were full when written to the queue.	
Percent of reads that came directly from the cache	The percentage of reads that came directly from the cache.	
Blocks full write total	Number of full blocks that an SQM thread wrote. Individual blocks can be written due either to block full state or to sysadmin command slow_queue (only one message per block).	
Blocks written total	Number of 16K blocks that an SQM thread wrote to a stable queue.	
Bytes written total	Bytes that an SQM thread wrote to a stable queue.	
Commands written total	Commands that an SQM thread wrote to a stable queue.	

Configure

KPI	Description	Data Collection
Command size total	Command size written to a stable queue.	
Segments active total	Active segments of an SQM queue: the number of rows in rs_segments for the given queue where used_flag = 1 .	
Queue segments allocated total	Segments allocated to a queue during the current statistical period.	
Queue segments deallocated total	Segments deallocated from a queue during the current statistical period.	
Time new segments allocated total	The elapsed time, in milliseconds, to allocate a new segment. Timer starts when a segment is allocated. Timer stops when the next segment is allocated.	
Updates to rs_oqid total	Updates to the RSSD. rs_oqid table by an SQM thread. Each new segment allocation may result in an update of rs_oqid value stored in rs_oqid for recovery purposes.	
Write requests total	Message writes that an SQM client requested.	
Write timer pop total	SQM writer thread initiated a write request due to timer expiration.	
Blocks read total	Number of 16 blocks that an SQM Reader (SQMR) thread read from a stable queue.	
Blocks read cached total	Number of 16K blocks that an SQMR thread read from a cache.	
Commands read total	Commands that an SQMR thread read from a stable queue.	
Time waiting for queue write total	The amount of time SQMR wait for a queue write.	
Unread blocks total	The number of blocks within a partially read segment that are yet to be read.	

KPI	Description	Data Collection
Unread segments total	The number of segments yet to be read.	
Read time total	The amount of time taken for SQMR to read a block.	
Memory used total	SQT thread memory use. Each command structure allocated by an SQT thread is freed when its transaction context is removed. For this reason, if no transactions are active in SQT, SQT cache usage is zero.	
Transactions added closed total	Transactions added to the Closed queue.	
Transactions removed closed total	Transactions removed from the Closed queue.	
Commands read total	Commands read from SQM. Commands include XREC_BEGIN , XREC_COMMIT , XREC_CHECKPT .	
Commands in tran total	Commands in transactions completely scanned by an SQT thread.	
Transactions added total	Transactions added to the Open queue.	
Transactions removed total	Transactions removed from the Open queue.	
Transactions added Read Q total	Transactions added to the Read queue.	
Transactions removed Read Queue total	Transactions removed from the Read queue.	
Cache lower bound size total	The smallest size to which SQT cache can be configured before transactions start being removed from cache.	
Closed transactions total	Current closed transaction count.	
Open transactions total	Current open transaction count.	
Read transactions total	Current read transaction count.	
Truncation queue transaction total	Current truncation queue transaction count.	

Configure

KPI	Description	Data Collection
Transactions removed total	Transactions whose constituent messages have been removed from memory. Removal of transactions is most commonly caused by a single transaction exceeding the available cache.	
Truncation queue transactions added total	Transactions added to the Truncation queue.	
Truncation queue transactions removed total	Transactions removed from the Truncation queue	
Partition		
Partition State	The state of the partition in the Replication Server.	Monitoring Statistics
Partition Usage	The percent of the partition in the Replication Server that is currently in use.	
RepAgent		
Replication Agent State	The state of the Replication Agent.	RA Monitoring Statistics
RepAgent VM		
VM maximum memory	Maximum memory (in bytes) available to the Java Virtual Machine (VM).	Performance Counters
VM total memory allocated	Total memory (in bytes) allocated to the Java VM at start-up.	
VM free memory	Memory (in bytes) allocated but not used by the Java VM.	
VM memory usage	Memory (in bytes) allocated and in use by the Java VM.	
VM % max memory used	Percentage of the the maximum memory available to the Java VM, currently in use by the Java VM.	
RepAgent LTI Statistics		
Number of LTL commands sent	Total number of LTL commands sent to Replication Server.	Performance Counters

KPI	Description	Data Collection
LTL Commands Sent per Second	LTL commands sent to Replication Server per second.	
Avg LTL command size	Average size (in bytes) of each LTL command sent to Replication Server.	
Avg LTL commands/sec	Average number of LTL commands sent per second to Replication Server.	
Total bytes sent	Number of bytes sent to Replication Server.	
Avg Bytes/second during transmission	Average bytes per second sent over connection to Replication Server.	
Avg LTL buffer cache time	Average time (in milliseconds) it takes between placing the LTL commands into the LTL buffer to the time it is actually sent to the Replication Server.	
Avg Rep Server turnaround time	Average time (in milliseconds) it takes Replication Server to acknowledge each LTL command buffer sent.	
Avg data arrival time	Average time (in milliseconds) log transfer interface (LTI) waits between receiving change sets from Log Reader.	
Avg time to create distributes	Average time (in milliseconds) LTI takes to convert a change-set into LTL.	
Avg LTL buffer size	Average size (in bytes) of each LTL buffer sent to Replication Server	
Avg LTM buffer utilization (%)	Average utilization (in percentage of LTL buffer size) of each LTL buffer sent to Replication Server.	
Avg LTL commands/buffer	Average number of LTL commands per buffer sent to Replication Server.	
LTI encoded column name cache size	Encoded column name cache size.	
LTI input queue size	Number of change sets in the LTI input queue	

Configure

KPI	Description	Data Collection
LTI output queue size	Current number of LTL commands in the LTI output queue.	
RepAgent Log Reader Statistics		
Total operations scanned	Number of operations read from the database transaction log.	Performance Counters
Operations Scanned per Second	Number of operations read from the database transaction log per second.	
Total operations processed	Number of operations processed from the database transaction log.	
Operations Processed per Second	Number of operations process from the database transaction log per second.	
Total operations skipped	Number of operations read from the database transaction log and skipped for any reason.	
Operations Skipped per Second	Number of operations read from the database transaction log and skipped for any reason per second.	
Current operation queue size	Current size (in bytes) of the Log Reader input queue.	
Current scan buffer size	Current size (in bytes) of the Log Reader scan buffer.	
Total maintenance user operations filtered	Number of maintenance-user operations read from the database transaction log and skipped.	
Total Maintenance User Operations Filtered per Second	Number of maintenance-user operations read from database transaction log and skipped per second.	
RepAgent LTM Statistics		
Items held in Global LRU Cache	Number of object references in the internal Least Recently Used cache.	Performance Counters
Adaptive Server RepAgent Thread		
ASE Replication Agent Thread Status	The status of the Adaptive Server RepAgent thread in the data server.	RepAgent Statistics

KPI	Description	Data Collection
ASE Replication Agent Thread Log	The difference between the current marker and the end of the log.	
ASE Replication Agent Thread Operations Scanned	The number of operations scanned from the Adaptive Server transaction log since the last monitoring cycle.	
ASE Replication Agent Thread Commands Processed	The number of commands processed and sent to the Replication Server since the last monitoring cycle.	

Table 17. Key Performance Areas

KPA	Description
KPA_RS_PERF_STATS	Replication Server common statistics.
KPA_RS_HEATCHART	Replication Server heat chart statistics.
kpa_path	Replication path statistics.
KPA_PATH_RS_TICKET	The latency times at various points in the path between the primary database and the replicate database.
kpa_connection	Replication Server connection statistics.
KPA_CONNECTION_PERF_STATS	Replication Server connection performance statistics.
kpa_logical_connection	Replication Server logical connection statistics.
kpa_route	Replication Server route statistics.
KPA_ROUTE_PERF_STATS	Replication Server route performance statistics.
kpa_queue	Replication Server queue statistics.
KPA_QUEUE_PERF_STATS	Replication Server queue performance statistics.
kpa_partition	Replication Server partition statistics.
Availability Statistics	Replication Agent availability statistics.
VM Statistics	Replication Agent performance statistics on JAVA VM component.
LTM component Statistics	Replication Agent performance statistics on log transfer manager (LTM) component.
LTI component Statistics	Replication Agent performance statistics on LTI component.

Configure

KPA	Description
LR component Statistics	Replication Agent performance statistics on log reader (LR) component.
kpa_ase	Adaptive Server statistics
kpa_rat_ctrl_availability	Adaptive Server RepAgent thread controller availability statistics.
kpa_rat_availability	Adaptive Server RepAgent thread availability statistics.

See also

- *Historical Performance Monitoring* on page 168
- *Heat Chart* on page 167
- *Replication Latency* on page 230
- *Availability Monitoring* on page 213
- *About Statistics* on page 121
- *Alerts* on page 175
- *Creating an Alert* on page 148
- *Assigning a Role to a Login or a Group* on page 103
- *Configuring the E-mail Server* on page 101

Setting Replication Parameters

Configure replication parameters to improve server performance. Sybase Control Center for Replication allows you to configure the parameters for Replication Server, Replication Agent, connection and logical connection, route, and Adaptive Server RepAgent thread.

See also

- *Setting Up Statistics Collection* on page 119
- *Creating an Alert* on page 148

Configuring Replication Server and Replication Agent Parameters

Configure Replication Server and Replication Agent parameters.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a server and click **Resource > Monitor**.
2. Click **Quick Admin**.
3. In the Configure tab, modify the current value of the server parameter in the Run Value column.

Note: Use the filter text box and check box options to search for the parameter you want to modify.

4. Click **Apply All**.
5. (Optional) Click **Reset All** to undo modifications that have not been applied.

Note: Not all of Replication Server and Replication Agent parameters can be configured. If a parameter cannot be configured, the text box in the administration view is greyed out.

Configuring Connection Parameters for Replication Server

Configure connection parameters that affect database connections originating from the Replication Server.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Connections**.
3. Select a connection and click **Quick Admin**.
4. In the Configure tab, modify the current value of the connection parameter in the Run Value column.

Note: Use the filter text box and check box options to search for the parameter you want to modify.

5. Click **Apply All**.
6. (Optional) Select **Recycle connection** before you click **Apply All** to allow Replication Server to suspend, then resume connection.

Note: This option allows the new parameters to be applied to the connection. If you click **Apply All** without selecting **Recycle connection**, the parameters are not applied to the connection until the connection is restarted.

7. (Optional) Click **Reset All** to undo modifications that have not been applied.

Configuring Logical Connection Parameters for Replication Server

Configure Replication Server parameters that apply to logical connections for warm standby applications.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Logical Connection**.
3. Select a logical connection and click **Quick Admin**.
4. In the Configure tab, modify the current value of the logical connection parameter in the Run Value column.

Note: Use the filter text box and check box options to search for the parameter you want to modify.

5. Click **Apply All**.
6. (Optional) Click **Reset All** to undo modifications that have not been applied.

Configuring Route Parameters for Replication Server

Configure parameters that affect routes from the current Replication Server to other Replication Servers.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Routes**.
3. Select a route and click **Quick Admin**.
4. In the Configure tab, modify the current value of the route parameter in the Run Value column.

Note: Use the filter text box and check box options to search for the parameter you want to modify.

5. Click **Apply All**.
6. (Optional) Select **Recycle route** before you click **Apply All** to allow Replication Server to suspend, then resume the route.

Note: This option allows the new parameters to be applied to the route. If you click **Apply All** without selecting **Recycle route**, the parameters are not applied to the route until the route is restarted.

7. (Optional) Click **Reset All** to undo modifications that have not been applied.

Configuring Adaptive Server RepAgent Thread Parameters

Configure Adaptive Server RepAgent thread parameters that affect replication performance.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select an Adaptive Server and click **Resource > Monitor**.
2. Select a RepAgent thread.
3. Click **Quick Admin**.
4. In the Configure tab, modify the current value of the RepAgent thread parameter in the Run Value column.

Note: Use the filter text box and check box options to search for the parameter that you want to modify.

5. Click **Apply All**.
6. (Optional) Click **Reset All** to undo modifications that have not been applied.

Note: Not all of Adaptive Server RepAgent thread parameters can be configured. If a parameter cannot be configured, the text box in the administration view is greyed out.

Configuring Replication Server to Gather Zero Values for Counters

Set the Replication Server parameter to gather zero values for performance counters when you run a collection job.

Prerequisites

Add all the servers to be monitored in the Perspective Resources, and verify that you have permission to perform this task.

Configure

Task

If you want zero values to be gathered for performance counters, run this command in the Replication Server:

```
configure replication server set stats_show_zero_counters to "on"
```

Creating an Alert

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

Prerequisites

- You must have administrative privileges (sccAdminRole) to perform this task.
- Specify an e-mail server for Sybase Control Center to use for alerts. You cannot create e-mail subscriptions to alerts without an e-mail server.
- Schedule data collections. Alerts for each product module are based on one or more data collections. If the correct collection or collections are not scheduled to run, the alert system cannot function and no alerts are generated. See the data collections topic for your product module for information on which collections you need to schedule to enable alerts.
- (Optional) If you want this alert to trigger the execution of a shell script, copy the script to a location on or accessible from the machine that hosts your Sybase Control Center server. Set permissions to make the script executable.

Warning! Use caution in writing scripts. A poorly designed script can cause a blocking situation, creating a deadlock in your Sybase Control Center server.

Task

1. In the Perspective Resources view, click the server or other resource and select **Resource > Properties** in the view's menu bar.
2. Select **Alerts** in the left pane and click **Add**.
The Add Alert Wizard opens. If the selected resource supports child alerts, the wizard opens to the Resource page. If the resource does not support child alerts, the wizard opens to the Type page—in that case, skip to step 5.
3. On the Resource page of the wizard, select the object on which to set the alert. Expand the folder representing the server or agent to select lower-level child objects.
4. Click **Next**.
5. On the Type page, select the alert type and click **Next**.
For this step and the next one, see the topic on key performance indicators for information on what this alert monitors and how it is triggered. (Each alert is based on a KPI.)
6. Based on the type of alert you selected, do one of the following:

- For a state-based alert – select a severity level for each alert state.

Note: You can associate only one severity level with each state.

- For a threshold-based alert – review and if necessary adjust the range of values that defines each severity.

7. Click **Next**.

8. (Optional) Enter the storm suppression period. Storm suppression blocks redundant alert notifications and script executions resulting from the same condition for the specified period of time. Enter this value in seconds, minutes, or hours and click **Next**.

9. (Optional) To configure this alert to trigger the execution of a script:

- a) **Alert Severity** specifies the severity level that triggers the script. Select **Critical**, **Warning**, or both.

Critical is typically more serious than Warning.

- b) Browse to the location of the script.

Note: In UNIX, make sure the script is executable. You cannot select a script unless it has execute permission.

- c) If the script requires parameter values, click **Select Parameters** to enter them in the **Execution Parameters** box.

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

Note: When you test a script, Sybase Control Center supplies test values for the **%Severity%** and **%Source_Application%** parameters (“Testing” and “TestScriptExecution,” respectively). Any test values you supply for these parameters are discarded. This prevents the test results from being confused with real script results after testing and in the SCC repository.

- d) (Optional) Click **Test** to perform a test execution of your script.

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

- e) Click **Next**.

If the selected resource has sibling resources (databases or devices of the same type, for example) that support this alert type, you see the Duplicates page. If the selected resource has no identical siblings, you see the Subscription page.

10. (Optional) On the Duplicates page, select any resources that should use this alert definition as a template for their own alerts. Click the box at the top of the list to select all the resources listed. Then click **Next**.

This step saves time when you need to configure similar alerts for several resources of the same type.

Configure

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 18. Alert subscription details

Option	Description
E-mail	To send an e-mail notification when this alert fires, click the E-mail Message box and enter the e-mail address of one user or list.
Escalation E-mail	To escalate this alert (by sending another e-mail notification if this alert has not been responded to after a specified period of time), click the Escalation E-mail box and enter the e-mail address of one user or list. You cannot enter an escalation address unless you enter an address for primary notification first.
Time Period	Specify how long to wait, following the initial alert notification, before Sybase Control Center sends an e-mail notification to the escalation address. (The same notification is sent again to the original notification address.) Select a time unit (hours, minutes, or seconds) and enter a number.

12. Click **Finish**.

If you are creating duplicate or child alerts, the **Cancel** button is activated; click it to interrupt the creation of further alerts. (The primary alert, at a minimum, is always created before the operation can be cancelled.) If you do not want to keep the duplicate or child alerts (if any) created before you cancelled the operation, drop them manually.

Note: Click **Cancel** to stop the creation of duplicate alerts.

Examples: Alert-triggered scripts

This sample script is a Windows .bat file. It outputs the parameter values you pass to it to a text file. Windows batch files support only nine arguments. (Arg0, the name of the script, is not counted.)

```
@echo off
@echo. >> stest.txt
@echo %date% %time% >> stest.txt
@echo arg0: %0 >> stest.txt
@echo arg1: %1 >> stest.txt
@echo arg2: %2 >> stest.txt
@echo arg3: %3 >> stest.txt
@echo arg4: %4 >> stest.txt
@echo arg5: %5 >> stest.txt
```

```
@echo arg6: %6 >> stest.txt
@echo arg7: %7 >> stest.txt
@echo arg8: %8 >> stest.txt
@echo arg9: %9 >> stest.txt
@echo. >> stest.txt
```

This is a sample execution parameter string for the script above:

```
Time:%Time%
Severity:%Severity%
Resource:%Resource%
Server:%Top_resource%
KPI:%KPI%
State:%Current_state%
URL:%SCC_URL%
```

The script's 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:
```

This is a UNIX script. It also outputs the parameter values you pass to it to a text file.

```
#!/bin/sh
outfile=c:/testing/latest/scriptTest.out
echo> $outfile
echo `date` >> $outfile
count=1
while [ "$1" ]
do
  echo arg$count: $1 >> $outfile
  shift
  count=`expr $count + 1`
done
echo --- DONE --- >> $outfile
```

See also

- *Setting Replication Parameters* on page 144
- *Optional Configuration Steps* on page 164
- *Alerts* on page 175
- *Replication Data Collections* on page 122
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123
- *Testing an Alert-Triggered Script* on page 178

Replication Alerts

Lists and describes alerts you can use to create alert instances for Replication.

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

Replication Objects	Alerts	Description	Alert Type	Data Collection Name
Replication Server	State Change	The alert to send when a Replication Server state changes.	State	Monitoring Statistics
	Memory Usage	The alert to send when a Replication Server memory usage is above the configured threshold.	Threshold	
	Partition Usage	The alert to send when a Replication Server partition's usage is above the configured threshold.	Threshold	
Paths	Path State Change	The alert to send when a replication path's state changes.	State	
	Change in Path Activity	The alert to send when the numbers of operations sent on a replication path is above the configured threshold.	Threshold	
Connections	Connection State Change	The alert to send when a Replication Server connection for the primary database changes state.	State	
	DSI Activity	The alert to send when the numbers of operations sent on a Replication Server DSI connection is above the configured threshold.	Threshold	
	DSI Activity per Second	The alert to send when the numbers of operations sent on a Replication Server DSI connection per second is above the configured threshold.	Threshold	

Replication Objects	Alerts	Description	Alert Type	Data Collection Name
Logical Connections	Logical Connection State Change	The alert to send when a Replication Server logical connection changes state.	State	
Routes	Route State Change	The alert to send when a Replication Server route's state changes.	State	
	Route Activity	The alert to send when the change in number of bytes sent through a Replication Server route is above the configured threshold.	Threshold	
	Route Activity per Second	The alert to send when the number of bytes sent through a Replication Server route per second is above the configured threshold.	Threshold	
Queues	Number of Commands Written	The alert to send when the number of commands written to a Replication Server queue is above the configured threshold.	Threshold	
	Number of Commands Written per Second	The alert to send when the number of commands written to a Replication Server queue per second is above the configured threshold.	Threshold	
	Number of Commands Read	The alert to send when the number of commands read from a Replication Server queue is above the configured threshold.	Threshold	
	Number of Commands Read per Second	The alert to send when the number of commands read from a Replication Server queue per second is above the configured threshold.	Threshold	

Configure

Replication Objects	Alerts	Description	Alert Type	Data Collection Name	
	Size of Commands Written	The alert to send when the number of commands bytes written to a Replication Server queue is above the configured threshold.	Threshold		
	Size of Commands Written per Second	The alert to send when the number of bytes written to a Replication Server queue is above the configured threshold.	Threshold		
	Queue Size	The alert to send when the size of a Replication Server queue is above the configured threshold.	Threshold		
Partitions	Partition Usage	The alert to send when a Replication Server partition's usage is above the configured threshold.	Threshold		
Replication Agent and Mirror Replication Agent	State Change	The alert to send when a Replication Agent state changes.	State		RA Monitoring Statistics
	Memory Usage	The alert to send when a Replication Agent memory usage is above the configured threshold.	Threshold		
	LTI Input Queue Size	The alert to send when the number of change sets in a Replication Agent Log Transfer Interface (LTI) input queue is above the configured threshold.	Threshold		
	LTI Output Queue Size	The alert to send when the number of change sets in a Replication Agent LTI output queue is above the configured threshold.	Threshold		

Replication Objects	Alerts	Description	Alert Type	Data Collection Name
	Number of LTL Commands Sent	The alert to send when the number of Replication Agent Log Transfer Language (LTL) commands sent to Replication Server is above the configured threshold.	Threshold	
	Number of LTL Commands Sent per Second	The alert to send when the number of Replication Agent LTL commands sent to Replication Server per second is above the configured threshold.	Threshold	
	Number of Operations Scanned	The alert to send when the number of operations scanned by a Replication Agent Log Reader is above the configured threshold.	Threshold	
	Number of Operations Scanned per Second	The alert to send when the number of operations scanned by a Replication Agent Log Reader per second is above the configured threshold.	Threshold	
	Number of Operations Processed	The alert to send when the number of operations processed by a Replication Agent Log Reader is above the configured threshold.	Threshold	
	Number of Operations Processed per Second	The alert to send when the number of operations processed by a Replication Agent Log Reader per second is above the configured threshold.	Threshold	

Configure

Replication Objects	Alerts	Description	Alert Type	Data Collection Name
	Number of Operations Skipped	The alert to send when the number of operations skipped by a Replication Agent Log Reader is above the configured threshold.	Threshold	
	Number of Operations Skipped per Second	The alert to send when the number of operations skipped by a Replication Agent Log Reader per second is above the configured threshold.	Threshold	
	Number of Maintenance-User Operations Skipped	The alert to send when the number of maintenance-user operations skipped by a Replication Agent Log Reader is above the configured threshold.	Threshold	
	Number of Maintenance-User Operations Skipped per Second	The alert to send when the number of maintenance-user operations skipped by a Replication Agent Log Reader per second is above the configured threshold.	Threshold	
	Log Reader Queue Size	The alert to send when the size (in bytes) of a Replication Agent Log Reader input queue is above the configured threshold.	Threshold	
	Log Reader Scan Buffer Size	The alert to send when the size (in bytes) of a Replication Agent Log Reader scan buffer is above the configured threshold.	Threshold	
Adaptive Server Enterprise	State Change	The alert to send when an Adaptive Server state changes.	State	Monitoring Statistics
Adaptive Server RepAgent Thread	RepAgent Thread State Change	The alert to send when an Adaptive Server RepAgent Thread changes state.	State	RepAgent Statistics

Replication Objects	Alerts	Description	Alert Type	Data Collection Name
	Transaction Log Size	The alert to send when the size of an Adaptive Server RepAgent Thread's transaction log is above the configured threshold.	Threshold	
	Number of Log Operations Scanned per Second	The alert to send when the number of log operations scanned per second by an Adaptive Server RepAgent Thread is above the configured threshold	Threshold	
	Number of Log Operations Processed per Second	The alert to send when the number of log operations processed per second by an Adaptive Server RepAgent Thread is above the configured threshold.	Threshold	

See also

- *Alert Types, Severities, and States for Replication* on page 157
- *Substitution Parameters for Scripts* on page 163
- *Alert-Triggered Scripts* on page 162

Alert Types, Severities, and States for Replication

Learn about the properties that define and control alerts.

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

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.

Configure

Severity	Description
Normal	No alert is issued.
Warning	A problem has given cause for concern. An alert is issued; you can subscribe to alerts that fire at the Warning level.
Critical	A serious problem exists. An alert is issued; you can subscribe to alerts that fire at the Critical level.

Alert states are used on the dashboards, heat chart, and topology icon decorators.

Replication Object	Object State	Alert State	Description
Replication Server	Active	Running	The Replication Server is running and actively replicating data.
	Quiesce		The Replication Server is running, but is not currently replicating data.
	Suspect	Warning	At least one of the Replication Server connections, routes, or queues is down.
	Rebuilding		The Replication Server is rebuilding queues. This state is returned by the admin health command.
	Hibernating		The Replication Server is in hibernation mode. This state is returned by the admin health command.
	Recovering		The Replication Server is in stand-alone mode and is rebuilding queues. This state is returned by the admin health command.
	Standalone		The Replication Server is in stand-alone mode. This state is returned by the admin health command.
	Down	Stopped	The Replication Server is unable to communicate with the RSSD. The server state is also set to DOWN if the user or password is incorrect.

Replication Object	Object State	Alert State	Description
	Unknown	Unknown	The state of the Replication Server cannot be determined. UNKNOWN can also indicate that the server is not part of the replication environment.
Replication Server Connection	Active	Running	The Replication Server DSI and RepAgent thread are actively replicating.
	Hidden		
	Suspended	Stopped	The Replication Server DSI or RepAgent thread is DOWN or SUSPENDED.
	Unknown	Unknown	The RepAgent for a primary connection is not part of the replication environment.
Replication Server Logical Connection	Active	Running	The RepAgent for the active physical connection and the DSI thread for the standby physical connection are both active.
	Suspended	Stopped	The logical connection can be suspended for these reasons: <ul style="list-style-type: none"> The active or standby physical connection is not defined for the logical connection. The RepAgent thread for the active connection is suspended. The DSI thread for the standby connection is suspended. The logical connection is in the process of switching the active and standby databases.
	Unknown	Unknown	The RepAgent thread for the active connection is unknown, or the DSI thread for the standby connection is unknown.

Configure

Replication Object	Object State	Alert State	Description
Replication Server Route	Active	Running	The route is open and data can pass from the origin to the destination Replication Server.
	Suspended	Stopped	The route is unavailable and data cannot pass between the Replication Servers. The description provides the reason for the suspended route; for example: <ul style="list-style-type: none"> • The route encountered an internal error. • The route is being created. • The route is suspended. • The route encountered an error at the destination. • The route is being dropped. • The route is being dropped with NOWAIT. • An indirect route is being changed to a direct route.
	Unknown	Unknown	The destination Replication Server is not part of the replication environment.
Replication Path	Active	Running	All objects along the Replication Path are actively replicating data.
	Suspect	Stopped	One or more objects along the Replication Path are not replicating data.
	Unknown	Unknown	The state of the objects in the Replication Path cannot be determined.
Replication Agent	Replicating	Running	Successfully connected to the Replication Agent. The agent is in replicating state. This state is returned by the ra_status command.

Replication Object	Object State	Alert State	Description
	Replicating (Resynchronization)	Running	The Replication Agent has been restarted and is resynchronizing the primary and replicate databases. This state is returned by the ra_status command.
	Admin	Stopped	Successfully connected to the Replication Agent. The agent is in administration state and is not currently replicating data. This state is returned by the ra_status command.
	Replication Down	Stopped	Cannot connect to the Replication Agent or the agent state is set to REPLICATION DOWN . This state is returned by the ra_status command.
	Unknown	Unknown	The state of the Replication Agent cannot be determined. Also may indicate that the agent is not part of the replication environment.
Adaptive Server RepAgent Thread	Active	Running	The RepAgent thread is enabled and started.
	Suspended	Stopped	The RepAgent thread is enabled, but stopped.
	Down	Stopped	The RepAgent thread is not enabled.
Replication-Only Adaptive Server	Active	Running	Successfully connected to the Adaptive Server and all RepAgent threads for connections within this environment are enabled and started.
	Down	Stopped	Cannot connect to the Adaptive Server. The server state is also set to DOWN if the user and password are incorrect.

Replication Object	Object State	Alert State	Description
	Unknown	Unknown	The state of the Replication-only Adaptive Server cannot be determined. This also may indicate that the server is not part of the replication environment.

See also

- *Replication Alerts* on page 152
- *Substitution Parameters for Scripts* on page 163
- *Alert-Triggered Scripts* on page 162

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\`

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

- *Replication Alerts* on page 152
- *Alert Types, Severities, and States for Replication* on page 157
- *Substitution Parameters for Scripts* on page 163

- *Testing an Alert-Triggered Script* on page 178
- *Alerts* on page 175

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 19. Substitution Parameters for State-Based Alerts

Parameter	Description
%Alert%	A three-part name supplied by the alert system. The parts are the name of this alert, the name of the resource, and the name of the key performance indicator (KPI) on which this alert is based.
%Current_state%	The current state of the resource on which this alert is configured.
%KPI%	The name of the KPI on which this alert is based.
%Resource%	The name of the resource with which this alert is associated.
%SCC_URL%	A link to Sybase Control Center, where more information about the alert may be available.
%Severity%	The severity of this alert: critical or warning.
%Source_application%	The SCC product module that generated this alert.
%Time%	The date and time at which the alert fired, in this format: Tue Sep 15 10:10:51 2009
%Server%	The name of the alerted resource's top-level parent resource—usually the server. This is valuable when the alerted resource is a component of a larger system (a database in a server, for example). If the alerted resource has no parent, %Server% and %Resource% have the same value.

Table 20. 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.

Configure

Parameter	Description
%Datapoint%	The current value, on the alerted resource, of the KPI on which this alert is based.
%KPI%	The name of the KPI on which this alert is based.
%Resource%	The name of the resource with which this alert is associated.
%SCC_URL%	A link to Sybase Control Center, where more information about the alert may be available.
%Severity%	The severity of this alert: critical or warning. (Critical is more serious.)
%Source_application%	The SCC product module that generated this alert.
%Threshold%	The threshold value at which this alert fires.
%Time%	The date and time at which the alert fired, in this format: Tue Sep 15 10:10:51 2009
%Server%	The name of the alerted resource's top-level parent resource. This is valuable when the alerted resource is a component of a larger system (a database in a server, for example). If the alerted resource has no parent, %Server% and %Resource% have the same value.

See also

- *Replication Alerts* on page 152
- *Alert Types, Severities, and States for Replication* on page 157
- *Alert-Triggered Scripts* on page 162
- *Testing an Alert-Triggered Script* on page 178
- *Modifying an Alert* on page 177

Optional Configuration Steps

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

Table 21. Configuration areas

Configuration area	Description	Topic
User authorization	Set up groups of users or assign roles. Make sure there are users with administrative privileges (sccAdminRole).	<i>User Authorization</i> on page 103
Authentication	Add authentication modules to allow Windows, UNIX, and LDAP users to log in to Sybase Control Center.	<i>Setting up Security</i> on page 83
Alerts	Modify alert thresholds and subscriptions and delete alerts.	<i>Alerts</i> on page 175
Data collection	Modify collection intervals and schedules, suspend and resume the schedule, and delete collection jobs.	<i>Job Scheduling</i> on page 170
Resources	Unregister resources, add them to perspectives, or remove them.	<i>Resources</i> on page 183
Perspectives	Create, remove, and rename perspectives.	<i>Perspectives</i> on page 186
Instances	Enable or disable shared-disk mode and deploy, remove, refresh, or convert SCC agent or server instances running from a shared disk.	<i>Instances</i> on page 189
Repository	Set purging options and schedule backups of the repository database.	<i>Repository</i> on page 199

See also

- *Creating an Alert* on page 148

Configure

Manage and Monitor

Manage and Monitor the servers and their components in the replication environment.

Heat Chart

The heat chart displays status and availability statistics for managed resources in the current perspective.

The heat chart displays the state of resources in your perspective—whether the resources are running, suspended, or down. In addition, the heat chart lists the type of each resource and provides statistical data, including the start time of the last data collection.

In the Perspective Heat Chart view, you can filter the resources that you want to see and search and sort the results by column. You can also select a resource and pull down its context menu to see monitoring and administrative options that vary based on the resource type.

Heat chart data is collected directly from managed servers, tagged with the date and time when it was collected, and stored in the Sybase Control Center repository.

See also

- *Historical Performance Monitoring* on page 168
- *Replication Latency* on page 230
- *Availability Monitoring* on page 213
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123

Displaying Resource Availability

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

1. From the application menu bar, select **View > Open > Heat Chart**.
2. (Optional) To display information about the status represented by an icon in the chart, hover the mouse over the icon.
 - Status column – icon tooltips describe the status of the resource (Running or Stopped, for example).
 - All columns to the right of Status – icon tooltips give the value of the KPI listed at the top of the column.
3. (Optional) To display tools for filtering (narrowing the list of resources in the heat chart) or changing the columns, select **View > Filter** from the Perspective Heat Chart menu bar. The Filter and Column tools appear in the left pane.

4. (Optional) To use filtering, select **View > Filter** from the view's menu bar and enter a search term in the **Filter string** field.
The search term can be any string that appears in the tabular portion of the heat chart, such as the name, or part of the name, of a server or a resource type (ASE Server, for example).
5. (Optional) Select a filtering setting:
 - **Match case** – search for resources whose displayed data includes the search term, including uppercase and lowercase letters; or
 - **Exact match** – search for resources whose displayed data includes an item identical to the search term.
6. (Optional) Select a column from the **Filter on** list to restrict your search to that column.
7. (Optional) Click **Columns** to customize your heat chart.
8. (Optional) Unselect any column that should not appear in your heat chart.
9. (Optional) Click the sorting arrow in the column headers to sort the column values in either ascending or descending order.
10. (Optional) Click the resource's row and pull down the menu to the right of the resource name to view options for the selected resource.
11. (Optional) To resize the Filter and Columns tools pane, move your mouse over the border between the tools pane and the resource table. When the mouse cursor changes to a resize icon, click and drag the border to the left or the right.
12. (Optional) To hide the Filter and Columns tools, unselect **View > Filter**.

See also

- *Graphing Performance Counters* on page 169
- *Configuring Repository Purging* on page 204

Historical Performance Monitoring

Monitor performance data to determine whether your environment is working efficiently.

Obtain detailed information about the status of the resources in your environment. You can create performance graphs that illustrate resource performance over a specified period of time.

See also

- *Heat Chart* on page 167
- *Replication Latency* on page 230
- *Availability Monitoring* on page 213
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123

Graphing Performance Counters

To show performance trends, generate a graph for any set of performance counters.

Prerequisites

Verify that statistical data to be graphed has been collected. To verify data collection, go to the Collection Jobs page of the Resource Properties view and check the History tab for a collection job. You can also look at the resource monitor: if data appears there, data is being collected.

Task

Tip: Data collections start running when a resource is authenticated. A recently authenticated resource might not have accumulated enough data to make a useful graph.

1. In the Perspective Resources view, click a resource and select **Resource > Launch Statistics Chart** in the view menu bar.
2. Expand the folders in the Statistics tab and select the key performance indicator (KPI) you want to graph.
3. Click **Graph Statistic** or drag the KPI onto the Chart tab.
The Chart tab displays the graphed data, while the KPI with its corresponding value and the date and time it was collected appear in the Data tab.
4. (Optional) Repeat to add KPIs to the graph.
5. (Optional) Use the slider at the bottom of the Chart tab to control the amount of time covered by the graph, ranging from a minute to a year.
6. (Optional) Use <<, <, >, and >> to move the displayed graph to an earlier or later time, depending on how the slider is set.

Tip: The statistics chart displays data covering a fixed period of time, and that period does not change automatically. If you are viewing the most recent statistics and want to keep the graph current, adjust the displayed time period as new statistics are collected.

7. (Optional) You can click the date/time labels that appear above the slider. Use these to change the start and end time and the chart time span.
8. (Optional) Click **Clear Graph** to remove all the graphed statistics and start anew.

Note: You can graph a maximum of five statistics with no more than two distinct units of measure. By default, only 24 hours of statistics are available; change the repository purge options to save statistics for a longer period.

See also

- *Configuring Repository Purging* on page 204
- *Displaying Resource Availability* on page 167

Manage Sybase Control Center

Manage and monitor Sybase Control Center for Replication through performance statistics and component activity logs.

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 119
- *Replication Data Collections* on page 122

Executing and Stopping a Data Collection Job

Use the Properties view to execute or stop a data collection job.

Most of the time, data collection jobs should run on a schedule; you should rarely need to start or stop a job manually.

1. In the Perspective Resources view, select the resource associated with the job and select **Resource > Properties**.
2. Select **Collection Jobs**.
3. Select the job and:
 - To execute a job immediately, click **Execute**.
 - To stop a job, click **Stop**, then click **Yes** to confirm.

See also

- *Deleting a Data Collection Job* on page 171
- *Resuming and Suspending a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172
- *Viewing or Deleting a Schedule* on page 173
- *Modifying the Data Collection Interval for a Job* on page 174
- *Resuming and Suspending the Scheduler* on page 174
- *Viewing the Job Execution History* on page 175

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 170
- *Resuming and Suspending a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172
- *Viewing or Deleting a Schedule* on page 173
- *Modifying the Data Collection Interval for a Job* on page 174
- *Resuming and Suspending the Scheduler* on page 174
- *Viewing the Job Execution History* on page 175

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 170
- *Deleting a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172
- *Viewing or Deleting a Schedule* on page 173
- *Modifying the Data Collection Interval for a Job* on page 174
- *Resuming and Suspending the Scheduler* on page 174
- *Viewing the Job Execution History* on page 175

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 170
- *Deleting a Data Collection Job* on page 171
- *Resuming and Suspending a Data Collection Job* on page 171
- *Viewing or Deleting a Schedule* on page 173
- *Modifying the Data Collection Interval for a Job* on page 174
- *Resuming and Suspending the Scheduler* on page 174
- *Viewing the Job Execution History* on page 175

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 170
- *Deleting a Data Collection Job* on page 171
- *Resuming and Suspending a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172
- *Modifying the Data Collection Interval for a Job* on page 174
- *Resuming and Suspending the Scheduler* on page 174
- *Viewing the Job Execution History* on page 175
- *Setting Up Statistics Collection* on page 119

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 170
- *Deleting a Data Collection Job* on page 171
- *Resuming and Suspending a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172
- *Viewing or Deleting a Schedule* on page 173
- *Resuming and Suspending the Scheduler* on page 174
- *Viewing the Job Execution History* on page 175

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 170
- *Deleting a Data Collection Job* on page 171
- *Resuming and Suspending a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172

- *Viewing or Deleting a Schedule* on page 173
- *Modifying the Data Collection Interval for a Job* on page 174
- *Viewing the Job Execution History* on page 175

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 170
- *Deleting a Data Collection Job* on page 171
- *Resuming and Suspending a Data Collection Job* on page 171
- *Adding a New Schedule to a Job* on page 172
- *Viewing or Deleting a Schedule* on page 173
- *Modifying the Data Collection Interval for a Job* on page 174
- *Resuming and Suspending the Scheduler* on page 174

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

- *Replication Data Collections* on page 122
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123
- *Creating an Alert* on page 148
- *Assigning a Role to a Login or a Group* on page 103
- *Configuring the E-mail Server* on page 101
- *Alert-Triggered Scripts* on page 162

Types, Severities, and States

Learn about the properties that define and control alerts.

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

Table 22. 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 23. 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 177
- *Modifying an Alert* on page 177
- *Testing an Alert-Triggered Script* on page 178
- *Deleting an Alert* on page 179
- *Alert Subscriptions* on page 179
- *Alert Notifications* on page 182

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 176
- *Modifying an Alert* on page 177
- *Testing an Alert-Triggered Script* on page 178
- *Deleting an Alert* on page 179
- *Alert Subscriptions* on page 179
- *Alert Notifications* on page 182

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 176
- *Viewing Alerts* on page 177
- *Testing an Alert-Triggered Script* on page 178
- *Deleting an Alert* on page 179
- *Alert Subscriptions* on page 179
- *Alert Notifications* on page 182

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 176
- *Viewing Alerts* on page 177
- *Modifying an Alert* on page 177
- *Deleting an Alert* on page 179
- *Alert Subscriptions* on page 179
- *Alert Notifications* on page 182
- *Alert-Triggered Scripts* on page 162
- *Substitution Parameters for Scripts* on page 163
- *Creating an Alert* on page 148

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 176
- *Viewing Alerts* on page 177
- *Modifying an Alert* on page 177
- *Testing an Alert-Triggered Script* on page 178
- *Alert Subscriptions* on page 179
- *Alert Notifications* on page 182

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 176
- *Viewing Alerts* on page 177
- *Modifying an Alert* on page 177
- *Testing an Alert-Triggered Script* on page 178
- *Deleting an Alert* on page 179
- *Alert Notifications* on page 182

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 24. 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 181
- *Enabling and Disabling Alert Subscription* on page 181

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 180
- *Enabling and Disabling Alert Subscription* on page 181

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:

Manage and Monitor

- 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 180
- *Unsubscribing from an Alert* on page 181

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 176
- *Viewing Alerts* on page 177
- *Modifying an Alert* on page 177
- *Testing an Alert-Triggered Script* on page 178
- *Deleting an Alert* on page 179
- *Alert Subscriptions* on page 179

Displaying Alert History and Resolutions

Use the Properties view to see historical information about resolved and unresolved alerts.

The History tab on the Alerts page of the Resource Properties view displays information about every time this alert has fired. Each row of the table represents a single notification generated by the selected alert.

The Resolutions tab displays information about alerts that have been resolved (closed) by a Sybase Control Center administrator.

The History and Resolutions tabs display the 100 most recent alerts or alerts for the last 24 hours, whichever is reached first.

1. In the Perspective Resources view, select a resource and select **Resource > Properties**.
2. Select **Alerts**.
3. Select the alert instance.
4. Click the **History** tab.

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

See also

- *Resolving Alerts* on page 183

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 182

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

Manage and Monitor

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

- *Compatibility Matrix* on page 3
- *Common Display Options* on page 6
- *Registering a Replication Server, Replication Agent, or Mirror Replication Agent* on page 112
- *Importing Resources for Batch Registration* on page 113

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 184
- *Removing a Resource from a Perspective* on page 185
- *Searching for Resources in the Resource Explorer* on page 186
- *Registering a Replication Server, Replication Agent, or Mirror Replication Agent* on page 112
- *Importing Resources for Batch Registration* on page 113

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 184
- *Removing a Resource from a Perspective* on page 185
- *Searching for Resources in the Resource Explorer* on page 186
- *Registering a Replication Server, Replication Agent, or Mirror Replication Agent* on page 112
- *Importing Resources for Batch Registration* on page 113

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 184
- *Adding a Resource to a Perspective* on page 184
- *Searching for Resources in the Resource Explorer* on page 186
- *Adding a Resource to a Perspective* on page 115

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 184
- *Adding a Resource to a Perspective* on page 184
- *Removing a Resource from a Perspective* on page 185

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 187
- *Renaming a Perspective* on page 187

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 187
- *Renaming a Perspective* on page 187

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 187
- *Removing a Perspective* on page 187

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.

Manage and Monitor

Each perspective includes these views:

- Perspective Resources
- Administration Console
- Heat chart
- Alert monitor
- Component log viewer
- Views that exist for each managed resource. These vary by resource type, but typically include the statistics chart, the properties view, and a monitoring view.

Managing a View

Open, close, minimize, maximize, or restore a view in the current perspective.

You can:

Task	Action
Open a view	Do one of the following: <ul style="list-style-type: none">• In the Perspective Resources view, click a resource, pull down its menu using the handle to the right of the resource name, and select the view to open.• In the application menu bar, select View > Open and choose a view.
Close a view	Select the view to close. In the application menu bar, select View > Close . You can also click the X in the view's upper right corner.
Maximize a view	Click the box in the view's upper right corner. The view enlarges to fill the entire perspective window. Click the box again to return the view to its former size.
Minimize a view	Click the _ in the view's upper right corner. The view shrinks to a small tab at the bottom of the perspective window.
Minimize all views	In the application menu bar, select View > Minimize All Views .
Restore a view	Click the box on the minimized tab to maximize the view. Click the box again to return the view to its former (smaller) size so you can see other views at the same time.
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 189

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 188

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.

Manage and Monitor

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 190
- *Refreshing or Converting an Instance* on page 192
- *Removing an Instance* on page 193
- *Shared-Disk Mode* on page 194
- *sccinstance Command* on page 195

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 189
- *Refreshing or Converting an Instance* on page 192
- *Removing an Instance* on page 193
- *Shared-Disk Mode* on page 194
- *sccinstance Command* on page 195

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
```


- (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 189
- *Deploying an Instance from a Shared Disk Installation* on page 190
- *Removing an Instance* on page 193
- *Shared-Disk Mode* on page 194
- *sccinstance Command* on page 195

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 189
- *Deploying an Instance from a Shared Disk Installation* on page 190
- *Refreshing or Converting an Instance* on page 192

- *Shared-Disk Mode* on page 194
- *sccinstance Command* on page 195

Shared-Disk Mode

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

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

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

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

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

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

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

See also

- *Enabling and Disabling Shared-Disk Mode* on page 189
- *Deploying an Instance from a Shared Disk Installation* on page 190
- *Refreshing or Converting an Instance* on page 192
- *Removing an Instance* on page 193
- *sccinstance Command* on page 195

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.)

Manage and Monitor

- **-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:

```
$ 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 189
- *Deploying an Instance from a Shared Disk Installation* on page 190
- *Refreshing or Converting an Instance* on page 192
- *Removing an Instance* on page 193
- *Shared-Disk Mode* on page 194
- *Instances* on page 189

Repository

The Sybase Control Center embedded repository stores information related to managed resources, as well as user preference data, operational data, and statistics.

You can back up the repository database on demand, schedule automatic backups, restore the repository from backups, and configure repository purging options. Full and incremental backups are available. A full backup copies the entire repository. An incremental backup copies the transaction log, capturing any changes since the last full or incremental backup.

By default, Sybase Control Center saves backups as follows:

- Each full backup is stored in its own subdirectory in `<SCC-install-directory>/backup`.
- Each incremental backup is stored in a file in `<SCC-install-directory>/backup/incremental`.

Sybase recommends that you periodically move backup files to a secondary storage location to prevent the installation directory from becoming too large.

Scheduling Backups of the Repository

Configure full and incremental backups of the repository to occur automatically.

Prerequisites

Determine your backup strategy, including when to perform full backups and incremental backups. For example, you might schedule incremental backups every day and a full backup every Saturday.

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

Task

A full backup copies the entire repository. An incremental backup copies the transaction log, capturing any changes since the last full or incremental backup.

1. From the main menu, select **Application > Administration**.
2. In the left pane, select **Repository**.
3. Click the **Full Backup** tab.
4. (Optional) To change the directory in which backups will be stored, click **Browse** and navigate to the desired directory.
5. Select **Schedule a Regular Backup**.
6. Specify the day you want scheduled backups to begin. Enter a **Start date** or click the calendar and select a date.
7. (Optional) Use the **Time** and **AM/PM** controls to specify the time at which backups occur.
8. Specify how often backups occur by setting the **Repeat interval** and selecting hours, days, or weeks.
9. (Optional) To purge the repository after each backup, select **Run a repository purge after the backup completes**.
10. If you include purging in the backup schedule, go to the **Size Management** tab and unselect **Automatically purge the repository periodically** to disable automatic purging.
11. Click **Apply** to save the schedule.
12. Click the **Incremental Backup** tab and repeat the steps above to schedule incremental backups to occur between full backups.

Next

Set purging options on the Size Management tab.

See also

- *Modifying the Backup Schedule* on page 201
- *Forcing an Immediate Backup* on page 201
- *Restoring the Repository from Backups* on page 202

- *Configuring Repository Purging* on page 204

Modifying the Backup Schedule

Suspend or resume repository backups or change the backup schedule.

Prerequisites

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

Task

1. From the main menu, select **Application > Administration**.
2. In the left pane, select **Repository**.
3. Choose the type of backup to modify:
 - Click the **Full Backup** tab, or
 - Click the **Incremental Backup** tab.
4. (Optional) To suspend or resume the backup schedule, select or unselect **Schedule a Regular Backup**.
When you unselect (uncheck) this option, the scheduling area is grayed out and scheduled backups no longer occur. However, the schedule is preserved and you can reinstate it at any time.
5. To change the backup schedule, edit the **Start date**, **Time**, **Repeat interval**, or units. You can also select or unselect **Run a repository purge after the backup completes**.
6. Click **Apply** to save the schedule.

See also

- *Scheduling Backups of the Repository* on page 200
- *Forcing an Immediate Backup* on page 201
- *Restoring the Repository from Backups* on page 202
- *Configuring Repository Purging* on page 204

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 200
- *Modifying the Backup Schedule* on page 201
- *Restoring the Repository from Backups* on page 202
- *Configuring Repository Purging* on page 204

Restoring the Repository from Backups

Load backup files into the repository database to revert undesirable changes or to recover from a catastrophic failure.

If you configured Sybase Control Center to store backups somewhere other than the default location, change the source directory in the copy commands in this procedure.

1. Shut down Sybase Control Center.
2. Copy the most recent full backup from <SCC-install-directory>/backup/<generated_directory_name> to <SCC-install-directory>/services/Repository. For example:

Windows:

```
copy C:\sybase\SCC-3_2\backup\repository.  
270110161105\scc_repository.db  
C:\sybase\SCC-3_2\services\Repository
```

UNIX:

```
cp <SCC-install-directory>/backup/repository.270110161105/  
scc_repository.db  
<SCC-install-directory>/services/Repository
```

3. If you have no incremental backups to load,
 - a) Also copy the log file from <SCC-install-directory>/backup/<generated_directory_name> to <SCC-install-directory>/services/Repository. For example:

Windows:

```
copy C:\sybase\SCC-3_2\backup\repository.  
270110161105\scc_repository.log  
C:\sybase\SCC-3_2\services\Repository
```

UNIX:

```
cp <SCC-install-directory>/backup/repository.270110161105/
scc_repository.log
<SCC-install-directory>/services/Repository
```

b) Skip to step 5 on page 203.

4. Start the repository database using the **-ad** option, which directs it to load transaction logs (incremental backups) from the `incremental` directory. (The database loads full backups automatically.) For example:

Windows:

```
cd <SCC-install-directory>\services\Repository
..\..\bin\sa\bin_<platform>\dbsrv11.exe scc_repository -ad
<SCC-install-directory>\backup\incremental
```

UNIX:

```
cd <SCC-install-directory>/services/Repository
../../bin/sa/bin_<platform>/dbsrv11 scc_repository -ad
<SCC-install-directory>/backup/incremental
```

The repository database loads the full backup and any subsequent incremental backups present in the `incremental` directory. Incremental backups are loaded in date order. After loading and saving, the database shuts down.

5. Start Sybase Control Center.

If you loaded incremental backups, SCC starts normally (that is, no further recovery occurs). If you copied a full backup to the `Repository` directory, the database recovers the repository from the full backup.

Example: Loading incremental backups into the repository database

These commands start SQL Anywhere® on a 32-bit Windows machine:

```
% cd C:\sybase\SCC-3_2\services\Repository
% ....\bin\sa\bin_windows32\dbsrv11.exe scc_repository -ad
C:\sybase\SCC-3_2\backup\incremental
```

These commands start SQL Anywhere on a 64-bit machine running Solaris:

```
$ cd /opt/sybase/SCC-3_2/services/Repository
$ ../../bin/sa/bin_sunsparc64/dbsrv11 scc_repository -ad
/opt/sybase/SCC-3_2/backup/incremental
```

See also

- *Scheduling Backups of the Repository* on page 200
- *Modifying the Backup Schedule* on page 201
- *Forcing an Immediate Backup* on page 201
- *Configuring Repository Purging* on page 204

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 200
- *Modifying the Backup Schedule* on page 201
- *Forcing an Immediate Backup* on page 201
- *Restoring the Repository from Backups* on page 202
- *Graphing Performance Counters* on page 169

- *Displaying Resource Availability* on page 167

Logging

Logging helps Sybase Control Center administrators identify and track errors and other system events by recording messages about the events in log files.

Sybase Control Center maintains these logs:

- The client log – captures messages about activities in the browser-based client components. These messages are generated by the component product modules to display information that is pertinent to the user but not critical enough to warrant a pop-up. Sybase also uses the client log to trace client browser operations.
- Server logs – capture messages about activities during the initialization sequence, such as starting services; auditing messages recording logins and logouts; errors such as missed scheduled events; and other events on the server. Server logs include:
 - Component logs, which record only events concerning individual product modules
 - The SCC agent log, which is a composite log. In an SCC server, the agent log records events in all product modules and in the Sybase Control Center framework. In an SCC agent, the agent log records events in the agent.
- The repository log – captures information about inserts and updates that have occurred in the Sybase Control Center repository, a SQL Anywhere database. This log is in `SCC-3_2\log\repository.log`.
- The alert services log – captures information about alert service status and events, including execution of alert-triggered scripts (start time, end time, and status and exit codes). This log is in `SCC-3_2\log>alert-server.log`.

Viewing the Replication Component Log

Launch the Sybase Control Center Log Viewer to display replication activity log.

Prerequisites

Verify that you are in a replication monitoring view.

Task

From any replication monitoring view, click the notebook icon in the left corner of the status bar. The status bar is at the bottom of a monitoring view.

See also

- *Viewing Sybase Control Center Server Logs* on page 206
- *Viewing the Sybase Control Center Client Log* on page 206
- *Changing the Logging Level* on page 207
- *Logging or Message Levels* on page 207
- *Changing Logging Configuration* on page 208

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 Replication Component Log* on page 205
- *Viewing the Sybase Control Center Client Log* on page 206
- *Changing the Logging Level* on page 207
- *Logging or Message Levels* on page 207
- *Changing Logging Configuration* on page 208

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 Replication Component Log* on page 205
- *Viewing Sybase Control Center Server Logs* on page 206
- *Changing the Logging Level* on page 207
- *Logging or Message Levels* on page 207
- *Changing Logging Configuration* on page 208

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 Replication Component Log* on page 205
- *Viewing Sybase Control Center Server Logs* on page 206
- *Viewing the Sybase Control Center Client Log* on page 206
- *Logging or Message Levels* on page 207
- *Changing Logging Configuration* on page 208
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71

Logging or Message Levels

Describes values you can use to control the types of events that are logged by Sybase Control Center.

These are the logging levels, from highest to lowest. The higher the level, the more serious an event must be to be logged. Each level includes all the levels above it—for example, if you set the logging level to WARN, you log events for the WARN, ERROR, and FATAL levels.

OFF	Nothing is logged. This is the highest level.
-----	---

FATAL	Logs only very severe error events that lead the server to abort. This is the highest level at which events are logged.
ERROR	Logs error events that might allow the server to continue running.
WARN	Logs potentially harmful situations. WARN is the default logging level during normal operation (that is, after system initialization).
INFO	Logs informational messages that track the progress of the server in a coarse-grained fashion. INFO is the default logging level during the system initialization process.
DEBUG	Logs a larger set of events that provides a finer-grained picture of how the server is operating. This level is recommended for troubleshooting.
ALL	Logs all loggable events. This is the lowest level.

See also

- *Viewing the Replication Component Log* on page 205
- *Viewing Sybase Control Center Server Logs* on page 206
- *Viewing the Sybase Control Center Client Log* on page 206
- *Changing the Logging Level* on page 207
- *Changing Logging Configuration* on page 208
- *scc Command* on page 78

Changing Logging Configuration

Edit the logging configuration file, `log4j.properties`, to modify Sybase Control Center logging.

You can change the names, locations, or maximum size of the log files as well as the number of log files backed up.

Options for the **scc** command let you change the overall Sybase Control Center log message level when you start SCC, but if you choose the DEBUG level, the large volume of log messages generated may be inconvenient. Editing the log properties file gives you finer control; you can set logging levels for each Sybase Control Center component separately. Sybase recommends making such changes only if you are familiar with log4j and you are working with Sybase technical support; DEBUG-level log messages are not likely to be meaningful to you. (If you have not used log4j before, a good place to start is <http://logging.apache.org/log4j/1.2/manual.html>.)

1. Shut down Sybase Control Center.
2. Make a backup copy of the `log4j.properties` file located in `<SCC-installation-directory>/conf`.
3. Open the `log4j.properties` file for editing.

4. Change values in the file to suit your needs. For example:

To	Modify
Change the name or location of a log file	<ul style="list-style-type: none"> • Agent log – log4j.appender.agent.File • Repository log – log4j.appender.repository.File • Collection statistics log – log4j.appender.collection-stats.File • Alert server log – log4j.appender.alert.File • Gateway log – log4j.appender.gateway.File
Change the maximum size that a log file can reach before Sybase Control Center creates a new file	<ul style="list-style-type: none"> • Agent log – log4j.appender.agent.MaxFileSize • Repository log – log4j.appender.repository.MaxFileSize • Collection statistics log – log4j.appender.collection-stats.MaxFileSize • Alert server log – log4j.appender.alert.MaxFileSize • Gateway log – log4j.appender.gateway.MaxFileSize
Change the number of log files Sybase Control Center backs up before deleting the oldest file	<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 Replication Component Log* on page 205
- *Viewing Sybase Control Center Server Logs* on page 206
- *Viewing the Sybase Control Center Client Log* on page 206
- *Changing the Logging Level* on page 207
- *Logging or Message Levels* on page 207
- *Starting and Stopping Sybase Control Center in Windows* on page 68
- *Starting and Stopping Sybase Control Center in UNIX* on page 71

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 67

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 210
- *shutdown command* on page 211
- *status Command* on page 212

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 210
- *shutdown command* on page 211
- *status Command* on page 212

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 210
- *info Command* on page 210
- *status Command* on page 212

status Command

Display the status of the SCC agent, plug-in, or service components of Sybase Control Center.

Syntax

```
status [-a | --agent]
[-h | --help]
[-p | --plugin [plugin-name]]
[-s | --service [service-name]]
```

Parameters

- **-a | --agent** – display the status of the Sybase Control Center agent component.
- **-h | --help** – display information about the **info** command.
- **-p | --plugin [plugin-name]** – display the status of the specified Sybase Control Center plug-in (for example, ASEMap, the Adaptive Server® management module). Omit the plugin-name argument to return a list of plug-ins.
- **-s | --service [service-name]** – display the status of the specified Sybase Control Center service (for example, the Alert service or the Messaging service). Omit the service-name argument to return a list of services.

Examples

- **Example 1** – displays status information on the Repository service:

```
status --service Repository
```

Permissions

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

See also

- *help Command* on page 210
- *info Command* on page 210
- *shutdown command* on page 211

Manage and Monitor the Replication Environment

Monitor the availability of the servers in the replication environment, control the flow of data and configure replication parameters to improve server performance.

Availability Monitoring

The statistics allow you to monitor the availability of each server and the state of the components, such as connections and routes, in a replication environment.

The status of a server is determined by its availability, whether it is up or down, and the state of the components within the server. The status of the server is a roll-up of the status of all the components within it.

Availability monitoring provides you with the information to determine not only the availability of the servers but also their effectiveness in the replication environment. There are a number of availability monitoring statistics that are gathered for replication. The process of gathering these statistics is turned on by default with the default interval time.

See also

- *Historical Performance Monitoring* on page 168
- *Heat Chart* on page 167
- *Replication Latency* on page 230
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123

Replication Topology

Replication topology is a graphical representation of how the servers in a replication environment are connected to each other. A topology or network diagram illustrates how data moves in a replication environment.

It provides a visual map of the availability of the replication environment. The lines on a topology diagram represent the servers, connections, and routes from the source to the target data server.

See also

- *Replication Latency* on page 230
- *Replication Monitors* on page 216

Replication Tracing

Replication tracing provides the information you need to determine the source, target, and the path of the data flow anywhere in the replication environment.

Replication tracing supports large, complex, and geographically dispersed networks and these component types:

- Sources – these are servers, databases, tables, stored procedures, and all components to which Replication Agents provide access.
- Targets – these are servers, databases, tables, and stored procedures.
- Replication Paths – these are connections, routes, and intermediate Replication Servers.

Displaying Replication Tracing from Target to Source

Select a database server and show the data flow path.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > View Replication Topology**.
2. Right-click a target database server and select **Trace Replication Sources**.
The traced servers, connections, and routes are highlighted in the topology view. The list of servers that participate in the trace are displayed in the Replication Tracing tab.

Note: There can be more than one source, depending on the replication environment.

3. (Optional) To remove the replication tracing, right-click a server and select **Clear Replication Tracing**.

See also

- *Displaying Replication Tracing from Source to Target* on page 214
- *Displaying Replication Tracing for a Replication Server* on page 215

Displaying Replication Tracing from Source to Target

Select a database server and show the data flow path, which includes the final target destination.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > View Replication Topology**.
 2. Right-click a source database server and select **Trace Replication Targets**.
The traced servers, connections, and routes are highlighted in the topology view. The list of servers that participate in the trace are displayed in the Replication Tracing tab.
-
- Note:** There can be more than one target, depending on the replication environment.
-
3. (Optional) To remove the replication tracing, right-click a server and select **Clear Replication Tracing**.

See also

- *Displaying Replication Tracing from Target to Source* on page 214
- *Displaying Replication Tracing for a Replication Server* on page 215

Displaying Replication Tracing for a Replication Server

Select a Replication Server and display all the data flow paths through the replication environment, including other Replication Servers, Replication Agents, and database servers.

Prerequisites

Add all the servers to be monitored to the Perspective Resources view, authenticate these servers, and verify that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > View Replication Topology**.
2. Right-click a Replication Server and select **Trace Replication Server**.
3. (Optional) To change the layout type and orientation of the topology, use the **View Controls** tab on the left. You can also use the additional options that control the way the topology appears.
4. (Optional) To view the list of nodes and edges included in the current topology, select the **Details** tab on the left.

To retain any changes you have made to the default layout settings of the replication topology, close Sybase Control Center without closing the Replication Topology view.

See also

- *Displaying Replication Tracing from Target to Source* on page 214
- *Displaying Replication Tracing from Source to Target* on page 214

Replication Monitors

Replication monitors display the availability and critical performance statistics of a particular component in the replication environment.

The replication monitors provide the status of Replication Servers, Replication Agents, and replication paths. Each monitor also provides a Quick Admin view to allow you to control the flow of data and configure replication parameters, a list of replication paths, and a status bar for each view.

The Replication Server Monitor includes these views:

- Overview – displays high-level information, such as the server name, version, and its status. It also displays the memory usage graph, and a bar chart for partitions and the Top 10 queues by size.
- Paths – displays a table of Replication Server paths, including their state, activity, latency, and latency monitoring status information. It also displays the latency history chart of a Replication Server.
- Connections – displays the information for Replication Server connections, which include the name, Replication Agent (RA) status, and Data Server Interface (DSI) status. It also shows these graphs for the selected connection:
 - DSI Activity and Stable Queue Transaction (SQT) Cache
 - Transactions Removed from SQT
 - Transaction Activity
 - Transaction Success Rate
 - Group and Batch Size
- Logical Connections – displays the Replication Server logical connection information in a warm-standby environment.
- Routes – displays the Replication Server route information, which includes the name, route status, and the route's activity graph. The name of a dedicated route is displayed with its corresponding primary connection enclosed in parenthesis.
- Queues – displays Replication Server queues information, which includes the name and type. The name of a dedicated route queue is displayed with its corresponding primary connection enclosed in parenthesis. It also shows these graphs for the selected queue :
 - Queue Reads and Writes
 - Queue Bytes and Queue Cache
 - SQT cache (only for inbound queues)
 - Backlog
 - Duplicates Detected
 - Transactions Removed from SQT (only for inbound queues)
 - Distributor activity (only for inbound queues)

The Replication Agent Monitor includes this view:

- Overview – displays high-level information, which includes the status, name, version number of the server, primary connection, primary host, Replication Server connection, and the information on when it was last started. It displays a list of paths, which you can click to open a path dashboard. It also displays these graphs:
 - Java Virtual Machine (JVM) Memory Usage
 - LTI Activity and Queue Sizes
 - Log Reader Activity, Queue sizes, and Scan Buffer

The Adaptive Server Replication Agent Thread Monitor has this view:

- Overview – displays high-level information, such as the name, status, and controlling Replication Servers for multiple RepAgent thread connections. It also displays the log and activity details in a graph format.

The Replication Path Dashboard has this view:

- Overview – displays a detailed information of a particular path, which includes its primary and replicate connection, status, data loss detection result, and these graphs:
 - Adaptive Server RepAgent thread commands sent per second
 - Inbound queue activity for the primary database
 - Outbound queue activity for the replicate database
 - Database Server Interface (DSI) Commands per Second
 - Latency
 - Latency History
 - Dedicated routes (if any)

See also

- *Replication Latency* on page 230
- *Replication Topology* on page 213

Data Flow Control

Controlling data flow comprises one part of troubleshooting in your replication environment.

The other part is tuning servers and components. You use either or both of the methods based on the problems in the environment.

See also

- *Replication Environment* on page 2
- *Replication Servers* on page 218
- *Replication Agents* on page 225
- *Replication Paths* on page 227
- *Compatibility Matrix* on page 3
- *Quiescing or Resuming Replication Server* on page 222
- *Suspending or Resuming a Connection's Replication Agent* on page 223

Manage and Monitor

- *Suspending or Resuming a Connection* on page 223
- *Suspending or Resuming a Route* on page 224
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227

Replication Servers

Replication Server maintains replicated data in multiple databases while ensuring the integrity and consistency of the data.

Replication Server provides clients using databases in the replication system with local data access, thereby reducing load on the network and centralized computer systems.

See also

- *Replication Environment* on page 2
- *Data Flow Control* on page 217
- *Replication Agents* on page 225
- *Replication Paths* on page 227
- *Compatibility Matrix* on page 3

Displaying Replication Server Status

View the most recent Replication Server statistics. Sybase Control Center automatically updates the page when new status and statistical values are collected.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have permission to perform this task.

Task

Note: You can also view the server status through the status badges on the server icons in the Perspective Resources, Perspective Heat Chart, Resource Explorer, and Replication Topology views.

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor** .
The Replication Server monitor view appears, with the Overview as the default display.
2. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

See also

- *Displaying Logical Connection Status* on page 219
- *Hiding Connection Thread Status* on page 220
- *Displaying Queue Status* on page 221
- *Displaying Route Status* on page 222

- *Logical Connection View Details* on page 219

Displaying Logical Connection Status

View the information of all logical connections of a Replication Server in a warm standby replication environment.

Prerequisites

Verify that the replication environment has been configured, and all servers have been registered and authenticated, and that the monitoring collection has been running.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Logical Connections**.
A list of all logical connection names along with their state information, active and standby connection names, their respective state information, and the inbound queue information for each logical connection appears in a table.
3. (Optional) Click the active or standby connection in the list to go to the Connections view, and see the status information of the physical connections that are associated with the logical connections.
4. (Optional) Click the inbound queue under its corresponding column to go to the Queues view.

See also

- *Displaying Replication Server Status* on page 218
- *Hiding Connection Thread Status* on page 220
- *Displaying Queue Status* on page 221
- *Displaying Route Status* on page 222
- *Logical Connection View Details* on page 219

Logical Connection View Details

The columns on the **Logical Connection** view show different information about the logical connections of a Replication Server.

Column	Description
Name	The DBID (database ID) for the current logical connection.

Column	Description
State	The state of the current logical connection. The possible values are: <ul style="list-style-type: none"> • Active • Switching Active Connection • Creating Standby Connection
Active Connection	The DBID, the data server, and the database name for the current active database.
Active State	The state of the active connection's RepAgent thread. The possible values are: <ul style="list-style-type: none"> • Active • Suspended
Standby Connection	The DBID, the data server, and the database name for the current standby database.
Standby State	Shows the Data Server Interface (DSI) state of the standby connection. The possible values are: <ul style="list-style-type: none"> • Active • Suspended
Inbound Queue	The name of the inbound queue for a logical connection.

See also

- *Displaying Replication Server Status* on page 218
- *Displaying Logical Connection Status* on page 219
- *Hiding Connection Thread Status* on page 220
- *Displaying Queue Status* on page 221
- *Displaying Route Status* on page 222

Hiding Connection Thread Status

Hide the connection status of a Replication Agent or Data Server Interface (DSI) thread.

Prerequisites

Verify that the replication environment has been configured and that all servers have been registered, authenticated, and replicating data correctly.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Connections**.
3. Select a Replication Server connection.
4. Click **Quick Admin**.
5. In the Hide Connection Status section of the **Control** tab, click either **Hide RepAgent Status** or **Hide DSI Status** to override the connection's current configuration.
6. Click **Apply**.

See also

- *Displaying Replication Server Status* on page 218
- *Displaying Logical Connection Status* on page 219
- *Displaying Queue Status* on page 221
- *Displaying Route Status* on page 222
- *Logical Connection View Details* on page 219

Displaying Queue Status

View the status information of Replication Server queues in a replication environment.

Prerequisites

Verify that the replication environment has been configured, all servers have been registered and authenticated, and that monitoring collection has been running.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Queues**.
3. Select a queue from the table to view its corresponding activity graphs.
4. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

See also

- *Displaying Replication Server Status* on page 218
- *Displaying Logical Connection Status* on page 219
- *Hiding Connection Thread Status* on page 220
- *Displaying Route Status* on page 222
- *Logical Connection View Details* on page 219

Displaying Route Status

View the status information of Replication Server routes in a replication environment.

Prerequisites

Verify that the replication environment has been configured, all servers have been registered and authenticated, and that monitoring collection has been running.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Routes**.
You can see a route's status in the Replication Server Routes table.
3. Select a route from the table to view its corresponding activity graph.
4. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

See also

- *Displaying Replication Server Status* on page 218
- *Displaying Logical Connection Status* on page 219
- *Hiding Connection Thread Status* on page 220
- *Displaying Queue Status* on page 221
- *Logical Connection View Details* on page 219

Quiescing or Resuming Replication Server

Shut down data flow from all sources and quiesce the current Replication Server or resume data flow from the up-stream replication path. You may need to quiesce all Replication Servers in your environment to recover databases, alter routes, and troubleshoot the system.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, verify that the default collection is running, and that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Quick Admin**.
3. In the Control tab, click **Quiesce** or **Resume All**.

See also

- *Suspending or Resuming a Connection's Replication Agent* on page 223
- *Suspending or Resuming a Connection* on page 223
- *Suspending or Resuming a Route* on page 224
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227
- *Data Flow Control* on page 217

Suspending or Resuming a Connection's Replication Agent

Suspend or resume a connection's Replication Agent that is associated with a Replication Server.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, verify that the default collection is running, and that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Connections**.
3. Select a connection.
4. Click **Quick Admin**.
5. In the Replication Agent section of the Control tab, click **Suspend** or **Resume**.

See also

- *Quiescing or Resuming Replication Server* on page 222
- *Suspending or Resuming a Connection* on page 223
- *Suspending or Resuming a Route* on page 224
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227
- *Data Flow Control* on page 217

Suspending or Resuming a Connection

Suspend or resume a Replication Server connection.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, verify that the default collection is running, and that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Connections**.
3. Select a connection.
4. Click **Quick Admin**.
5. In the Data Server Interface section of the Control tab, click **Suspend** or **Resume**.
6. (Optional) Click **Suspended Immediately**.
7. (Optional) Select one of these options when resuming a connection:
 - Skip the first transaction. This option is available only for versions of Replication Server 15.0 or earlier.
 - Resume after skipping transactions. Provide the number of transactions to skip. This option is available only for versions of Replication Server 15.1 and later.
 - Execute sytem transaction.

See also

- *Quiescing or Resuming Replication Server* on page 222
- *Suspending or Resuming a Connection's Replication Agent* on page 223
- *Suspending or Resuming a Route* on page 224
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227
- *Data Flow Control* on page 217

Suspending or Resuming a Route

Suspend or resume a route associated to a Replication Server.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, verify that the default collection is running, and that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Routes**.
3. Select a route.
4. Click **Quick Admin**.
5. In the Control tab, click **Suspend** or **Resume**.
6. (Optional) Click **Skip transaction with large message** when resuming a route.

Note: This option applies only to direct routes where the site version at the replicate site is 12.1 or earlier.

See also

- *Quiescing or Resuming Replication Server* on page 222
- *Suspending or Resuming a Connection's Replication Agent* on page 223
- *Suspending or Resuming a Connection* on page 223
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227
- *Data Flow Control* on page 217

Replication Agents

Replication Agents enable the movement and sharing of data where and when it is needed - into data warehouses, data marts, OLTP applications and mobile workforce.

Replication Agents allows heterogeneous data servers to act as primary data servers in a replication system based on Sybase replication technology.

See also

- *Replication Environment* on page 2
- *Data Flow Control* on page 217
- *Replication Servers* on page 218
- *Replication Paths* on page 227
- *Compatibility Matrix* on page 3

Displaying Replication Agent Status

Review high-level server availability by viewing the Replication Agent or Mirror Replication Agent status.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Agent and click **Resource > Monitor** .

The monitor view appears displaying high-level information, which includes the status of the selected Replication Agent.

2. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

See also

- *Displaying RepAgent Thread Status in Adaptive Server* on page 226
- *Suspending or Resuming a Replication Agent or Mirror Replication Agent* on page 226
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227

Displaying RepAgent Thread Status in Adaptive Server

View the RepAgent thread status in an Adaptive Server, either single instance or shared disk cluster, for replication.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have permission to perform this task.

Task

1. From the Perspective Resources view, select an Adaptive Server and click **Resource > Monitor** .
The Replication Agent Thread Monitor appears. You can view the status of an Adaptive Server RepAgent thread in the displayed table with its name and the controlling Replication Server information.
2. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

See also

- *Displaying Replication Agent Status* on page 225
- *Suspending or Resuming a Replication Agent or Mirror Replication Agent* on page 226
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227

Suspending or Resuming a Replication Agent or Mirror Replication Agent

Suspend or resume a Replication Agent or a Mirror Replication Agent in a replication environment.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, verify that the default collection is running, and that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Agent or a Mirror Replication Agent and click **Resource > Monitor**.

2. Click **Quick Admin**.
3. In the Control tab, click **Suspend** or **Resume**.

See also

- *Displaying Replication Agent Status* on page 225
- *Displaying RepAgent Thread Status in Adaptive Server* on page 226
- *Starting or Stopping an Adaptive Server RepAgent Thread* on page 227

Starting or Stopping an Adaptive Server RepAgent Thread

Start or stop a RepAgent thread of an Adaptive Server.

1. From the Perspective Resources view, select an Adaptive Server and click **Resource > Monitor**.
2. Select a RepAgent thread (ASE Server, Replication Only) or click **Replication Agent** in the left pane, then select a RepAgent thread (ASE Server).
3. Click **Quick Admin**.
4. In the Control tab, click **Start** or **Stop**.
5. (Optional) Click **Stop immediately** before clicking **Stop**.

See also

- *Quiescing or Resuming Replication Server* on page 222
- *Suspending or Resuming a Connection's Replication Agent* on page 223
- *Suspending or Resuming a Connection* on page 223
- *Suspending or Resuming a Route* on page 224
- *Data Flow Control* on page 217
- *Displaying Replication Agent Status* on page 225
- *Displaying RepAgent Thread Status in Adaptive Server* on page 226
- *Suspending or Resuming a Replication Agent or Mirror Replication Agent* on page 226

Replication Paths

A replication path is the set of all servers and internal components where transactions travel through when moving from the primary to the replicate database. Replication may occur over a single path from primary to replicate, or over multiple paths.

Single Path Replication

The traditional replication in a Replication Server is over a single path. It is a single connection from primary database to primary Replication Server, a single connection from primary Replication Server to replicate Replication Server, and a single connection from replicate Replication Server to the replicate database. The single replication path always maintains transaction commit order.

MultiPath Replication

To increase replication throughput and performance, and reduce contention, Replication Server 15.7 and later also use multiple replication paths. Multi-Path Replication™ supports the replication of data from a single primary database through different paths or streams, while still maintaining transaction commit order within a path, but not across paths.

A replication path encompasses all the components and modules between the Replication Server and the primary and replicate databases. In multipath replication, there are multiple primary replication paths for multiple Replication Agent connections from a single primary database to one or more Replication Servers, and multiple replicate paths from one or more Replication Servers to a single replicate database.

Multiple RepAgent Thread Connections from the Primary Database

Multiple primary replication paths can exist from the primary database to one or more Replication Servers. Each primary path consists of a RepAgent path and an associated Replication Server primary connection. You can view multiple RepAgent threads in the Adaptive Server Monitor's Replication Agent view, where it shows all the controlling Replication Servers for a RepAgent thread.

Multiple Dedicated Routes

Multiple dedicated routes can exist between Replication Servers to increase the throughput of the Distributor (DIST) thread. A dedicated route distributes only transactions for a specific primary connection.

You can view the dedicated routes in the Replication Topology view and in the Replication Server Monitor Routes view.

Multiple Connections to the Replicate Database

Multiple replicate connections can exist to the same replicate database. You can view these connections in the Replication Topology view and in these Replication Server Monitor views:

- Overview
- Quick Admin
- Paths and Replication Path Dashboard
- Connections

See *Replication Server Administration Guide Volume 2 > Performance Tuning > Multi-Path Replication*.

See also

- *Replication Environment* on page 2
- *Data Flow Control* on page 217
- *Replication Servers* on page 218
- *Replication Agents* on page 225

- *Compatibility Matrix* on page 3

Displaying Replication Path Status

View the status of the servers, connections, routes, and queues that make up a replication path.

Prerequisites

Register and add all the servers to be monitored to the Perspective Resources view, authenticate them, and verify that you have permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server or a Replication Agent and click **Resource > Monitor**.
2. Click a replication path.
The Replication Path Dashboard for the selected replication path appears.
3. (Optional) To maximize or minimize a window within the monitor view, click the maximize or restore window icon in its right corner.

If new statistics are collected while the replication path dashboard is displayed, Sybase Control Center automatically updates the information on the current page.

See also

- *Displaying Latency* on page 230
- *Latency Monitoring Process* on page 119

Displaying Server License Information

View server licensing information from the Replication resource monitors.

Prerequisites

Verify that you have sa privilege.

Note: Licensing information is available only for Replication Server version 15.7 and later.

Task

1. From the Perspective Resources view, select a server and click **Resource > Monitor**.
2. On the left side of the screen, click **License Information**.

Replication Latency

Replication latency is the amount of time it takes for a transaction that occurs in the primary database to be applied to the replicate database.

The time includes Replication Agent processing, Replication Server processing, and network usage. Replication latency appears on the path dashboard.

See also

- *Historical Performance Monitoring* on page 168
- *Heat Chart* on page 167
- *Availability Monitoring* on page 213
- *Key Performance Indicators and Key Performance Areas for Replication* on page 123
- *Setting Up the Latency Monitoring Process* on page 117
- *Replication Topology* on page 213
- *Replication Monitors* on page 216

Displaying Latency

View latency information to determine the amount of time it takes for a transaction to be applied from the primary database to the replicate database.

Prerequisites

- Add all the resources to be monitored in the Perspective Resources view.
- Verify that you have permission to perform this task and that the latency monitoring process has already been started.

Task

1. From the Perspective Resources view, select a server and click **Resource > Monitor**.
2. Click **Paths**.
3. Select or click a path from the path list.
When you select a path, its latency information appears below the path list. When you click a selected path, its dashboard opens, displaying the latency history graph.
4. (Optional) Click the bar graph icon in the top right corner of the Latency window to toggle between the overall latency and the latency monitoring history of the path.

See also

- *Setting Up the Latency Monitoring Process* on page 117
- *Stopping the Latency Monitoring Process* on page 231
- *Latency Chart Details* on page 231
- *Displaying Replication Path Status* on page 229

- *Latency Monitoring Process* on page 119

Stopping the Latency Monitoring Process

Stop the latency monitoring process.

Prerequisites

Verify that all servers to be monitored have been registered, authenticated, and added to the Perspective Resources, that latency monitoring process is running, and that you have the permission to perform this task.

Task

1. From the Perspective Resources view, select a Replication Server and click **Resource > Monitor**.
2. Click **Paths**.
3. Select a path where the "Monitoring" value is true.
4. Click **Quick Admin** to display the latency monitoring process administration view for the selected path.
5. Click **Stop**.

See also

- *Setting Up the Latency Monitoring Process* on page 117
- *Displaying Latency* on page 230
- *Latency Chart Details* on page 231
- *Latency Monitoring Process* on page 119

Latency Chart Details

Replication latency is the amount of time it takes for a transaction to be applied to the replicate database from the primary database. The latency threads table describes the threads that represent the latency chart in the path dashboard.

Table 25. Latency Threads

Latency Thread	Description
Replication Executor Thread (EXEC)	The number of seconds from the start, to time the ticket passes through the EXEC thread.

Latency Thread	Description
Distribution (DIST) Thread	The number of seconds from the primary Replication Server time to the time the ticket passes through the DIST thread. For each primary database managed by a Replication Server, there is a DIST, which in turn uses Stable Queue Transaction (SQT) to read from the inbound queue, and Stable Queue Manager (SQM) threads to write transactions to the outbound queue.
Replication Server Interface (RSI) Thread	The number of seconds from the DIST thread time to the time the ticket passes through the RSI thread. RSI threads send messages from one Replication Server to another. There is one RSI thread for each destination Replication Server.
Data Server Interface (DSI) Thread	The number of seconds from the RSI thread time to the time the ticket passes through the DSI thread. The DSI thread translates the transaction modifications into remote procedure calls (RPC) or the language as specified by the function strings in the function-string class assigned to the destination database. Replication Server starts the DSI threads to submit transactions to a replicate database to which it has a connection.
Replicate Database	The number of seconds from the DSI thread time to the time the ticket arrived at the replicate database.
Total Latency	The number of seconds between the time at origin and destination for the transaction committed.

See also

- *Setting Up the Latency Monitoring Process* on page 117
- *Stopping the Latency Monitoring Process* on page 231
- *Displaying Latency* on page 230

Troubleshoot Sybase Control Center for Replication

Troubleshoot problems that occur in Sybase Control Center for Replication.

Collection Job for Adaptive Server Fails

Problem: A collection job for Adaptive Server may fail when the **number of open databases** is too low.

Solution: Modify the value of **number of open databases** by using either the Server Configuration screen of the Adaptive Server monitor, or these steps:

1. Log in to the Adaptive Server.
`isql -S<server_name> -U<sa user name> -P<sa password>`
2. Run this command to display the current configuration value:
`sp_configure 'number of open databases'`
3. Run this command to change the current configuration value:
`sp_configure 'number of open databases', <number>`
Add 10 to the current configuration value and substitute this number for *<number>*.

Data on Dashboards or Charts Is Missing

Problem: There is no data displayed on the dashboards or on the charts.

Solution: Create collection jobs for Replication Server statistics collections.

See also

- *Setting Up the Latency Monitoring Process* on page 117

No Data Is Returned for Performance Counters Collection Job

Problem: No data is returned for the Performance Counters collection job after completing its execution.

Solution: Let the job execution complete for three cycles. The Replication Server Performance Counters collection job returns data in the Statistics Chart only after it completes execution of approximately three cycles.

Same Data Is Displayed for Two or More Adaptive Server Collection Job Executions

Problem: If the interval time of the Adaptive Server RepAgent thread collection job is set to less than 60 seconds, Sybase Control Center still works but you may get the same data for two or more continuous executions of the job.

Solution: Do not set the interval time for the Adaptive Server RepAgent thread collection job to less than 60 seconds. This is because the collection data are updated only once during 60 seconds for better performance.

Servers Are Missing in the Topology View

Problem: There are servers missing when you display the topology view.

Solution: Verify that all the servers have been added to the current perspective because the topology view can only display servers from current perspective.

Servers, Routes, or Connections Display an Unknown Status

Problem: "Unknown" status is displayed for servers, routes, or connections.

Solution: Verify that a collection job is scheduled and running to display the correct status in the Replication Topology view.

Some Connector Lines Do Not Display a Status Icon

Problem: Some connector lines do not display a status icon.

Solution: Not all connector lines are associated with a managed object. Those associated with a managed object are the only ones that display a status icon. Currently, status icons are displayed only on Replication connections and routes.

Topology Display Is Incorrect When Two Replication Servers With the Same Name Are Registered

Problem: When you register two Replication Servers with the same name in a given replication environment, the replication topology view is displayed incorrectly.

Solution: Do not use the same name for two replication resources. The servers in a replication environment must have unique names.

Topology View Displays Blank Canvas

Problem: The Replication Topology view displays blank canvas.

Solution: From the Replication Topology view, click Reset to restore the default layout. This problem is usually due to a layout that is actually displayed off-screen.

Topology View Takes a Long Time to Display

Problem: The Replication Topology view takes a long time to display.

Solution: The Replication Topology view requires data from each server to be displayed. For large environments, this may take 10 to 15 seconds.

Unable to Retrieve Topology Relationship Data or Server Objects

Problem: An error message "Unable to retrieve topology relationship data" or "Unable to retrieve topology server objects" is displayed.

Solution: These errors are displayed for a number of reasons. Perform these:

- Verify that you have the repMonitorRole for each server in the topology assigned to your login.
- Verify that each server in the topology has been authenticated.
- The server relationship data cannot be retrieved from the Replication Management plug-in or one or more servers cannot be retrieved from the Sybase Control Center.

Check the plug-in and the Unified Agent Framework (UAF) agent log files for errors:

For Windows

- %SYBASE%\SCC-3_1\plugins\RMAP\log\rmap.log
- %SYBASE%\SCC-3_1\log\agent.log

For UNIX

- `$(SYBASE)/SCC-3_1/plugins/RMAP/log/rmap.log`
- `$(SYBASE)/SCC-3_1/log/agent.log`

Problems with Basic Sybase Control Center Functionality

Troubleshoot problems that involve basic features like starting and stopping, authentication, alerts, and scheduling.

Cannot Log In

Problem: Cannot log in to Sybase Control Center Web console.

Solution: Make sure that Sybase Control Center has been configured:

- To allow logins through the operating system
- To grant appropriate roles to your login account

Ask the Sybase Control Center administrator to help you check.

See also

- *User Authorization* on page 103
- *Setting Up Security* on page 83

Sybase Control Center Fails to Start

Problem: The Sybase Control Center server does not start.

Solution 1: Port conflict

Solution: SCC might be using one or more ports that are also being used by another server or application on this machine. To check for port conflicts:

1. Execute this command:

```
scc --info ports
```

The command lists all the ports on which Sybase Control Center and its services listen, indicates whether each port is in use, and shows the service running on each port. If SCC is not running, any port shown to be in use represents a conflict.

2. If you discover a conflict, use **scc --port** to change the port used by the Sybase Control Center service.

Solution 2: Insufficient memory

You might see this error why you try to start: Could not create the Java Virtual machine. Increase the maximum memory setting.

See also

- *Configuring Ports* on page 98

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 119

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 204
- *Graphing Performance Counters* on page 169

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 174

Memory Warnings at Startup

Problem: When Sybase Control Center starts, you see warnings about system memory or heap memory allocation.

Solution: Increase the maximum memory setting (`SCC_MEM_MAX` or `jvmopt=-Xmx`).

See also

- *Configuring Memory Usage* on page 75

SCC Out of Memory Errors

Problem: Sybase Control Center generates OutOfMemory errors.

Solution:

- If the OutOfMemory error says that Sybase Control Center is out of heap space, increase the maximum memory setting (*SCC_MEM_MAX* or `jvmopt=-Xmx`).
- If the OutOfMemory error says that Sybase Control Center is out of permanent generation space, increase the permanent memory setting (*SCC_MEM_PERM* or `jvmopt=-XX:MaxPermSize`).
- Repeated OutOfMemory errors may indicate a memory leak. OutOfMemory errors generate heap dumps:
 - When Sybase Control Center runs as a service in Windows:
C:/windows/system32
 - When Sybase Control Center runs as a service in UNIX:
<SCC-install-directory>/SCC-3_2/binSend the heap dump files to Sybase technical support for analysis.

See also

- *Configuring Memory Usage* on page 75

Glossary: Sybase Control Center for Replication

Definitions of Sybase Control Center terms related to Replication.

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.

component – A server object in a replication environment. Examples of components are connections, routes, and Replication Agent threads.

connection – A connection from a Replication Server to a database.

database – A set of related data tables and other objects that are organized and presented to serve a specific purpose.

data server – A server that provides the functionality necessary to maintain the physical representation of a replicated table in a database. An example of a data server is Sybase Adaptive Server.

Data Server Interface (DSI) – Replication Server threads corresponding to a connection between a Replication Server and replicate database. DSI threads submit transactions from the DSI outbound queue to a replicate database. They consist of a scheduler thread and one or more executor threads. The scheduler thread groups the transactions by commit order and dispatches them to the executor threads. The executor threads map functions to function strings and execute the transactions in the replicate database. DSI threads use an Open Client connection to a database. See also connection.

edge – A topology object representing a relationship between two nodes, displayed in the form of a line that connects the topology node icons. Edges are usually defined with a direction.

Enterprise Connect Data Access (ECDA) – An Open-Server-based software gateway that supports DB-Library and CT-Library application program interfaces and Open Database Connectivity (ODBC) protocols. It serves as a fundamental building block for database middleware applications that provide access to non-Sybase databases.

environment – A set of servers that interact to perform replication. It can be a subset of the replication domain identified by the ID server.

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.

ID Server – A Replication Server in a replication system. It assigns unique ID numbers to every Replication Server and database in the system and maintains the version information for the system.

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.

latency – The measure of the time it takes to distribute a data modification operation from a primary database to a replicate database. The time includes Replication Agent processing, possibly ECDA processing—if the replicate database is non-Sybase—Replication Server processing, and network overhead.

Mirror Activator – Provides an integrated, zero-loss disaster recovery solution when used in conjunction with a storage replication system. Replication Server, Mirror Replication Agent, and Enterprise Connect Data Access (ECDA) are the components of Mirror Activator.

Mirror Replication Agent – A component of a Mirror Activator disaster recovery system that reads the mirror transaction log of the primary database and sends transaction data as LTL commands for replicated tables and replicated stored procedures to the primary Replication Server, which converts the LTL commands into SQL and applies the SQL to the replicate database.

node – A topology object representing a server or other entity type, displayed visually in the form of an icon.

partition – A raw disk partition or operating system file that a Replication Server uses for stable queue storage. See also stable queue.

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.

RepAgent thread – The Replication Agent for a primary Adaptive Server. It reads the transaction log and sends the transaction data as LTL commands for replicated tables and replicated stored procedures to the primary Replication server, which converts the LTL commands into SQL and applies the SQL to the replicate database.

replication – A process by which the changes to the data in one database—including creation, updating, and deletion of records—are also applied to the corresponding records in another database.

Replication Agent – Reads the transaction log of a primary, non-Sybase database, and sends transaction data as LTL commands for replicated tables and replicated stored procedures to the primary Replication Server, which converts the LTL commands into SQL and applies the SQL to the replicate database.

replication path – The set of all servers and internal components, such as threads and queues, that transactions travel through when moving from the primary to the replicate database.

repository – a database in Sybase Control Center that stores information related to managed resources, along with user preference data, operational data, and performance statistics.

resource – a server, agent, or other entity that can be monitored or administered by Sybase Control Center. Resources SCC can manage include Adaptive Server, Replication Server, Replication Agent, Mirror Replication Agent, Sybase IQ, and certain subcomponents.

route – A one-way message stream from a source Replication Server to a destination Replication Server. Routes carry data modification commands (including those for RSSDs) and replicated functions or stored procedures between Replication Servers.

SCC-enabled login account – a user account that has been granted privileges in Sybase Control Center by mapping appropriate Sybase Control Center roles. (Roles are typically mapped to a group to which the account belongs rather than to the account itself.) The user account and group can be native to Sybase Control Center or created in the operating system or the LDAP directory service to which Sybase Control Center authentication is delegated. You must use an SCC-enabled account to log in to Sybase Control Center.

SCC agent – a Sybase Control Center agent that runs on a managed server and enables Sybase Control Center to manage it. The SCC agent is installed automatically as part of the Sybase server.

schedule – the definition of a task (such as the collection of a set of statistics) and the time interval at which Sybase Control Center executes the task.

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 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.

stable queue manager (SQM) – A thread within Replication Server that manages the stable queues. Replication Server uses one SQM thread for each stable queue, whether inbound or outbound.

stable queue – A store-and-forward queue where Replication Server stores messages destined for a route or database connection. Replication Server builds stable queues using its disk partitions.

topology – A graphical representation of how the servers in a replication environment are connected to each other. It is a network diagram that provides a visual map of the availability of the replication environment.

view – a window in a perspective that displays information about one or more managed resources. Some views also let you interact with managed resources or with Sybase Control Center itself. For example, the Perspective Resources view lists all the resources managed by the current perspective. Other views allow you to configure alerts, view the topology of a replication environment, and graph performance statistics.

Index

- Xmx maximum memory option 35, 77
- XX:MaxPermSize permanent memory option 35, 77

A

- accessibility 12
- Adaptive Server collection job
 - same data collected for Replication, troubleshooting 234
- Adaptive Server shared disk cluster
 - RepAgent thread status 226
- Administration Console
 - column filtering 8
 - display tools and options 6
- Adobe Flex 12
- alert notifications 241
- alert storm 241
- alert storm suppression 241
- alert subscriptions
 - disable 181
 - enable 181
- alert type 241
- alert-triggered scripts
 - example 150
- alerts 241
 - about 175
 - configured, deleting 179
 - configured, modifying 177
 - configured, viewing 177
 - configuring duplicate alerts 149
 - configuring e-mail server 54, 101
 - configuring escalations 150
 - configuring storm suppression 149
 - configuring subscriptions 150
 - configuring to execute scripts 149
 - creating 148
 - displaying history 182
 - displaying resolutions 182
 - effects of repository purging on history 204
 - escalations 179
 - log 162
 - modifying subscriptions 180
 - monitoring 177
 - not being generated 237

- notifications, about 182
- notifications, viewing 177
- Replication 152
- resolving 183
- scripts executed by 162
- setting triggering states and thresholds 148
- subscribing to 180
- subscriptions 179
- substitution parameters for scripts 163
- testing 178
- triggering scripts, about 162
- types, states, and severities 176
- unsubscribing from 181
- alerts, Replication
 - types, states, and severities 157
- ALL logging level 207
- authenticate 241
- authenticating
 - Adaptive Server 117
 - Replication Server 117
 - SCC with a managed resource 17, 116
- authentication
 - about 38, 84
 - configuring for LDAP 41, 88
 - configuring for UNIX 40, 86
 - configuring for Windows 39, 85
- authorization 55, 103
- availability 241
- availability, Replication
 - monitoring 213

B

- background, running SCC or SCC agent in 30, 71
- backups
 - about 199
 - changing the schedule 201
 - forcing 201
 - restoring from 202
 - scheduling 200
 - suspending and resuming 201
- badges, status 5

Index

C

- charts, Replication
 - missing data, troubleshooting 233
- client log, viewing 206
- collection job
 - for Adaptive Server in Replication, troubleshooting 233, 234
- collection jobs
 - Replication Performance Counters, troubleshooting 233
- collections 241
- column filtering in SCC 8
- columns
 - sorting by 6
- compatibility, Replication
 - server version 3
- component log
 - replication 205
- components, Replication 241
- configuration
 - optional 164
- configuration parameters, Replication
 - configuring 144
 - setting up 144
- configuring
 - logical connection parameters 146
 - RepAgent thread parameters 147
 - replication parameters 144
 - Replication Server parameter 147
- connections, Replication 241
 - configuring parameters 145
 - resuming 223
 - suspending 223
 - unknown status, troubleshooting 234
- connector lines
 - do not display a status icon for Replication, troubleshooting 234
- console
 - about 209
 - commands 210
- conventions, style and syntax 10
- csi_config.xml file 41, 88

D

- dashboards, Replication
 - missing data, troubleshooting 233
- data collection jobs
 - adding 119

- adding schedules 172
- creating 119
- deleting 171
- displaying history 175
- executing 170
 - not saving data 119
- removing schedules 173
- resuming 171
- stopping 170
- suspending 171
 - viewing schedules 173
- data collection schedules
 - adding 119
 - modifying 174
- data collections
 - replication 122
 - troubleshooting timeouts 238
- data server interface, Replication 242
- data server, Replication 241
- databases, Replication 241
- DEBUG logging level 207
- defined 241–244
- display options in Sybase Control Center 6
- displaying
 - queue status 221
 - replication log 205
 - replication topology view, troubleshooting 235
 - route status 222
- displaying Replication Server status 18, 218
- displaying, Replication
 - latency 230
 - Replication Agent status 19, 225
 - topology 215
- drivers
 - ODBC, registering 14, 68
- DSI
 - hiding connection status 220

E

- e-mail server, configuring for alerts 54, 101
- edges, Replication 242
- Enterprise Connect Data Access 242
- environment variables
 - SCC_MEM_MAX 34–36, 75–77
 - SCC_MEM_PERM 34–36, 75–77
- environment, Replication 242
- ERROR logging level 207

errors

- OutOfMemory 239
- REJECT_DUPLICATE_RESOURCE_AND_COLLECTION policy 238
- timeouts for data collections 238

evaluation

- quick start instructions 13

events 242

- expiration dates for login accounts 108

F

- F11 (browser full screen mode toggle) 8

F5 (browser refresh)

- logging out of Sybase Control Center 237

FATAL logging level 207

Flash Player 15

- foreground, running SCC or SCC agent in 30, 71

full backups 200

full screen mode 8

G

- getting started after installing 15

glossaries

- SCC for Replication terms 241

graphing statistics 169

grid format, using 6

groups 57, 109

- adding login accounts 57, 106
- assigning monitoring and administration roles 56, 103
- creating 56, 105
- in LDAP, mapping to SCC roles 49, 96
- in OS, mapping to SCC roles 49, 96
- remove login 106
- removing 105
- removing roles 104
- SCC Administrator 49, 96
- sybase 49, 96

groups, Replication

- assigning monitoring and administration roles 115

H

heartbeat process

- monitoring 119

heat chart 167, 242

- customizing columns 17, 167

display tools and options 6

displaying 17, 167

filtering resources displayed 17, 167

icons 5

launch icon 4

help command (console) 210

hiding

- RepAgent connection status 220

historical performance monitoring 168

history displays for alerts 182

I

icons

- for server status 5
- in SCC toolbar 4
- minimize/maximize sections of a view 8

ID Server, Replication 242

incremental backups 200

info command (console) 210

INFO logging level 207

instances 242

- about 22, 62, 194
- converting 192
- deploying 20, 60, 190
- deploying and managing 23, 63, 195
- file locations 21, 61, 191
- refreshing 192
- removing 193

interfaces files, importing resources from 113

J

Java system properties

- displaying information about 210

jobs 242

- modifying collection intervals 174
- resuming 174
- suspending 174

jvmopt memory options for Windows services 34, 35, 75, 77

K

key performance areas

- Replication 123

key performance indicators 242

- getting values from heat chart icons 17, 167
- Replication 123, 152

Index

keyboard shortcuts 9

keyboard shortcuts for Adobe Flex 12

KPIs 242

L

latency

chart details 231

displaying 230

monitoring 119

replication 230

stop monitoring 231

viewing 230

latency monitoring

setting up 117

stopping 231

latency, Replication 242

LDAP

configuration properties 43, 90

configuring authentication 41, 88

configuring to authenticate SCC logins 38, 84

log4j.properties file 208

logging in to Sybase Control Center 36, 82

troubleshooting 236

logging in to Sybase Control Center - first user 15

logging levels 207

logging out of Sybase Control Center 83

unintentionally, using F5 browser refresh 237

logical connection

configuring parameters 146

details 219

displaying 219

viewing 219

login accounts

assigning monitoring and administration roles
56, 103

authenticating 17, 116, 117

creating automatically (UNIX) 40, 86

creating automatically (Windows) 39, 85

default 107

expiration date, imposing 108

granting privileges with roles and groups 49,
96

modifying 108

native SCC, adding 107

predefined 57, 109

removing 108

removing roles 104

suspending 108

login accounts, default

about 15

login accounts, Replication

assigning monitoring and administration roles
115

login modules 38, 84

login session timeout 83

setting 55, 102

logs

agent log, viewing 206

alert services 162

alert services log, about 205

changing the logging level 207

client log, about 205

client log, viewing 206

component logs, about 205

configuring 208

repository log, about 205

repository log, viewing 206

SCC agent log, about 205

script execution log, about 205

server logs, about 205

server logs, viewing 206

M

managed resources 183, 243

managed servers

See managed resources

managing

replication environment 213

memory

configuring 34, 75

displaying information about 210

warnings at startup 238

memory leak 239

memory, insufficient 236

message levels 207

minimize/maximize icon 8

Mirror Activator 243

Mirror Replication Agent 243

monitoring

heartbeat process 119

latency 119

performance 168

replication environment 213

monitoring, Replication 1

server availability 213

server status 213

N

- new features
 - Sybase Control Center for Replication 1
- nodes, Replication 243

O

- ODBC drivers
 - registering 14, 68
- online help
 - resetting 237
- operating system
 - configuring to authenticate SCC logins 38, 84
- OutOfMemory errors, SCC 239

P

- parameters
 - configuring logical connection 146
 - routes, configuring 146
- parameters for scripts 163
- parameters, Replication
 - configuring 144
 - connection 145
- partitions, Replication 243
- passencrypt utility 51, 98
- passwords
 - encrypting 51, 98
 - for repository database dba account, changing 78
 - for SCC default login account 15
- paths
 - multipath replication 227
 - multiple connections from the primary database 227
 - multiple connections to the replicate database 227
 - multiple dedicated routes 227
- performance counters, Replication
 - no data, troubleshooting 233
- permissions
 - Replication, configuring monitoring accounts 116
- Perspective Heat Chart view 167
- Perspective Resources view
 - about 183, 186
 - display tools and options 6
 - icons 5

- show/hide icon 4
- perspectives 243
 - about 186
 - adding resources 115, 184
 - creating 115, 187
 - removing 187
 - removing a resource 185
 - renaming 187
- pluggable authentication modules for UNIX
 - authentication 40, 86
- port conflicts 236
- ports
 - changing 78
 - configuring 51, 98
 - default 78
 - displaying information about 210
- postinstallation tasks 15
- production environment, setting up SCC in 19

Q

- queues
 - displaying status 221
- quick start instructions 13
- quiescing
 - Replication Server 222

R

- registration
 - about 183
- REJECT_DUPLICATE_RESOURCE_AND_COLLECTION policy errors 238
- RepAgent
 - hiding connection status 220
- RepAgent connection status
 - hiding 220
- RepAgent thread
 - configuring parameters 147
 - status in Adaptive Server shared disk cluster 226
- RepAgent thread, Replication
 - displaying status 226
 - starting 227
 - stopping 227
- RepAgent threads 243
- replication 243
 - data collections 122
 - heartbeat 230

Index

- latency 230
- rs_ticket 230
- servers 2
- topology 213
- tracing 214
- Replication Agent
 - about 225
 - configuring parameters 144
 - resuming 223, 226
 - suspending 223, 226
- Replication Agents 243
- Replication Agents, Replication
 - displaying status 19, 225
- replication environment
 - managing 213
 - monitoring 213
- replication monitors
 - connections 216
 - logical connections 216
 - overview 216
 - paths 216
 - queues 216
 - RepAgent threads 216
 - Replication Agents 216
 - Replication Servers 216
 - routes 216
- replication parameters
 - configuring 144
 - setting up 144
- replication path 243
- replication paths
 - displaying status 229
 - multipath replication 227
 - multiple connections from the primary
 - database 227
 - multiple connections to the replicate database 227
 - multiple dedicated routes 227
- Replication Server
 - configuring parameters 144
 - displaying status 18, 218
 - quiescing 222
 - register, troubleshooting 235
 - resuming 222
 - viewing status 18, 218
- Replication Servers
 - about 218
- replication sources
 - displaying 214
 - tracing 214
 - viewing 214
- replication targets
 - displaying 214
 - tracing 214
 - viewing 214
- replication topology 213
 - displaying 215
 - viewing 215
- repository 199, 243
 - backing up 201
 - changing backup schedule 201
 - changing database dba password 78
 - configuring purging 204
 - restoring from backup 202
 - scheduling backups 200
- resource explorer
 - launch icon 4
- Resource Explorer
 - about 183
 - display tools and options 6
 - searching in 186
- resources 243
 - about 183
 - adding to a perspective 115, 184
 - authenticating 17, 116, 117
 - displaying availability 17, 167
 - filtering by column in Admin Console 6
 - importing in batch 113
 - modifying data collection schedules 174
 - registering with Sybase Control Center 16, 112
 - removing from a perspective 185
 - searching for 186
 - unregistering 184
- restarts
 - configuring in UNIX 30, 71
 - configuring in Windows 27, 68
- resuming
 - connection 223
 - Replication Server 222
 - route 224
- role-mapping.xml file 49, 96
- roles
 - assigning to users and groups 56, 103
 - mapping SCC roles to LDAP or OS groups 49, 96
 - predefined 57, 109
 - product level 55, 103

- removing 104
- Replication, configuring monitoring accounts
 - 116
 - system level 55, 103
- roles, Replication
 - assigning to users and groups 115
- route
 - resuming 224
 - suspending 224
- routes
 - configuring parameters 146
 - displaying status 222
 - viewing status 222
- routes, Replication 244
 - unknown status, troubleshooting 234
- RSSD user name, using to authenticate 17, 116
- S**
- Save data collected from this job checkbox 120
- SCC Administrator group 49, 96
- SCC agent 244
 - deploying and managing instances 23, 63, 195
 - deploying instances from a shared disk 20, 60, 190
 - shared-disk mode 22, 61, 189
 - starting in UNIX 30, 71
 - starting in UNIX as a service 30, 71
 - starting in Windows 27, 68
 - starting in Windows as a service 27, 68
 - stopping in UNIX 30, 71
 - stopping in Windows 27, 68
- scc command 78
 - using to launch Sybase Control Center 14, 67
- SCC_MEM_MAX 34–36, 75–77, 238
- SCC_MEM_PERM 34–36, 75–77
- SCC-enabled login account 244
- scc.bat 14, 27, 68
- scc.sh 30, 71
- sccadmin account
 - about 15
- sccAdminRole 57, 109
- sccd shell script 30, 71
- sccinstance command 23, 63, 195
- sccUserRole 57, 109
- scheduler
 - resuming 174
 - suspending 174
- schedules 170, 244
 - adding to a job 172
 - creating for a data collection job 119
 - removing from a job 173
 - viewing 173
- screens
 - maximizing 8
 - maximizing and minimizing sections of a view 8
- scripts
 - example 150
 - substitution parameters 163
 - triggered by alerts 162
- security 38, 84
 - configuring 37, 83
- security providers
 - configuring 38, 84
- server license information, Replication
 - displaying 229
 - viewing 229
- server logs, viewing 206
- server status, Replication
 - monitoring availability 213
- server version, Replication
 - compatibility 3
- servers
 - authenticating 17, 116, 117
 - displaying availability 17, 167
 - importing in batch 113
 - modifying data collection schedules 174
 - registering with Sybase Control Center 16, 112
 - replication environment 2
 - searching for 186
 - unregistering 184
- servers, Replication
 - missing in topology view, troubleshooting 234
 - unknown status, troubleshooting 234
- services
 - enabling and disabling 78
 - listing 210
- services, UNIX
 - configuring SCC memory options for 36, 77
 - running SCC or SCC agent as 30, 71
- services, Windows
 - configuring SCC memory options for 35, 77
 - running SCC or SCC agent as 27, 68
- setting up
 - latency monitoring 117
 - replication parameters 144

Index

- severities for alerts 176
 - severities for alerts, Replication 157
 - shared-disk mode 244
 - about 22, 62, 194
 - enabling and disabling 22, 61, 189
 - shutdown command (console) 211
 - sorting by column 6
 - sql.ini files, importing resources from 113
 - stable queue manager, Replication 244
 - stable queues, Replication 244
 - start up
 - automatic, configuring in UNIX 30, 71
 - automatic, configuring in Windows 27, 68
 - starting Sybase Control Center 14, 67
 - states, Replication 157
 - statistics
 - about 121
 - availability 121
 - performance 121
 - statistics chart
 - displaying data for a longer period 237
 - effects of repository purging on 204
 - graphing performance counters 169
 - troubleshooting 237
 - stats_show_zero_counters, Replication 147
 - status command (console) 212
 - status icons and badges for resources 5
 - storm suppression for alerts 149
 - substitution parameters for scripts 163
 - suspending
 - connection 223
 - route 224
 - Sybase Control Center
 - accessibility 12
 - connecting a browser to 15
 - console commands 210
 - deploying and managing instances 23, 63, 195
 - deploying instances from a shared disk 20, 60, 190
 - display tools and options 6
 - failure to start 236
 - keyboard shortcuts 9
 - log files 206
 - logging in 36, 82
 - logging out 83
 - logging out unintentionally with F5 237
 - shared-disk mode 22, 61, 189
 - starting 14, 67
 - starting in UNIX 30, 71
 - starting in UNIX as a service 30, 71
 - starting in Windows 27, 68
 - starting in Windows as a service 27, 68
 - stopping in UNIX 30, 71
 - stopping in Windows 27, 68
 - Sybase Control Center for Replication
 - enhancements 1
 - new features 1
 - Sybase Control Center, Replication
 - About 1
 - managing 170
 - monitoring 170
 - sybase group 49, 96
 - system properties
 - displaying information about 210
 - system-wide features
 - configuring 37, 83
- ## T
- terms
 - SCC for Replication 241
 - text conventions 10
 - timeout
 - errors on data collections 238
 - setting for login sessions 55, 102
 - toolbar icons 4
 - topology
 - displaying 215
 - replication 213
 - server 213
 - viewing 215
 - topology data
 - unable to retrieve for Replication 235
 - topology view, Replication
 - blank canvas, troubleshooting 235
 - incorrect display, troubleshooting 235
 - topology, Replication 244
 - tracing
 - replication path 214
 - replication source 214
 - replication target 214
 - troubleshooting
 - controlling data flow 217
 - Replication 233
 - types of alerts 176
 - types of alerts, Replication 157

U

UNIX

- configuring authentication 40, 86
 - running SCC or SCC agent in the background 30, 71
 - running SCC or SCC agent in the foreground 30, 71
 - starting, stopping SCC or SCC agent 30, 71
- user accounts
- default 107
 - native SCC, adding 107
 - native SCC, not using 38, 84
- user information
- modifying 108
- user interface, about 4

V

- view layouts
- cascade 189
 - close all 189
 - horizontal tiling 189
 - minimize all 189
 - restore all 189
 - vertical tiling 189

View menu 8

viewing

- logical connection status 219

- queue status 221

- replication log 205

- route status 222

viewing Replication Server status 18, 218

viewing, Replication

- latency 230

- Replication Agent status 19, 225

- topology 215

views 244

- about 187

- bringing to front of perspective 188

- closing 188

- icons for managing 4

- maximizing 188

- maximizing and minimizing sections 8

- minimizing 188

- opening 188

- restoring 188

W

WARN logging level 207

Windows

- configuring authentication 39, 85

- starting, stopping Sybase Control Center or SCC agent 27, 68

