

SYBASE®

**Sybase Control Center for Adaptive Server®
Enterprise**

Sybase Control Center 3.1

DOCUMENT ID: DC01265-01-0310-01

LAST REVISED: February 2010

Copyright © 2010 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.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. 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.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc. 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 Adaptive Server	1
Display and Copy Options in Adaptive Server monitors	1
New Features in Sybase Control Center for Adaptive Server Enterprise	3
Configure	5
Configuring Adaptive Server for Monitoring	6
Registering an Adaptive Server	7
Importing Resources	9
Creating a Perspective	10
Adding a Resource to a Perspective	10
Role Assignment in Sybase Control Center for Adaptive Server	11
Authenticating a Login Account for a Monitored Resource	12
Setting Up Statistics Collection	12
About Statistics	14
Adaptive Server Data Collections	14
Key Performance Indicators for Adaptive Server	15
Setting Display Options for Adaptive Server Performance Data	20
Setting Adaptive Server Parameters in the Configuration File	21
Configuration Parameters for Adaptive Server	22
Creating an Alert	23
Adaptive Server Alert Templates	26
Alert Types, Severities, and States for Adaptive Server	30
Alert-triggered Scripts	30

Substitution Parameters for Scripts	31
Registering the Unified Agent for an Adaptive Server	33
Optional Configuration Steps	34
Manage and Monitor	35
Heat Chart	35
Displaying Resource Availability	35
Historical Performance Monitoring	36
Graphing Performance Counters	36
Manage Sybase Control Center for Adaptive Server	37
Settings	37
Viewing the Adaptive Server Component Log	38
Administer Adaptive Server	38
Registering the Unified Agent for an Adaptive Server	38
Authenticating the Unified Agent	39
Starting an Adaptive Server	40
Stopping an Adaptive Server	41
Viewing and Filtering the Adaptive Server Error Log	41
Manage and Monitor the Adaptive Server Environment	42
Clusters	43
Displaying the Performance Overview	52
Caches	55
Databases	62
Devices	66
Engines	69
Processes	70
Segments	76
Transactions	78
Server Configuration Values	79
SQL Activity	81
Replication Agents	81
Settings	83

Statistics	83
Troubleshoot Sybase Control Center for Adaptive Server	
.....	87
Error: Unable to Format the Date String	87
KPI is Not Updated	87
Invalid Connection Profile	88
Cannot Monitor Adaptive Server or Display Statistics	
Chart	88
Data on Screens or Charts Is Missing	88
Adaptive Server Is Responding Slowly	89
Error: No Result Set for this Query	90
Collection Job for Adaptive Server Fails	90
Cannot Authenticate Server Configured with a	
Multibyte Character Set	91
Glossary: Sybase Control Center for Adaptive Server	93
Index	97

About Sybase® Control Center for Adaptive Server

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

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

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

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

Display and Copy Options in Adaptive Server monitors

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

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

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

While monitoring Adaptive Server cluster configurations, you can also:

About Sybase® Control Center for Adaptive Server

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

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

New Features in Sybase Control Center for Adaptive Server Enterprise

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

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

Feature	Topics
New script to start and stop Sybase Control Center: <ul style="list-style-type: none"> In Windows: scc.bat In UNIX: scc.sh 	<i>Launching Sybase Control Center</i>
Use the scc script to start and stop servers, configure ports, enable and disable services, and more.	<i>scc Command</i>
Import and register servers in a batch from an interfaces file.	<i>Importing Resources</i> on page 9
Sybase Control Center automatically assigns user roles for monitoring.	<i>Role Assignment in Sybase Control Center for Adaptive Server</i> on page 11
Sybase Control Center automatically schedules default collections to gather data from managed resources.	<i>Adaptive Server Data Collections</i> on page 14
Tools for backing up and restoring the Sybase Control Center data repository.	<i>Repository</i>
Write a script and configure Sybase Control Center to execute it when a particular alert is triggered.	<i>Alert-Triggered Scripts</i> on page 30 <i>Creating an Alert</i> on page 23
Configure alerts on child objects (such as databases and caches) that are components of managed resources.	<i>Creating an Alert</i> on page 23
Monitor instances of an Adaptive Server cluster configuration including memory usage, device I/O, engine CPU utilization, and connection information.	<i>Cluster Instances</i> on page 43
Monitor interprocess communication in an Adaptive Server cluster configuration.	<i>Cluster Interconnect</i> on page 45
Monitor workload for an Adaptive Server cluster configuration including load profiles, load scores, routes, connections, and states.	<i>Workload Management</i> on page 48

New Features in Sybase Control Center for Adaptive Server Enterprise

Feature	Topics
Administrative options for Adaptive Server Enterprise that allow you to register and authenticate a running Unified Agent, start and stop your server, and view the server error log.	<i>Administer Adaptive Server</i> on page 38
Register and authenticate a Unified Agent for Adaptive Server	<i>Registering the Unified Agent for an Adaptive Server</i> on page 38 <i>Authenticating the Unified Agent</i> on page 39
Start and stop an Adaptive Server	<i>Starting an Adaptive Server</i> on page 40 <i>Stopping an Adaptive Server</i> on page 41
View the Adaptive Server error log	<i>Viewing and Filtering the Adaptive Server Error Log</i> on page 41
Monitor in-memory caches and databases	<i>Monitoring Adaptive Server In-memory Storage</i> on page 61
Identify and terminate blocking processes	<i>Identifying Blocked and Blocking Processes</i> on page 71 <i>Terminating Blocking Processes</i> on page 72

Configure

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

1. *Configuring Adaptive Server for Monitoring*

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

2. *Registering an Adaptive Server*

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

3. *Importing Resources*

Import and register multiple servers from an interfaces or `sql.ini` file.

4. *Creating a Perspective*

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

5. *Adding a Resource to a Perspective*

Add a resource to the current perspective.

6. *Role Assignment in Sybase Control Center for Adaptive Server*

With Sybase Control Center version 3.1, you no longer need to grant special roles for administrative or monitoring privileges on the Adaptive Server.

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

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

8. *Setting Up Statistics Collection*

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

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

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

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

11. *Creating an Alert*

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

12. *Registering the Unified Agent for an Adaptive Server*

Use the Adaptive Server Administration Dashboard to register the Unified Agent by providing the host name and port number.

13. *Optional Configuration Steps*

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

Configuring Adaptive Server for Monitoring

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

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

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

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

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

3. Use one of these methods to set monitoring configuration options in Adaptive Server.
 - Option 1: Use **`sp_configure`** to set the monitoring options to the values shown in the example below. (For information on using **`sp_configure`**, see the chapter on setting configuration parameters in the Adaptive Server *System Administration Guide*, Volume 1.)
 - Option 2: Edit the Adaptive Server configuration file manually:
 - a) Shut down Adaptive Server.
 - b) In a text editor, open the server's configuration file, found at:
 - Windows: `%SYBASE%\<Adaptive-Server-name>.cfg`
 - UNIX: `$(SYBASE)/<Adaptive-Server-name>.cfg`
 - c) Save a backup copy of the configuration file.
 - d) Search for the Monitoring section of the file.
 - e) Set the monitoring options to the values shown in the example below.
 - f) Save the file and exit.
 - g) Start Adaptive Server.

Example

This example shows the monitoring section of the configuration file. Set these options either using **`sp_configure`**, or by manually editing the file. The Adaptive Server Monitor uses all these parameters, and notifies you if any of these options is not enabled. You may have to

increase the values of **sql text pipe max messages** and **errorlog pipe max messages** depending on the level of activity on the monitored Adaptive Server.

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

Next

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

See also

- *Registering an Adaptive Server* on page 7
- *Authenticating a Login Account for a Monitored Resource* on page 12
- *Setting Adaptive Server Parameters in the Configuration File* on page 21
- *Displaying Configuration Values* on page 79
- *Setting Up Statistics Collection* on page 12
- *Role Assignment in Sybase Control Center for Adaptive Server* on page 11
- *Adding a Resource to a Perspective*

Registering an Adaptive Server

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

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

Table 2. 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 (15.0.2.0) – monitor Adaptive Server 15.0.2.0 or later. Choose this type for full Adaptive Server monitoring capabilities. • ASE Server, Replication Only (12.5.0.0) – monitor only the RepAgent threads for an Adaptive Server that is older than version 15.0.2.0. Choose this type for an Adaptive Server that is part of a replication environment.
Description	A brief description to help you identify the resource.

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

Table 3. New resource connection details

Field	Description
Server Host Name/Host Name	Local host name
Port Number	Local host port number
Character Set	Character set configured on Adaptive Server Note: If the Adaptive Server is configured to use a language that requires a multibyte character set such as Chinese, make sure to specify the correct character set in the connection profile.
Language	Language configured on Adaptive Server

5. Click **Finish**.

See also

- *Configuring Adaptive Server for Monitoring* on page 6
- *Importing Resources* on page 9

- *Resources*
- *Importing Resources* on page 9

Importing Resources

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.

An `interfaces` (UNIX) or `sql.ini` file (Windows) is a list of Sybase servers and their ports; it may contain other connection information as well. The file is created during the installation of a Sybase server:

- Windows: `%SYBASE%\ini\sql.ini`
- Unix: `$SYBASE/interfaces`

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

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

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

1. In the application menu, select **View > Open > Resource Explorer**.
2. In the Resource Explorer, select **Resources > Import**.
The Import Resources wizard opens; **Interfaces file** is already selected.
3. Click **Next**.
The Directory Service Connection page appears.
4. Click **Browse** and navigate to the `interfaces` file you want to import from.
You cannot type in the **File name** field.
5. Click **Next**.
6. On the Import Resource Type page, select the type of server you want to import.
7. On the Resource Selection page, click to select the servers you want to import.
Select only servers of the type you chose on the Import Resource Type page. If you import servers with incorrect types, Sybase Control Center will not be able to monitor or manage them properly.

- 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.
- Click **Next**.
The Confirmation page displays a list of the resources you have selected.
- Click **Finish** if you are ready to import, or click **Back** to return to the previous screen and change your selections.
When you click **Finish**, Sybase Control Center imports and registers the resources and displays a summary page.
- Click **Close** to finish the wizard.

The newly imported resources appear in the Resource Explorer.

Next

Add the new resources to a perspective so you can monitor and manage them.

See also

- *Registering an Adaptive Server* on page 7
- *Creating a Perspective* on page 10
- *Resources*

Creating a Perspective

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

- From the application menu bar, select **Perspective > Create**.
- Enter a name for your perspective. The name can contain up to 255 characters.
- Click **OK**.

See also

- *Importing Resources* on page 9
- *Adding a Resource to a Perspective* on page 10

Adding a Resource to a Perspective

Add a resource to the current perspective.

Add a server or other resource to a perspective so you can monitor and manage it along with other resources in the same perspective.

- From the Sybase Control Center toolbar, click the **Launch resource explorer** icon.

2. Select the resources to add to your perspective. Select multiple resources by pressing the **Ctrl** key while you select. Then 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

- *Creating a Perspective* on page 10
- *Role Assignment in Sybase Control Center for Adaptive Server* on page 11

Role Assignment in Sybase Control Center for Adaptive Server

With Sybase Control Center version 3.1, you no longer need to grant special roles for administrative or monitoring privileges on the Adaptive Server.

Sybase Control Center version 3.1 automatically assigns **aseMonitorRole** privileges to users who have **mon_role** privileges on an Adaptive Server, and lets them perform monitoring tasks on the server.

Sybase Control Center version 3.1 automatically assigns **sccAdminRole** privileges to users who have **sa_role** privileges on an Adaptive Server, and lets them perform administrative tasks on the server.

Note: You cannot grant **sa_role** and **mon_role** privileges through the Sybase Control Center interface.

Sybase Control Center checks for role validation every 30 minutes. The **monitor** option is greyed out on the Perspectives window until the resource is authenticated. If you try to authenticate a resource without having **mon_role** privileges, Sybase Control Center displays an error message.

Note: If a role is revoked on the Adaptive Server, Sybase Control Center will not register the change till the next role-check occurs. However, as the monitoring or administrative role has been revoked on the Adaptive Server, the user will not be able to successfully execute such a task through Sybase Control Center.

See also

- *Adding a Resource to a Perspective* on page 10
- *Authenticating a Login Account for a Monitored Resource* on page 12
- *Logins, Roles, and Groups*
- *Assigning a Role to a Login or a Group*

Authenticating a Login Account for a Monitored Resource

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

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

Note: You can also authenticate a server when you create a collection job.

1. Connect a browser to Sybase Control Center and log in.
2. In the Perspective Resources view, right-click the icon for your resource and select **Authenticate**.
3. Select **Use my current SCC login** or **Specify different credentials**.
4. If you chose **Specify different credentials**, enter the login and password for Sybase Control Center to use to connect to your resource.
5. If the selected server is a Replication Server, also enter the RSSD user name and password.
6. Click **OK** to exit the dialog.

See also

- *Role Assignment in Sybase Control Center for Adaptive Server* on page 11
- *Setting Up Statistics Collection* on page 12
- *Configuring Adaptive Server for Monitoring* on page 6
- *Setting Adaptive Server Parameters in the Configuration File* on page 21
- *Setting Up Statistics Collection* on page 12
- *User Authorization*

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.

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 60 seconds or more.
1. In the Perspective Resources view, select a resource and select **Resource > Properties**.
 2. Select **Collection Jobs**.

3. Click **Create Job**.

The Create Collection Job window opens.

4. If this resource has not been authenticated yet, the Authentication page appears. Enter a user name and password that Sybase Control Center can use to log in to the resource. Then click **Authenticate** to allow Sybase Control Center to verify your credentials.

5. Select the data collection that this job will run and click **Next**.

6. (Optional) If you do not want to create a schedule yet, click to uncheck **Create a schedule for this job**.

7. Specify details for the new schedule:

Table 4. New schedule details

Field	Description
Name	A name for this schedule
Description	A description of the schedule
Start date	The day when the schedule will take effect
Time	The time when the schedule will take effect

Note: Enter dates and times using the server's local time.

8. Choose an option to specify the duration of this schedule:

- **Run once**
- **Repeat until**

For **Repeat until**, specify these details:

Field	Description
Repeat interval	Interval, in seconds, between recurrences to be added to the schedule
Repeat until	End date
Time	Time when the job is to end

- **Repeat indefinitely**

For **Repeat indefinitely**, specify these details:

Field	Description
Repeat interval	Interval, in seconds, between recurrences to be added to the schedule

9. Click **Finish**.

See also

- *Authenticating a Login Account for a Monitored Resource* on page 12

- *Setting Display Options for Adaptive Server Performance Data* on page 20
- *Job Scheduling*

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 your Sybase Control Center product module for information on data collections and the KPIs contained in them.

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

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

Adaptive Server Data Collections

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

When an Adaptive Server is first authenticated, Sybase Control Center sets up a default collection of data called **collection_ase_availability**. The data in the default collection is

gathered in 60 second intervals. The KPIs used in the default collection include Server Percent CPU Utilization, Number of Blocked Processes, Number of Suspended Processes, and Server Availability State.

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

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

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

Set up jobs in the scheduler to collect these preconfigured statistics collections:

- **collection_ase_all_client_kpis** – Collects the data for historical charts in the Adaptive Server component, including those on the Overview, Devices, Engines, and Segments screens in the Adaptive Server monitor. Schedule this collection to see real time charting for these resources.
- **collection_ase_histmon** – Collects historical statistics for an Adaptive Server. These statistics are not displayed in the Adaptive Server monitor. To view these statistics, launch the 'Statistics Chart' window from the context menu. Schedule this collection to activate alerts and view them in a statistics chart.
- **collection_ase_rat** – Collects RepAgent Threads metrics for the charts on the Adaptive Server component's Replication Agent screen. Include this collection only if you are planning to monitor replication from this primary database.

Schedule the collections for every 60 seconds. See [Creating a Data Collection Job](#) (linked below) for instructions on using the scheduler.

See also

- *Key Performance Indicators for Adaptive Server* on page 15
- *Role Assignment in Sybase Control Center for Adaptive Server* on page 11
- *Configuring Adaptive Server for Monitoring* on page 6
- *Authenticating a Login Account for a Monitored Resource* on page 12
- *Setting Adaptive Server Parameters in the Configuration File* on page 21

Key Performance Indicators for Adaptive Server

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

Each Adaptive Server data collection includes a subset of the KPIs listed here.

Adaptive Server Objects	KPI	Description
Cluster Instances	Active Connections in Cluster Instance	Number of active connections to a cluster instance. This KPI is collected separately for each cluster instance.
	Workload Load Score in Cluster Instance	Load score in each cluster instance. This KPI is collected separately for each cluster instance.
	Number of Bytes Received in Cluster Instance	Number of bytes received during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Bytes Sent in Cluster Instance	Number of bytes sent during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of CIPC Messages Received in Cluster Instance	Number of CIPC messages received during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of CIPC Messages Sent in Cluster Instance	Number of CIPC messages sent during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Committed Transactions in Cluster Instance	Number of committed transactions in the cluster instance during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Packets Received in Cluster Instance	Number of packets received during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Packets Sent in Cluster Instance	Number of packets sent during the current collection cycle. This KPI is collected separately for each cluster instance.
	Device IO Rate in Cluster Instance	This KPI is collected separately for each cluster instance and is used to generate the Instance connection chart on the Cluster Instance screen.

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

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

Adaptive Server Objects	KPI	Description
	Number of Packets Received in Network IO	Number of packets received during the current collection cycle.
	Number of Packets Sent in Network IO	Number of packets sent during the current collection cycle.
	Number of Page Locks	Number of page-level locks server-wide.
	Number of Row Locks	Number of row-level locks server-wide.
	Number of Suspended Processes	Number of processes that are currently suspended. (The Heat Chart uses this metric to display server status.)
	Number of Table Locks	Number of table-level locks server-wide.
	Number of Transactions	<p>Total number of transactions during the current collection cycle.</p> <hr/> <p>Note: This KPI is available only on Adaptive Server versions 15.0.3 Cluster Edition and 15.0.3 ESD #3, and later.</p> <hr/>
	Number of User Connections	Current number of user connections on the server.
	Procedure Cache Hit Ratio	Hit ratio in the procedure cache.
	Resource State	Status of the Adaptive Server. Values of most interest are STOPPED and RUNNING.
	Server CPU Utilization	Average CPU utilization percentage across all active Adaptive Server engines on the server.
	Server Device IO Rate	Total number of I/O operations performed by all devices on the server during the current collection cycle.

Adaptive Server Objects	KPI	Description
	sp_who Response Time	Time in milliseconds the sp_who stored procedure takes to return a response. sp_who is called each time collection_ase_histmon is executed to collect Adaptive Server performance statistics.
	Statement Cache Hit Ratio	Hit ratio in the statement cache during the current collection cycle.
	Server tempdb Free Space	Amount, in megabytes, of free space in the tempdb database.
	Server tempdb Space Used	Amount, in megabytes, of space used in the tempdb database.
TempDBs Activity	Cluster Temp DB IO Rate	The rate of tempdb IO activity, per second, in a cluster.

See also

- *Adaptive Server Alert Templates* on page 26

Setting Display Options for Adaptive Server Performance Data

Change the screen refresh interval, chart trend period, alert list size, historical SQLs size, and historical SQLs trend period for this Adaptive Server. These options are located on the Settings screen.

Follow these steps to set the options:

1. Right-click the server you want to configure in the Perspective Resources view and select **Monitor**.
2. Select **Settings** from the left panel.
3. Enter a new value in the **Screen Refresh Interval** field. Refreshing a screen redraws it with the most recent available data. The screen refresh interval is the period between refreshes, with a default of 30 seconds.
4. Enter a new value in the **Chart Trend Period** field. The chart trend period is the amount of time covered by historical charts in the Adaptive Server component. The default is 15 minutes.
5. Enter a new value in the **Alert list size** field. The Alerts table, in the Overview window, contains a list of all alerts configured for a server. The maximum number of rows in the Alerts table is indicated by the alert list size, with a default of 100.

6. Enter a new value in the **Historical SQLs size** field. The active SQLs table, in the SQL Activity window, contains a list of active SQL statements. The maximum number of statements in this table is denoted by the historical SQLs size, with a default of 500.
7. Enter a new value in the **Historical SQLs trend period** field. The list of active SQL statements in the active SQLs table are displayed for a maximum period of time denoted by the historical SQLs trend period. The default value for this setting is 5 minutes.
8. Click **Apply Settings**.

See also

- *Setting Up Statistics Collection* on page 12
- *Setting Adaptive Server Parameters in the Configuration File* on page 21

Setting Adaptive Server Parameters in the Configuration File

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

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

```
Windows: %SYBASE%\SCC-3_1\plugins\ASEMap\config.properties
UNIX: $SYBASE/SCC-3_1/plugins/ASEMap/config.properties
```

1. Open the `config.properties` file with a text editor.
2. Save a backup copy of the file.
3. (Optional) To change the number of times Sybase Control Center tries to reopen a broken JDBC connection to Adaptive Server, edit the value of `attempts_reopen_con`.
4. (Optional) To change the interval (in seconds) between successive attempts to reopen a broken JDBC connection, edit the value of `time_between_reattempts`.
5. (Optional) To change the interval (in hours) between refreshes of the list of monitored objects in the Sybase Control Center repository, edit the value of `revalidation_frequency`.
6. (Optional) To change the time interval (in seconds) for JDBC internal queries to execute before they time out, edit the value of `jdbc_internal_query_timeout`.
7. Save and exit the file.
8. To make the new settings take effect, restart Sybase Control Center.

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

See also

- *Setting Display Options for Adaptive Server Performance Data* on page 20
- *Creating an Alert* on page 23

Configuration Parameters for Adaptive Server

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

Sybase Control Center provides these configuration parameters on Adaptive Server:

- `attempts_reopen_con` – Changes the number of times Sybase Control Center tries to reopen a broken JDBC connection to Adaptive Server. The default value is 5. This parameter only tries to re-establish broken connections—it does not control retries when a first attempt to connect to the server fails.
- `time_between_reattempts` – Changes the interval (in seconds) between successive attempts to reopen a broken JDBC connection. The default value is 6. This parameter only tries to re-establish broken connections—it does not control retries when a first attempt to connect to the server fails.

- `revalidation_frequency` – Changes the interval (in hours) between refreshes of the list of monitored objects in the Sybase Control Center repository.

The repository stores a list of the caches, devices, engines, and segments for all monitored Adaptive Servers. Sybase Control Center does not collect statistics for objects unless they are listed in the repository. The `revalidation_frequency` parameter controls how often Sybase Control Center refreshes (revalidates) the list of objects associated with each monitored server. The default value is 24.

- `jdbc_internal_query_timeout` – Changes the time interval (in seconds) for JDBC internal queries to execute before they time out. The default value is 15. This prevents server connection attempts and other operations from hanging when resources on the server have been exhausted. You may want to modify this if your ASE server is up, but performance is so slow that the ASEMAP agent fails to successfully establish a connection to the server. This failure usually indicates a configuration problem on the ASE server and not in the SCC ASEMAP component. Sybase recommends that you contact Sybase Customer Support to determine the cause of connection problems before modifying this parameter.
- `query_timeout` – The default query timeout for any statements created on this connection. The default value is 30 seconds.
- `query_timeout_spaceused` – The query timeout for **sp_spaceused**, which is used to determine the space usage parameters for a database. The default value is 180 secs,

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

Creating an Alert

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

Prerequisites

- You must have administrative privileges (sccAdminRole) to perform this task.
- Specify an e-mail server for Sybase Control Center to use for alerts. If an e-mail server is not configured, you cannot create alert subscriptions.
- 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.

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

Note: Only alerts on Adaptive Server or Replication Server resources can trigger script execution.

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.
3. On the Resource page of the wizard, select the object on which you want to set the alert. Expand the folder representing the server or agent to select lower-level child objects.
4. Click **Next**.
The Type page of the wizard appears.
5. 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**.

The Storm Suppression page of the wizard appears.

8. Enter the storm suppression period, which suppresses redundant alerts resulting from the same condition (an alert storm) for the specified period of time. (The default is 60 seconds.) Enter this value in seconds, minutes, or hours in **Storm Suppression Period** and click **Next**.

9. (Optional; not available for alerts on Sybase IQ servers) To configure this alert to trigger the execution of a script:

a) **Type Configuration** specifies the alert severity level that triggers the script. Select **Critical**, **Warning**, or both.

Critical is more serious than Warning.

b) Enter or browse to the location of the script.

c) Enter parameter values in the **Execution Parameters** box.

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

d) Click **Next**.

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

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

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

11. (Optional) On the Subscription page, specify e-mail addresses if you want this alert to issue e-mail notifications when it fires.

The e-mail addresses default to the address in your user profile, but you can override the defaults.

For both critical and warning alerts:

Table 5. 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.

Option	Description
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. Note that you cannot enter an escalation address unless you enter an address for primary notification first.
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.

12. Click **Finish**.

Example: An alert-triggered script

This sample script is a Windows .bat file:

```
@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:
```

See also

- *Setting Adaptive Server Parameters in the Configuration File* on page 21
- *Registering the Unified Agent for an Adaptive Server* on page 33
- *Assigning a Role to a Login or a Group*
- *Configuring the E-mail Server*
- *Alert-triggered Scripts* on page 30
- *Alerts*
- *Adaptive Server Data Collections* on page 14
- *Adaptive Server Alert Templates* on page 26
- *Substitution Parameters for Scripts* on page 31
- *Key Performance Indicators for Adaptive Server* on page 15
- *Alert Types, Severities, and States for Adaptive Server* on page 30

Adaptive Server Alert Templates

Lists and describes alert templates you can use to create alert instances for Adaptive Server.

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

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

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

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

Adaptive Server Objects	Alert Template	Description
	Number of CIPC Messages Sent in Cluster Instance	Number of messages sent during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Packets Received in Cluster Instance	Number of packets received during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Packets Sent in Cluster Instance	Number of packets sent during the current collection cycle. This KPI is collected separately for each cluster instance.
	Number of Committed Transactions in Cluster Instance	Number of committed transactions in the cluster instance during the current collection cycle. This KPI is collected separately for each cluster instance.
	Workload Load Score in Cluster Instance	Load score in each cluster instance. This KPI is collected separately for each cluster instance.
Cluster Workload	CPU Busy Value for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the CPU busy chart on the Workload screen.
	IO Load Value for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the IO load chart on the Workload screen.
	Load Score for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the load score chart on the Workload screen.
	Run Queue Length for Logical Cluster Workload	This KPI is collected separately for each cluster workload and is used to generate the run queue length chart on the Workload screen.
Data Caches	Cache Hit Ratio	Hit ratio in the data cache during the current collection cycle.
Devices	Device IO Rate	Rate of I/O operations per second on this device.
	Device IO Response Time	Response time, in milliseconds, for I/O operations performed on this device.
	Device Free Space	Total amount of free space, in megabytes, on this device.

Adaptive Server Objects	Alert Template	Description
Engines	Engine CPU Utilization	Percentage of CPU cycles used by this Adaptive Server engine.
Logical Clusters	Active Connections in Logical Cluster	Number of active connections using the Logical Cluster at the time of collection. This KPI is collected separately for each Logical Cluster.
	Number of Failover Instances in Logical Cluster	Number of instance that are failed-over in the currently active Logical Cluster. This KPI is collected separately for each logical cluster.
Replication Agent	RepAgent Thread State Change	The alert to send when an Adaptive Server RepAgent Thread changes state.
	Transaction Log Size	The size of an Adaptive Server RepAgent Thread's transaction log.
	Number of Log Operations Scanned per Second	The number of operations scanned by an Adaptive Server RepAgent thread.
	Number of Log Operations Processed per Second	The number of operations processed by an Adaptive Server RepAgent thread.
Segments	Segment Free Space	Amount of free space in the segment. This KPI is collected separately for each segment.
Server	Average Blocked Process Wait Time	Average time, in milliseconds, that the current blocked processes have waited.
	Number of Address Locks	Number of address-level locks server-wide.
	Number of Blocked Processes	Number of currently blocked processes that have been blocked for more than 5 seconds. The Heat Chart uses this metric to display server status.
	Number of Bytes Received in Network IO	Number of bytes received during the current collection cycle.
	Number of Bytes Sent in Network IO	Number of bytes sent during the current collection cycle.

Adaptive Server Objects	Alert Template	Description
	Number of Deadlocks	Number of deadlocks on the server since the most recent execution of the collection.
	Number of Locks	Total number of active locks of all types on the server.
	Number of Packets received in Network IO	Number of packets received during the current collection cycle.
	Number of Packets sent in Network IO	Number of packets sent during the current collection cycle.
	Number of Page Locks	Number of page-level locks server-wide.
	Number of Row Locks	Number of row-level locks server-wide.
	Number of Suspended Processes	Number of processes that are currently suspended. The Heat Chart uses this metric to display server status.
	Number of Table Locks	Number of table-level locks server-wide.
	Number of Transactions	Total number of transactions during the current collection cycle.
	Number of User Connections	Current number of user connections on the server.
	Procedure Cache Hit Ratio	Hit ratio in the procedure cache.
	Resource State	Status of the Adaptive Server. Values of most interest are STOPPED and RUNNING.
	Server CPU Utilization	Average CPU utilization percentage across all active Adaptive Server engines on the server.
	Server Device IO Rate	Total number of I/O operations performed by all devices on the server during the current collection cycle.
	Server tempdb Free Space	Amount of free space in the tempdb database, in megabytes.
	Server tempdb Space Used	Amount of space used by the tempdb database, in megabytes.

Adaptive Server Objects	Alert Template	Description
	sp_who Response Time	Time in milliseconds the sp_who stored procedure takes to return a response. sp_who is called each time collection_ase_histmon is executed to collect Adaptive Server performance statistics.
	Statement Cache Hit Ratio	Hit ratio in the statement cache during the current collection cycle.

See also

- *Key Performance Indicators for Adaptive Server* on page 15

Alert Types, Severities, and States for Adaptive Server

Table 6. Adaptive Server States

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

See also

- *Alert-triggered Scripts* on page 30
- *Alerts*
- *Assigning a Role to a Login or a Group*
- *Configuring the E-mail Server*
- *Creating an Alert* on page 23

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 e-mail 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.

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

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 7. Substitution Parameters for State-based Alerts

Parameter	Description
%Alert%	A three-part name supplied by the alert system. It includes the names of this alert, the resource, and the key performance indicator 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 key performance indicator 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 application on the managed resource that generated this alert.
%Time%	The date and time at which the alert fired, in this format: Tue Sep 15 10:10:51 2009

Parameter	Description
%Top_resource%	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, %Top_resource% and %Resource% have the same value.

Table 8. Substitution Parameters for Threshold-based Alerts

Parameter	Description
%Alert%	A three-part name supplied by the alert system. It includes the names of this alert, the resource, and the key performance indicator on which this alert is based.
%Datapoint%	The current value, on the alerted resource, of the KPI on which this alert is based.
%KPI%	The name of the key performance indicator 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 application on the managed resource 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
%Top_resource%	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, %Top_resource% and %Resource% have the same value.

Registering the Unified Agent for an Adaptive Server

Use the Adaptive Server Administration Dashboard to register the Unified Agent by providing the host name and port number.

You must register and authenticate the Unified Agent in order to use Sybase Control Center to perform any administrative tasks such as starting or stopping the Adaptive Server, or viewing the Adaptive Server error log.

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

For information on configuration options for the Unified Agent, see the chapter on configuring Unified Agent in the *Unified Agent and Agent Management Console Users Guide*.

For information on security features for the Unified Agent, see the chapter on security in the *Unified Agent and Agent Management Console Users Guide*.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard opens and displays the administrative options.
2. Click **Register Agent**.
3. Enter the host name and port number for the Unified Agent.
4. Click **OK**.
You see the Administration Dashboard where you can authenticate the Unified Agent.
5. (Optional) Click **Edit Agent Registration** to change the host name or port number for the Unified Agent.

See also

- *Creating an Alert* on page 23
- *Optional Configuration Steps* on page 34
- *Authenticating the Unified Agent* on page 39
- *Starting an Adaptive Server* on page 40
- *Stopping an Adaptive Server* on page 41
- *Viewing and Filtering the Adaptive Server Error Log* on page 41

Optional Configuration Steps

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

Table 9. 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>
Authentication	Add authentication modules to allow Windows, UNIX, or LDAP users to log in to Sybase Control Center.	<i>Setting Up Security</i>
Alerts	Modify alert thresholds and subscriptions and delete alerts.	<i>Alerts</i>
Data collection	Modify collection intervals and schedules, suspend and resume the schedule, and delete collection jobs.	<i>Job Scheduling</i>
Resources	Unregister resources, add them to perspectives, or remove them.	<i>Resources</i>
Perspectives	Create, remove, and rename perspectives.	<i>Perspectives</i>
Repository	Set purging options and schedule backups of the repository database.	<i>Repository</i>

See also

- *Registering the Unified Agent for an Adaptive Server* on page 33

Manage and Monitor

Manage and monitor the Adaptive Server, Unified Agent and Replication Agents.

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. You can also search and sort the results by column. From within the Perspective Heat Chart, you can right-click a resource to see a menu of monitoring and administrative options that vary based on the resource type.

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

Displaying Resource Availability

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

1. From the application menu bar, select **View > Open > Heat Chart**.
2. (Optional) To display tools for filtering (narrowing the list of resources in the heat chart) or changing the columns, select **View > Filter** from the Perspective Heat Chart menu bar. The Filter and Column tools appear in the left pane.
3. (Optional) To use filtering, select **View > Filter** from the view's menu bar and enter a search term in the **Filter string** field.

The search term can be any string that appears in the tabular portion of the heat chart, such as the name, or part of the name, of a server or a resource type (ASE Server, for example).

4. (Optional) Select a filtering setting:
 - **Match case** – search for resources whose displayed data includes the search term, including uppercase and lowercase letters; or
 - **Exact match** – search for resources whose displayed data includes an item identical to the search term.
5. (Optional) Select a column from the **Filter on** list to restrict your search to that column.
6. (Optional) Click **Columns** to customize your heat chart.

7. (Optional) Unselect any column that should not appear in your heat chart.
8. (Optional) Click the sorting arrow in the column headers to sort the column values in either ascending or descending order.
9. (Optional) Right-click the resource's row to view a menu of options for the selected resource.
10. (Optional) To resize the Filter and Columns tools pane, move your mouse over the border between the tools pane and the resource table. When the mouse cursor changes to a resize icon, click and drag the border to the left or the right.
11. (Optional) To hide the Filter and Columns tools, unselect **View > Filter**.

Historical Performance Monitoring

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

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

Graphing Performance Counters

Generate a graph for any set of performance counters to show performance trends.

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 is displayed, data is being collected.

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 additional 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.
7. (Optional) You can click the date/time that appear above the slider. Use these to change the start and end date/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.

Manage Sybase Control Center for Adaptive Server

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

Settings

Learn about Adaptive Server monitoring controls on the Settings screen.

Table 10. Controls on the Settings screen

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

Viewing the Adaptive Server Component Log

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

Prerequisites

Verify that you are in an Adaptive Server monitoring view.

Use either of these methods to launch the Sybase Control Center Log Viewer:

- From the Adaptive Server 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.
- Click **View > Open > Log Window**

Administer Adaptive Server

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

To start or stop the server, and to view the error log, you must be granted the **uafagent** users role on the host machine. Use the options in the Adaptive Server Administration Dashboard to manage the Unified Agent and Adaptive Server.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard opens and displays the administrative options.
2. To register the Unified Agent, click **Register Agent**. You must specify connection information that the Unified Agent needs to manage your resources remotely.
3. To authenticate the user with the Unified Agent, click **Authenticate Agent**.
4. To start the Adaptive Server using a server RUN script, click **Start ASE Server**.
5. To stop a running Adaptive Server, click **Shutdown ASE**.
6. To view the error log of an Adaptive Server, click **View ASE Errorlog**. You can further choose to filter the error log.

Registering the Unified Agent for an Adaptive Server

Use the Adaptive Server Administration Dashboard to register the Unified Agent by providing the host name and port number.

You must register and authenticate the Unified Agent in order to use Sybase Control Center to perform any administrative tasks such as starting or stopping the Adaptive Server, or viewing the Adaptive Server error log.

You must register a Unified Agent for each Adaptive Server you have configured. The Unified Agent is configured on the same host as the Adaptive Server that it manages. When you

register the Unified Agent, you are updating Sybase Control Center with information on the machine and port number on which the Unified Agent is configured.

For information on configuration options for the Unified Agent, see the chapter on configuring Unified Agent in the *Unified Agent and Agent Management Console Users Guide*.

For information on security features for the Unified Agent, see the chapter on security in the *Unified Agent and Agent Management Console Users Guide*.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard opens and displays the administrative options.
2. Click **Register Agent**.
3. Enter the host name and port number for the Unified Agent.
4. Click **OK**.
You see the Administration Dashboard where you can authenticate the Unified Agent.
5. (Optional) Click **Edit Agent Registration** to change the host name or port number for the Unified Agent.

Authenticating the Unified Agent

Authenticate the Unified Agent using the **Authenticate Agent** option in the Adaptive Server Administration Dashboard

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

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard displays the administrative options.
2. Click **Authenticate Agent**.
You see the ASE Agent Authentication window.
3. Enter the agent user name and password (optional) for the Unified agent.
You must have configured the Unified agent to allow the Sybase Control Center user to authenticate. The Sybase Control Center user must also have the necessary permissions to manage the Adaptive Server using the Unified agent.

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

4. Click **OK**.
5. (Optional) Click **Edit Agent Authentication** to change the username or password, or to clear the authentication.

See also

- *Registering the Unified Agent for an Adaptive Server* on page 33
- *Starting an Adaptive Server* on page 40
- *Stopping an Adaptive Server* on page 41
- *Viewing and Filtering the Adaptive Server Error Log* on page 41

Starting an Adaptive Server

Start an Adaptive Server using a RUN server script.

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

Use the **Start Adaptive Server** option in the Adaptive Server Administration Dashboard to start the Adaptive Server.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard displays the administrative options.
2. Click **Start Adaptive Server**.

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

3. Select the existing RUN server script to start Adaptive Server, or specify the location of a new RUN server file.
You must have permissions to execute the RUN server script.
4. (Optional) Click to show the server error log.
If you do not have permissions to read the error log, or the error log does not exist, you see an error message. If you have clicked on the **Show server log** option, you see the server start-up status.
5. Click **Finish**.

See also

- *Registering the Unified Agent for an Adaptive Server* on page 33
- *Authenticating the Unified Agent* on page 39
- *Stopping an Adaptive Server* on page 41

- *Viewing and Filtering the Adaptive Server Error Log* on page 41

Stopping an Adaptive Server

Shut down an Adaptive Server using the process ID.

To shut down the Adaptive Server, use the **Shut Down Adaptive Server** option in the Adaptive Server Administration Dashboard.

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

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard displays the administrative options.
2. Click **Shut Down Adaptive Server**.
You see the table of current server processes.

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

3. Select the option to shut down the server immediately, or to remind you, after a certain period of time, if the server has not yet shut down.
You must have permissions to see the error log.
4. Click **Yes**.

See also

- *Registering the Unified Agent for an Adaptive Server* on page 33
- *Authenticating the Unified Agent* on page 39
- *Starting an Adaptive Server* on page 40
- *Viewing and Filtering the Adaptive Server Error Log* on page 41

Viewing and Filtering the Adaptive Server Error Log

Apply filters and view the Adaptive Server error log.

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

Use the **View Adaptive Server's Error Log** option in the Adaptive Server Administration Dashboard to start the Adaptive Server.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Administer**.
The Adaptive Server Administration Dashboard displays the administrative options.
2. Click **View Adaptive Server's Error Log**.
You see the error log of the Adaptive Server.
3. (Optional) Select the **Filter** option.
You see the Server Log Properties window.

Note: User-defined regular expressions are only valid in the session that they are defined in. When the Sybase Control Server is restarted, you will have to define the regular expressions again.

4. Select how you want to retrieve log entries from the error log. You can retrieve all log entries, a specified number of log entries last accessed, all log entries from a specified time period, or all log entries matching a regular expression.
The pre-defined regular expressions provided by the dialog are generally sufficient for filtering. They contain the most commonly accessed keywords in the error log such as **Error**, **deadlock**, **Warning**, and so on.
However if you wish to define your own custom regular expression, click on **add**, and specify a descriptive name for the new regular expression, and the regular expression clause. For example, specify "All Lines that have configuration" for **Name** and "configuration" for **Regular Expression**.
5. (Optional) Select a different display icon to match different regular expressions.
6. Click **OK**.

See also

- *Registering the Unified Agent for an Adaptive Server* on page 33
- *Authenticating the Unified Agent* on page 39
- *Starting an Adaptive Server* on page 40
- *Stopping an Adaptive Server* on page 41

Manage and Monitor the Adaptive Server Environment

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

Clusters

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

Displaying the Cluster Overview

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

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

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

See also

- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47
- *Workload Management Statistics and Details* on page 50
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49
- *Performance Overview Statistics and Details* on page 53
- *Displaying the Performance Overview* on page 52

Cluster Instances

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

See also

- *Displaying the Cluster Overview* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47

- *Workload Management Statistics and Details* on page 50
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49

Monitoring Cluster Instances in Adaptive Server Cluster Configurations

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

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

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

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

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

See also

- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49
- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47

- *Workload Management Statistics and Details* on page 50

Cluster Instances Statistics and Details

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

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

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

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

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

Table 11. Tabs on the Cluster Instances window

Details	Displays the Memory Usage and Processes charts.
Advanced	Displays the Device I/O per Sec , Active Connections , and Engine CPU Utilization graphs.

See also

- *Cluster Interconnect Statistics and Details* on page 47
- *Workload Management Statistics and Details* on page 50
- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49

Cluster Interconnect

Monitor interprocess communication in an Adaptive Server cluster configuration.

See also

- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Workload Management* on page 48
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47
- *Workload Management Statistics and Details* on page 50
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49

Monitoring Interprocess Communication in Adaptive Server Cluster Configurations

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

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

1. In the Perspective Resources view, select a shared-disk cluster and click **Monitor**. This opens the Adaptive Server monitor.
2. To show the Cluster Interconnect screen, select **Cluster Interconnect** from the left pane. The Cluster Interconnect window displays information pertaining to each cluster instance, including instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches. Additional tabs show more information about CIPC (cluster interprocess communication) links, messages, channels, and endpoints.
3. (Optional) Click a tab at the top of the screen to display more CIPC information on individual instances. The tab **CIPC** is selected by default.
 - **CIPC** – displays the current status for each CIPC link. Status information for each instance includes the instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches.
 - **CIPC Links** – displays the current status for each CIPC link. Status information for each instance includes the ID, local and remote interfaces, passive and active states, and ages of passive and active states.
 - **CIPC Mesh** – displays cluster instance name, channel name, far-end cluster instance, received and transmitted message counts, and dropped, re-sent and retried messages. There is a Message Send Queue summary with send queue and sent queue information. The send queue includes current messages waiting to be sent to the instance, while the sent queue includes sent messages whose notification has not been processed.
 - **CIPC EndPoints** – displays cluster instance name, endpoint name, received and transmitted message counts, and received and transmitted byte counts. There is also a

Message Received Summary with received queue and done queue information. The received queue includes messages queued for this logical endpoint. The done queue includes messages for this logical endpoint that have been processed, and await further action.

See also

- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49
- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47
- *Workload Management Statistics and Details* on page 50

Cluster Interconnect Statistics and Details

Interpret the Interconnect window for Adaptive Server cluster configurations.

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

Table 12. Tabs on the Cluster Interconnect window

CIPC	Displays the current status for each CIPC link. Status information for each instance includes the instance name, count of received, transmitted, multicast and retransmitted messages, and count of successful and failed switches.
CIPC Links	Displays the current status for each CIPC link. Status information for each instance includes the ID, local and remote interfaces, passive and active states, and ages of passive and active states.

CIPC Mesh	Displays cluster instance name, channel name, far-end cluster instance, received and transmitted message counts, and dropped, re-sent and retried messages. There is a Message Send Queue summary with send queue and sent queue information. The send queue includes current messages waiting to be sent to the instance, while the sent queue includes sent messages whose notification has not been processed.
CIPC Endpoints	Displays cluster instance name, endpoint name, received and transmitted message counts, and received and transmitted byte counts. There is also a Message Received Summary with received queue and done queue information. The received queue includes messages queued for this logical endpoint. The done queue includes messages for this logical endpoint that have been processed, and await further action.

See also

- *Cluster Instances Statistics and Details* on page 45
- *Workload Management Statistics and Details* on page 50
- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49

Workload Management

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

See also

- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47

- *Workload Management Statistics and Details* on page 50
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49

Monitoring Workloads in Adaptive Server Cluster Configurations

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

1. In the Perspective Resources view, select a shared-disk cluster and click **Monitor**. This opens the Adaptive Server monitor. Select **Workload Management** from the left pane.
2. (Optional) Click a tab at the top of the screen to display more workload information on individual instances. The tab **Logical Clusters** is selected by default.
 - **Logical Clusters** – displays the logical cluster name, state, active connections, base instances, active base instances, failover instances, active failover instances, and load profile name. The possible values for state are: online, offline, failed, inactive, and time_wait. The type may be application, login, or alias. The General, Base Instances, Failover Instances, and Routes tabs at the bottom of the screen provide access to further information about a selected logical cluster.
 - (Optional) To see the system view, start-up mode, failover mode, down-routing mode, login distribution mode, and logical cluster role for the selected logical cluster, select **General**.
 - (Optional) To see the ID, name, and state of the selected logical cluster select **Base Instances**.
 - (Optional) To see the ID, name, state, and failover group of the selected logical cluster, select **Failover Instances**.
 - (Optional) To see the name and type of route associated with the selected logical cluster, select **Routes**.
 - **Workloads** – displays a list of workloads with information on instance, load profile, load score, user connections, CPU busy, run queue length, I/O load, engine deficit, and user score.
 - (Optional) To see charts depicting load score, and percentage use of the CPU, select **Details** for a selected workload.
 - (Optional) To see charts depicting queue length, and I/O load, select **Advanced** for a selected workload.
 - **Load Profiles** – displays a list of load profiles with information on name, type, minimum load score, login redirection, and dynamic migration.
 - (Optional) To see a chart of metrics, select **Metric Weight** for a selected load profile. The metrics include user connections, CPU busy, run queue length, and so on, and the corresponding weights associated with the metric.

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

See also

- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44
- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47
- *Workload Management Statistics and Details* on page 50

Workload Management Statistics and Details

Interpret the Workload Management window for Adaptive Server cluster configurations.

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

Table 13. Tabs on the Workload Management window

<p>Logical Clusters</p>	<p>Displays the logical cluster name, state, active connections, base instances, active base instances, failover instances, active failover instances, and load profile name. The possible values for state are: online, offline, failed, inactive, and time_wait. The type may be application, login, or alias.</p> <p>Routes depicts the name and type of route associated with the selected logical cluster.</p> <p>Failover Instances depicts the ID, name, state, and failover group of the selected logical cluster.</p> <p>Base Instances depicts the ID, name, and state of the selected logical cluster.</p> <p>General depicts the system view, start-up mode, failover mode, down-routing mode, login distribution mode, and logical cluster role for the selected logical cluster.</p>
--------------------------------	--

<p>Workloads</p>	<p>Displays a list of workloads with information on instance, load profile, load score, user connections score, CPU busy score, run queue length score, I/O load score, engine deficit score, and user score.</p> <p>Details displays charts depicting load score, and percentage use of the CPU, for a selected workload.</p> <p>Advanced displays charts depicting run queue length score, and I/O load score, for a selected workload.</p> <hr/> <p>Note: Elements of the Adaptive Server monitor screens may be updated at different intervals depending on the mechanism that is used to collect the data. For example, the historical charts are updated based on the frequency of the all_client_kpis collection while the data in the tables on the screen are updated with the frequency that the user has set for the screen refresh interval on the Settings window.</p>
<p>Load Profiles</p>	<p>Displays a list of load profiles with information on name, type, minimum load score, login redirection, and dynamic migration.</p> <p>Metric Weight for a selected load profile displays user connections, CPU busy, run queue length, and so on, and the corresponding weights associated with the metric.</p>
<p>Routes</p>	<p>Displays a list of routes with the name of the route, the logical cluster it is defined on, and the type of route such as alias, application, or login.</p>

See also

- *Cluster Instances Statistics and Details* on page 45
- *Cluster Interconnect Statistics and Details* on page 47
- *Displaying the Cluster Overview* on page 43
- *Cluster Instances* on page 43
- *Cluster Interconnect* on page 45
- *Workload Management* on page 48
- *Monitoring Cluster Instances in Adaptive Server Cluster Configurations* on page 44

- *Monitoring Interprocess Communication in Adaptive Server Cluster Configurations* on page 46
- *Monitoring Workloads in Adaptive Server Cluster Configurations* on page 49

Displaying the Performance Overview

The Overview screen shows Adaptive Server performance status.

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

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

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Licenses** - displays a list of licenses that are currently checked out by the server or cluster instances. There is also information about the number, type, status (expirable, permanent and so on), and expiration date of each license.
 - **Alerts** - displays a list of all fired alerts configured at server, cluster or cluster instance levels. For each alert, there is information about the time at which the alert was fired, severity, current statistic and threshold.
7. (Optional - Adaptive Server Cluster configurations) Mouse over the **Cluster Instances** graph to display precise information for points on the bar graph.
 8. (Optional - Adaptive Server Cluster configurations) If data collections are running, mouse over the **Logical Cluster** graph to display precise information for points on the bar graph.

See also

- *Displaying the Cluster Overview* on page 43
- *Interpreting Statistics* on page 85
- *Performance Overview Statistics and Details* on page 53

Performance Overview Statistics and Details

Learn about the information presented on the Overview screen.

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

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

<p>Engine CPU Utilization</p>	<p>Displays aggregate CPU utilization for all engines on this server. For information on individual Adaptive Server engines, see the Engines screen. (Because all I/O for a process goes through one engine, CPU usage is not always evenly distributed across engines.) For cluster configurations, the graph shows the aggregate CPU utilization for each instance of the cluster.</p>
<p>Device IO/Sec</p>	<p>(For Adaptive Server cluster configurations, this information is on the Cluster Instances window) Displays device I/O per second aggregated across all devices on this server. For information on individual devices, see the Devices screen.</p>
<p>Memory</p>	<p>(For Adaptive Server cluster configurations, this information is on the Cluster Instances window) Displays memory usage statistics, including:</p> <ul style="list-style-type: none"> • The amounts of physical and logical memory in use • The amount of unused memory • The size of the procedure, statement, and data caches
<p>Processes</p>	<p>Displays process statistics, including:</p> <ul style="list-style-type: none"> • Max User Processes – number of processes this server is configured for • High Water Mark – highest number of processes that ran concurrently since this server started • Active – processes running now • (Adaptive Server cluster configurations only) Blocked Processes – processes that are waiting for a resource or for another process to finish <p>For more on processes, see the Processes screen.</p>
<p>Details tab</p>	<p>Displays information about this server, including number of days it has been running, number of deadlocks, data cache hit rate, procedure cache stalls, page and device sizes, maximum number of online engines, number of open databases, and the dates and times of the server's most recent restart and of the clearing of these counters. For cluster configurations, the Details tab also displays information about software product and version, hardware platform, server edition, number of deadlocks, page and device sizes, and when the server was last restarted, and the counters last cleared.</p>

Configured Resources tab	<p>Displays usage statistics for many of the configured resources for this server or cluster instance:</p> <ul style="list-style-type: none"> • Current – amount of this resource the server is using now • Run value – configured maximum for this resource • Percentage – percentage of the configured maximum represented by the current use of this resource • High Water Mark – maximum amount of this resource that has been used since this server started <p>Use the Percentage and High Water Mark columns to identify resources that might be over-configured or under-configured.</p> <p>For information on configured resources, see the Adaptive Server Enterprise documentation.</p>
Licenses tab	<p>Displays information about software licenses for Adaptive Server Enterprise on this server or cluster instance.</p>
Alerts tab	<p>Displays, for this server, cluster, or cluster instance, all alert notifications that have occurred since the Adaptive Server monitor was opened. If any alerts have occurred since you last looked at the Alerts tab, a yellow warning icon appears on the tab.</p> <p>You can control the number of alerts displayed using the Alert List Size property on the Settings screen.</p>

See also

- *Device Statistics and Details* on page 67
- *Engine Statistics and Details* on page 69
- *Process Statistics and Details* on page 74

Caches

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

Monitoring Adaptive Server Data Caches

Display information about data caches.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.
2. In the left pane, select **Caches**.
3. Click the **Data Caches** tab.

4. Select a cache in the **Data Caches** table.
The tabs at the bottom of the screen are populated with information about the selected cache.
5. (Optional for Adaptive Server cluster configurations) Select **Global** for information about global data caches. Select **Local** for information on data caches in each cluster instance.
6. (Optional for Adaptive Server cluster configurations) Select a cache in the **Global Data Caches** or **Local Data Caches** table.
The tabs at the bottom of the screen are populated with information about the selected cache.
7. Click the tabs to display **Pool Information** or **Cached Objects** details.
8. (Optional for Adaptive Server cluster configurations) Click the tab to display **Distribution** details for a selected global data cache.

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

See also

- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Monitoring Adaptive Server In-memory Storage* on page 61
- *Data Cache Statistics and Details* on page 57
- *Procedure Cache Statistics and Details* on page 59
- *Statement Cache Statistics and Details* on page 60
- *In-memory Storage Statistics and Details* on page 61

Modifying Data Cache Sizes

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

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

1. In the Adaptive Server monitor, select **Caches**.
2. From the Data Caches table, select the cache to configure.
3. Right-click the selected cache to display the **Resize** option that allows you to modify the size of the data cache.
You see the current size of the data cache, available space, and partition information.
4. Enter the new size of the data cache.
An increase in the data cache size is immediately effected in the server. A decrease in the data cache size requires the server to be re-started for the change to take effect.
5. (Optional) Enter a new value for the partitions in the data cache.

6. (Optional) Use **Calculate Overhead** to calculate the amount of memory required to resize the data cache with the new input size.
7. Click **Save**.
The dialog box closes if the operation succeeds, else Sybase Control Center displays an error and the dialog box stays open.

Adding Data Cache Buffer Pools

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

1. In the Adaptive Server monitor, select **Caches**.
2. From the Data Caches table, select the cache to configure.
3. Right-click the selected cache to display the **Add Buffer Pool** option that allows you to modify the configuration of the data cache. Alternately, click **Add Buffer Pool** in the Pool Information table.
4. To configure the new buffer pool, enter values for I/O buffer size, amount in pool, wash size, and a prefetch limit. You may also select the pool from which space will be allocated for the changed configuration.
5. Click **Save**.
The dialog box closes if the operation succeeds, else Sybase Control Center displays an error and the dialog box stays open.

Data Cache Statistics and Details

Interpret the Data Cache window for Adaptive Server.

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

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

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

Table 14. Tabs on the Data Cache screen

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

See also

- *Procedure Cache Statistics and Details* on page 59
- *Statement Cache Statistics and Details* on page 60
- *In-memory Storage Statistics and Details* on page 61
- *Monitoring Adaptive Server Data Caches* on page 55
- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Monitoring Adaptive Server In-memory Storage* on page 61

Monitoring the Adaptive Server Procedure Cache

Display information about the procedure cache.

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

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

See also

- *Monitoring Adaptive Server Data Caches* on page 55
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Monitoring Adaptive Server In-memory Storage* on page 61
- *Data Cache Statistics and Details* on page 57
- *Procedure Cache Statistics and Details* on page 59
- *Statement Cache Statistics and Details* on page 60
- *In-memory Storage Statistics and Details* on page 61

Procedure Cache Statistics and Details

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

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

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

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

See also

- *Data Cache Statistics and Details* on page 57
- *Statement Cache Statistics and Details* on page 60
- *In-memory Storage Statistics and Details* on page 61
- *Monitoring Adaptive Server Data Caches* on page 55
- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Monitoring Adaptive Server In-memory Storage* on page 61

Monitoring the Adaptive Server Statement Cache

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

1. In the Perspective Resources view, right-click the server and select **Monitor**.

2. In the left pane, select **Caches**.
3. Click the **Statement Cache** tab.

The Statement Cache tab is disabled if statement cache monitoring is not enabled. Statement cache monitoring is controlled by the **enable stmt cache monitoring** Adaptive Server configuration option.
4. (Optional for Adaptive Server cluster configurations) Select a specific cluster instance. The Cached Statements table shows all the statements for the selected instance.
5. Select a statement in the **Cached Statements** table.

The SQL query appears at the bottom of the screen.
6. Click the tabs to display pool information or details about cached objects.

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

See also

- *Monitoring Adaptive Server Data Caches* on page 55
- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring Adaptive Server In-memory Storage* on page 61
- *Data Cache Statistics and Details* on page 57
- *Procedure Cache Statistics and Details* on page 59
- *Statement Cache Statistics and Details* on page 60
- *In-memory Storage Statistics and Details* on page 61

Statement Cache Statistics and Details

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

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

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

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

See also

- *Data Cache Statistics and Details* on page 57
- *Procedure Cache Statistics and Details* on page 59
- *In-memory Storage Statistics and Details* on page 61
- *Monitoring Adaptive Server Data Caches* on page 55

- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Monitoring Adaptive Server In-memory Storage* on page 61

Monitoring Adaptive Server In-memory Storage

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

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

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.
2. In the left pane, select **Caches**.
3. Click the **In-memory storage** tab.
4. Select a cache in the **In-memory storage** table.
The tabs at the bottom of the screen are populated with information about the selected in-memory storage cache.
5. Click the tabs to display **In-memory Devices**, **In-memory Databases**, or **Cached Objects** details.

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

See also

- *Monitoring Adaptive Server Data Caches* on page 55
- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Data Cache Statistics and Details* on page 57
- *Procedure Cache Statistics and Details* on page 59
- *Statement Cache Statistics and Details* on page 60
- *In-memory Storage Statistics and Details* on page 61

In-memory Storage Statistics and Details

Interpreting the data on the In-memory Storage window.

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

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

Table 15. Tabs on the In-memory Storage window

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

See also

- *Data Cache Statistics and Details* on page 57
- *Procedure Cache Statistics and Details* on page 59
- *Statement Cache Statistics and Details* on page 60
- *Monitoring Adaptive Server Data Caches* on page 55
- *Monitoring the Adaptive Server Procedure Cache* on page 58
- *Monitoring the Adaptive Server Statement Cache* on page 59
- *Monitoring Adaptive Server In-memory Storage* on page 61

Databases

Monitor Adaptive Server databases.

Determining the Backup Status of a Database

Find out when a database's last backup started, whether the last backup failed, whether a backup is currently in progress, and more.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Databases**.
You can also display the Databases window by clicking a **Databases** link on another window in the Adaptive Server monitor.
3. Locate your database in the Databases table.

The table shows:

- The date and time at which the last backup started
- Whether a backup is currently in progress
- Whether the last backup failed
- Whether the transaction log is full

- Whether there are suspended processes associated with this database

For more information on backups, see the chapters on developing a recovery plan and on backing up and restoring user databases in the Adaptive Server *System Administration Guide*, Volume 2.

See also

- *Displaying Resources Used by a Database* on page 63
- *Displaying Information About Segments Used by a Database* on page 65
- *Database Statistics and Details* on page 64

Displaying Resources Used by a Database

View disk usage, running processes, unused indexes, and devices and segments associated with an Adaptive Server database.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Databases**.

You can also display the Databases screen by clicking a **Databases** link from a different window in the Adaptive Server monitor.

3. Select a database in the Databases table.

The tabs at the bottom of the screen are populated with information on the database you selected.

Note: When you select a database, Sybase Control Center calculates space usage before displaying any data, and this may take 30 seconds or more for a large database.

4. Click a tab to see information about the resource you are interested in:
 - **Details** – shows disk usage, including the size of reserved and unreserved data and log segments.
 - **Running Processes** – shows the server process ID (spid), login, host, command, and transaction for each process that is currently using the database. Each spid number is a link; click it to see more information about that process.
 - **Devices Used** – shows the size and usage allocation (data or log) for each device that provides storage for this database. Each device name is a link; click it to see more information about that device.
 - **Segments Used** – shows the size and unused portion of each segment, in megabytes. Each segment name is a link; click it to see more information about that segment.
 - **Unused Indexes** – shows the name and table of each unused index.
 - **Frequently Used Tables** – shows usage statistics for frequently used tables.

See also

- *Determining the Backup Status of a Database* on page 62
- *Displaying Information About Segments Used by a Database* on page 65

- *Database Statistics and Details* on page 64

Modifying Database Sizes

Sybase Control Center allows you to increase the size of your databases from the Database window.

1. In the Adaptive Server monitor, select **Databases**.
2. From the Databases table, select the database to configure.
3. Right-click a row and select **Extend**.
4. Select the **Extend** menu item from the context menu.
You see the **Extend Database Size** wizard .
5. (Optional) Select a **Device Name** on which to extend the database.
6. (Optional) Specify the amount of space you want to allocate to the log and data segments.
7. Click **OK**.
The dialog box closes if the operation succeeds, else Sybase Control Center displays an error and the dialog box stays open.

Database Statistics and Details

The Databases window shows a variety of detailed statistics, including the status, of active Adaptive Server databases.

Note: In Adaptive Server cluster configurations, the Global Databases table provides information about global databases, and the Local Databases tables provides information about local, temporary databases, grouped by cluster instance.

The Databases table lists the databases in the current Adaptive Server by name. If a database is unavailable, for example, because it is quiesced or is offline, the Name column includes the reason.

For Adaptive Server 15.5 and later, the Databases table lists the type, durability, and DML logging status for each database. The Databases table also includes for each database, the ID, and current status information, backup status, whether the transaction log is full, and whether there are suspended processes. Processes may be suspended when the transaction log fills up.

The type of database is indicated for temporary, in-memory, proxy and archive databases, and left blank for all other databases. The Durability column indicates if a database is recoverable.

The tabs at the bottom of the screen display information about the selected database. For a large database, it might take 30 seconds or more for the information to appear.

Table 16. Tabs in the Databases View

Details	Displays information about space usage, with pie charts for data segments and log segments. If this database does not have a log segment, the pie chart on the right shows combined data and log segment usage.
---------	---

Running Processes	<p>Displays information about processes that are currently using this database, including the process ID, login, host, command, and transaction name.</p> <p>Click a process ID in the SPID column to switch to the Processes view's information about that process.</p>
Devices Used	<p>Displays information about devices on which this database stores its data, including the device name, the amount of space used on that device, and the usage allocation (data or log).</p> <p>Click a device in the Name column to switch to the Devices view's information about that device.</p>
Segments Used	<p>Displays information about segments used by this database, including the segment name, the size of the segment in megabytes, and the amount of free space in the segment.</p> <p>Click a segment in the Name column to switch to the Segments view's information about that segment.</p>
Unused Indexes	<p>Lists indexes in this database that have not been used since the Adaptive Server was last restarted.</p>
Frequently Used Tables	<p>Displays information about tables in this database that have been used since the Adaptive Server was last restarted, including the table name, index ID, logical and physical reads, lock requests and waits, and contention statistics.</p>

The Local Temporary Databases screen provides information about local, temporary databases grouped by cluster instances. For each cluster instance, Sybase Control Center lists the ID and current status information, including backup status; whether the transaction log is full, and whether there are suspended processes.

The tabs at the bottom of the Local Temporary Databases screen display information about the selected database. Click the Temporary Database Activity tab to see information about log requests and wait status, device read and write values, and charts depicting I/O activity in the selected temporary database.

See also

- *Determining the Backup Status of a Database* on page 62
- *Displaying Resources Used by a Database* on page 63
- *Displaying Information About Segments Used by a Database* on page 65

Displaying Information About Segments Used by a Database

For a database, get details about space usage, devices that make up each segment, and tables and indexes that use it.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Databases**.
3. Select a database in the Databases table at the top of the window.
The tabs at the bottom of the window are populated with information on the database you selected.
4. Click the **Segments Used** tab.
5. Each segment name is a link; click it to see more information about that segment.
The Segments window appears. In the Segments window, the tabs at the bottom of the screen are populated with information about the selected segment.
6. Click the tabs to see information about space usage on the segment, devices that make up the segment, and tables and indexes that are allocated on the segment.
7. If the database uses more than one segment, return to the Databases screen to identify and click through to the remaining segments.

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

See also

- *Determining the Backup Status of a Database* on page 62
- *Displaying Resources Used by a Database* on page 63
- *Database Statistics and Details* on page 64
- *Determining the Space Used by a Table on a Segment* on page 76
- *Extending a Segment in Adaptive Server* on page 76
- *Segment Statistics and Details* on page 77

Devices

Monitor the devices used by Adaptive Server.

Determining Device I/O Response and I/O per Second

Find out how long a device is taking to respond to I/O requests and what its I/O rate is.

High response time can indicate problems in the functioning of the physical device or the storage layer, problems with the configuration of the storage layer, or that the device is busy.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.
2. In the left pane, select **Devices** .
You can also display the Devices screen by clicking a **Devices** link on another window in the Adaptive Server monitor.
3. (For shared-disk clusters only) Select **Global** to display global devices, and **Local** to display devices for instances of Adaptive Server clusters.

4. In the Devices table, select the device to monitor.
5. The **IO Response Time** column and the **Device IO/Sec** graph on the Details tab display details on the I/O activity on the selected device..

See also

- *Modifying Device Sizes* on page 67
- *Device Statistics and Details* on page 67

Modifying Device Sizes

Sybase Control Center allows you to increase the size of your devices from the Devices window.

1. In the Adaptive Server Monitor, select **Devices**.
Sybase Control Center displays the Devices table.
2. Select the device to configure.
3. Right-click and select **Resize**.
Sybase Control Center displays the device resize dialog with the name of the selected device, allocated size, input field for increased size, unit of size, and an option that allows you to specify whether to initialize the device.
4. Input the amount by which to increase the device size.
The dialog box now displays the new device size that is calculated based on the input. If there is an error, it is indicated in the dialog box.
5. Click **OK**.
For more information on devices, see the Adaptive Server *System Administration Guide*, Volume 2.

See also

- *Determining Device I/O Response and I/O per Second* on page 66
- *Device Statistics and Details* on page 67

Device Statistics and Details

Interpret the Devices screen for Adaptive Server.

Note: For Adaptive Server Cluster configurations, the **Devices** table is called **Global Devices** when you select the **Global** tab, and **Private Database Device** when you select the **Local** tab. Information in the Private Database Device table is grouped by cluster instance.

The Devices table displays information about all devices that store databases for this Adaptive Server. A device can be a whole disk drive, or any part of a disk or file system. The charts are populated by data from collection_ase_all_client_kpis, covering the current trend period.

The **Devices** table includes device semaphore statistics. The device semaphore controls access to device I/O; a high ratio of **Device Semaphore Waits** to **Device Semaphore Requests**

indicates contention. If **IO Wait Time** is high enough to cause concern, you may want to redistribute the data on the physical devices.

The tabs at the bottom of the screen show information about the device selected in the **Devices** table.

Table 17. Tabs on the Devices Screen

<p>Details</p>	<p>Displays two charts:</p> <ul style="list-style-type: none"> • A pie chart showing space usage on the selected device. Includes used and unused space, in megabytes, and as percentages of all the available space on the device. The title above the chart indicates the total available space. • Device IO/Sec – a line graph showing the rate of I/O per second on the selected device over the current trend period. The graph shows the sum of reads, writes, and asynchronous prefetch (APF) reads. Because the graph shows a rate, and the read, write, and APF read figures in the table are changes since the last refresh, the values do not correspond. <hr/> <p>Note: Sybase Control Center does not display the Device IO/Sec graph for in-memory devices.</p>
<p>Advanced</p>	<p>Displays two charts:</p> <ul style="list-style-type: none"> • Device IO Response Time – a line graph showing the response time, in milliseconds, for I/O operations performed on the selected device. • Device APF Reads/Sec – a line graph showing the rate of asynchronous prefetch read operations, per second, on the selected device. APF reads indicate that table scans are taking place. <hr/> <p>Note: Sybase Control Center does not display the Advanced tab for in-memory devices.</p>
<p>IO Distribution</p>	<p>Appears only for Global Devices on Adaptive Server shared-disk clusters. Details for master device I/O activities for each instance include I/O wait time, response time, APF reads and request and wait times for device semaphores.</p>

See also

- *Engine Statistics and Details* on page 69
- *Process Statistics and Details* on page 74
- *Determining Device I/O Response and I/O per Second* on page 66
- *Modifying Device Sizes* on page 67
- *Setting Up Statistics Collection* on page 12

Engines

Monitor Adaptive Server engines.

Displaying Engine CPU Utilization

Find out how heavily Adaptive Server engines are being used.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Engines**.
3. In the Engines table, select the engine you want to monitor.
The I/O Processing and Garbage Collection tables and the Engine CPU Utilization graph at the bottom of the screen are populated with data for the selected engine.

See also

- *Engine Statistics and Details* on page 69

Engine Statistics and Details

Interpret the Engines screen for Adaptive Server.

The Engines screen displays information about all processing engines for this Adaptive Server. The charts are populated by data from the collection_ase_all_client_kpis, covering the current trend period.

The Engines table identifies engines by number, and gives CPU utilization percentages, status, start date and time, number of connections, and the operating system process identifier (**OS PID**) for each one.

Note: For Adaptive Server shared-disk clusters, the Engines table has information that is grouped by cluster instances.

The area at the bottom of the screen shows information about the selected engine selected.

IO Processing table	Provides counts of disk I/O checks, checks without waits, polls, and completed operations over the current trend period.
Garbage Collection table	For the current trend period, provides the garbage collector's maximum queue size and counts of pending items, high water mark items, and overflows.
Engine CPU Utilization graph	A line graph showing CPU utilization for this engine as a percentage. If Adaptive Server is performing poorly, use the information from this graph to determine how busy the engines are.

See also

- *Displaying Engine CPU Utilization* on page 69

Processes

Monitor Adaptive Server processes.

Identifying Resource-Intensive Processes

Find the user processes that are consuming the most system resources on the selected Adaptive Server.

You can choose a system resource (CPU, disk I/O, incoming network traffic, or outgoing network traffic) and display information about the user processes on the selected server that are using the chosen resource most intensively. For each system resource, you can rank the processes by cumulative or most recent activity values. Each bar in the graph represents the value of the selected metric for a process.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Processes**.
The **All Processes** tab is selected. A bar graph shows the five user processes that are using the most cumulative CPU cycles.
3. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
4. Use the menu to the right of the bar graph to change the system resource. You can choose:
 - **CPU Cumulative** – cumulative CPU activity since the Adaptive Server started or the counter wrapped.
 - **CPU Activity** – CPU activity, per second, since the last screen refresh.
 - **Disk I/O Cumulative** – cumulative disk I/O since the Adaptive Server started or the counter wrapped.
 - **Disk I/O Activity** – I/O activity per second since the last screen refresh.
 - **Incoming Network Traffic Cumulative** – cumulative incoming network traffic since the Adaptive Server started or the counter wrapped.
 - **Incoming Network Traffic Activity** – incoming network traffic per second since the last screen refresh.
 - **Outgoing Network Traffic Cumulative** – cumulative outgoing network traffic since the Adaptive Server started or the counter wrapped.
 - **Outgoing Network Traffic Activity** – outgoing network traffic per second since the last screen refresh.
5. (Optional) Move your mouse over a bar in the graph to display the server process ID (SPID) and the value of the selected system resource metric for the process.
6. (Optional) Click a bar in the graph to highlight information about that process in the table below.

- (Optional) Click a bar in the graph or a row in the table to display information about that process in the Details, SQL, and Wait Events tabs at the bottom of the screen.

See also

- *Displaying Wait Events for a Process* on page 74
- *Displaying the SQL Query Associated with a Process* on page 73
- *Identifying the Lead Blocker in a Chain* on page 72
- *Identifying Blocked Processes and Blocking Processes* on page 71
- *Terminating Blocking Processes* on page 72
- *Process Statistics and Details* on page 74

Identifying Blocked Processes and Blocking Processes

Find user processes that are blocked and the processes that are blocking them.

- In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
- In the left pane, select **Processes**.
- Click the **All Processes** tab.
- (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
- Check the table below the bar graph for rows highlighted in red, and with a "lock" icon, which indicate blocked processes.
The Blocked by SPID column identifies the blocking process. Blocking processes are also shown in the table, highlighted in yellow.
- Click a red table row to display information about the blocked process in the the Details, SQL, and Wait Events tabs at the bottom of the screen.
- Click the row for the blocking process (yellow) to display its information in the tabs.

Note: Identifying the lock held by a blocking process is not always straightforward. For example, the blocking process does not necessarily hold a page lock; it might hold a table lock. For this reason, Sybase Control Center shows the lock request process that is blocking another process, not the blocking lock.

- Click the **Blocked Processes** tab to display additional information about blocked processes, including details about the lock, the row number, and the page number.

Next

For information on handling blocked processes, see the locking reports chapter of the *Adaptive Server Performance and Tuning Series: Locking and Concurrency Control*.

See also

- *Displaying Wait Events for a Process* on page 74

- *Displaying the SQL Query Associated with a Process* on page 73
- *Identifying the Lead Blocker in a Chain* on page 72
- *Identifying Resource-Intensive Processes* on page 70
- *Terminating Blocking Processes* on page 72
- *Process Statistics and Details* on page 74

Terminating Blocking Processes

Terminate a blocking process from the Processes window.

Sybase Control Center allows you to terminate a blocking process.

1. In the Adaptive Server monitor, select **Processes**.
2. Select the blocking process, or a set of blocking processes, that you wish to terminate.
3. Right-click the selected row and select **Terminate**.
4. Select **Yes** to relay the terminate request to the Adaptive Server, or **No** to close the dialog box without performing the terminate operation.

See also

- *Displaying Wait Events for a Process* on page 74
- *Displaying the SQL Query Associated with a Process* on page 73
- *Identifying the Lead Blocker in a Chain* on page 72
- *Identifying Resource-Intensive Processes* on page 70
- *Identifying Blocked Processes and Blocking Processes* on page 71
- *Process Statistics and Details* on page 74

Identifying the Lead Blocker in a Chain

Find a process that is blocking several other processes.

When Process A blocks Process B, which blocks Process C—and so on—the blocking processes form a chain.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Processes**.
3. Click on **Blocked Processes**.

You see a table with information about blocked and blocking processes, including the lock requests on the basis of which processes are blocked.

4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
5. The table on the Blocked Processes tab has an entry for each lead blocker; click the arrow to expand the entry and show all the blocked processes in the chain.

Note: Identifying the lock held by a blocking process is not always straightforward. For example, the blocking process does not necessarily hold a page lock; it might hold a table

lock. For this reason, Sybase Control Center shows the lock request process that is blocking another process, not the blocking lock.

6. Select a process to populate the tabs at the bottom of the screen with information about that process.
7. Click the tabs to see:

Options	Description
Details	Details about the selected process, including the initiating program, transaction information, and network statistics
SQL	The SQL statement and query plan for the selected process
Wait Events	Information about wait events for the selected process, including number of waits, wait times, and wait descriptions

8. (Optional) Click the **All Processes** tab at the top of the window.
In the table below the bar graph, rows highlighted in yellow are blocking processes. Rows highlighted in red are blocked processes.

Next

For information on handling blocked processes, see the *Adaptive Server Performance and Tuning Series: Locking and Concurrency Control*. You can find Adaptive Server documentation on the Sybase Product Manuals Web site at <http://sybooks.sybase.com>.

See also

- *Displaying Wait Events for a Process* on page 74
- *Displaying the SQL Query Associated with a Process* on page 73
- *Identifying Resource-Intensive Processes* on page 70
- *Identifying Blocked Processes and Blocking Processes* on page 71
- *Terminating Blocking Processes* on page 72
- *Process Statistics and Details* on page 74

Displaying the SQL Query Associated with a Process

See the SQL statement and query plan for an Adaptive Server user process.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.
2. In the left pane, select **Processes**.
3. Click either the **All Processes** tab or the **Blocked Processes** tab.
4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**. Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.

5. To select a process, click a row in the table.
6. At the bottom of the window, click the **SQL** tab.
The tab shows the SQL query and query plan for this process.

See also

- *Displaying Wait Events for a Process* on page 74
- *Identifying the Lead Blocker in a Chain* on page 72
- *Identifying Resource-Intensive Processes* on page 70
- *Identifying Blocked Processes and Blocking Processes* on page 71
- *Terminating Blocking Processes* on page 72
- *Process Statistics and Details* on page 74

Displaying Wait Events for a Process

Get information about wait events that are affecting an Adaptive Server process.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Processes**.
3. Click **All Processes**.
Each row of the table below the bar graph displays information about a process.
4. (Optional for Adaptive Server cluster configurations) Click **Select User Processes by**.
Choose **All Instances** to depict information for the entire cluster, or select a specific instance to see information for only the selected instance of the cluster.
5. To select a process, click a row in the table.
6. Click **Wait Events**.

For more information on wait events, see the wait events chapter in the Adaptive Server *Performance and Tuning Series: Monitoring Tables* guide.

See also

- *Displaying the SQL Query Associated with a Process* on page 73
- *Identifying the Lead Blocker in a Chain* on page 72
- *Identifying Resource-Intensive Processes* on page 70
- *Identifying Blocked Processes and Blocking Processes* on page 71
- *Terminating Blocking Processes* on page 72
- *Process Statistics and Details* on page 74

Process Statistics and Details

Interpret Adaptive Server process information.

Lock icons in the SPID column of the Processes table identify processes that are blocked (a grayed-out lock) or blocking (a gold lock) other processes. Other columns of the Processes

table include the family ID which is the parent SPID value, processes blocked by SPID, CPU activity, CPU cumulative activity, disk I/O activity and disk I/O cumulative activity.

On the Blocked Processes tab, Sybase Control Center shows the lock request process that is blocking another process, not the blocking lock itself. A yellow warning icon appears on the Blocked Processes tab label when there are any blocked processes.

Table 18. Color indicators in the Processes table

Color	Indicates that the process is...
Blue	Selected
Green	Executing a query
Yellow	Blocking another process
Red	Blocked by a lock held by another process

Table 19. Tabs

Details	Displays information about the selected process, including program name, current SQL command, client machine IP address and name, transaction name and start time, physical and logical reads, and CPU utilization.
SQL	If the selected process is active, displays the SQL statement and query plan for the query that the process is executing.
Wait Events	<p>Displays a list of all events the selected process has waited for. Wait events are internal Adaptive Server states that represent conditions that cause a process to stop. Common wait events include waiting:</p> <ul style="list-style-type: none"> • On the Adaptive Server scheduler runnable queue for a CPU to become available • For disk I/O to complete • For a lock on a table to be released

See also

- *Displaying Wait Events for a Process* on page 74
- *Displaying the SQL Query Associated with a Process* on page 73
- *Identifying the Lead Blocker in a Chain* on page 72
- *Identifying Resource-Intensive Processes* on page 70
- *Identifying Blocked Processes and Blocking Processes* on page 71
- *Terminating Blocking Processes* on page 72

Segments

Monitor the segments used by Adaptive Server databases.

Determining the Space Used by a Table on a Segment

Find reserved space figures for tables on a segment.

You can sort tables by their reserved sizes, which simplifies planning for a reorganization or rebuild.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.
2. In the left pane, select **Segments**.
You can also display the Segments window by clicking a **Segments** link on another window in the Adaptive Server monitor.
3. (Optional for Adaptive Server cluster configurations) Select **Global** for information about segments on the global database. Select **Local** for information about segments on local, temporary databases.
4. In the Segments table, select the segment to monitor.
5. Click the **Used Tables** tab.
The tab displays the name and amount of space reserved, in kilobytes, for each table.
6. To sort a table by reserved size, select the table and click **Space Reserved**.

For more information on segments, see the chapter on creating and using segments in the *Adaptive Server System Administration Guide, Volume 2*.

See also

- *Extending a Segment in Adaptive Server* on page 76
- *Displaying Information About Segments Used by a Database* on page 65
- *Segment Statistics and Details* on page 77

Extending a Segment in Adaptive Server

Sybase Control Center allows you to extend a segment on a specific device.

1. In the Adaptive Server monitor, select **Segments**.
You see the Segments table.
2. Select the segment to configure.
3. Right-click the selected row and select **Extend**.
4. Select the **Extend** menu item from the context menu.
You see the **Extend** dialog, which includes the name, size, and unused size of the device.
5. Click the device name onto which to extend the segment.

Note: The device list is empty if the selected segment is using all the devices configured on the server. In this case, add a new device to the server in order to extend the segment.

6. Click **OK**.

See also

- *Determining the Space Used by a Table on a Segment* on page 76
- *Displaying Information About Segments Used by a Database* on page 65
- *Segment Statistics and Details* on page 77

Segment Statistics and Details

Interpret the Segment information for Adaptive Server.

The Segment Statistics and Details screen displays information about Adaptive Server segments. In Adaptive Server cluster configurations, selecting the **Global** tab displays information about segments on global databases. Selecting the **Local** tab displays information about segments on local, temporary databases, grouped by cluster instances.

The Segments screen displays information about all segments for this Adaptive Server. The charts on this screen are populated by data from the collection_ase_all_client_kpis, covering the current trend period.

The Segments table lists each segment used by this Adaptive Server and gives the name of the database that uses the segment, the database's size and unused space on the segment, and the number of thresholds.

The tabs at the bottom of the screen show information about the segment selected in the Segments table.

Table 20. Tabs on the Segments Screen

Details	<p>Displays two charts:</p> <ul style="list-style-type: none"> • A pie chart showing current space usage on the selected segment. Includes used and unused space, in megabytes, and as percentages of the available space on the segment. The title above the chart indicates the total available space. • Space Usage – a line graph showing changes in space usage on the selected segment over the current trend period.
Devices Used	<p>Displays devices included in the selected segment and the size of each device, in megabytes. Click a name in the Device column to switch to the Devices monitoring view's information for that device.</p>
Used Tables	<p>Displays tables allocated on the selected segment and the reserved size of each table, in kilobytes.</p>

Used Indexes	Displays indexes allocated on the selected segment and the table associated with each index.
--------------	--

For more information on segments, see the chapter on creating and using segments in the Adaptive Server *System Administration Guide*, Volume 2.

See also

- *Determining the Space Used by a Table on a Segment* on page 76
- *Extending a Segment in Adaptive Server* on page 76
- *Displaying Information About Segments Used by a Database* on page 65
- *Setting Up Statistics Collection* on page 12

Transactions

Monitor active Adaptive Server transactions.

Identifying a Transaction's Process

Get information about a currently running transaction, including the process that initiated the transaction.

1. In the Perspective Resources view, right-click the server you want to monitor and select **Monitor**.
2. In the left pane, select **Transactions**.
You can also display the Transactions screen by clicking a SPID link on Running Processes tab of the Databases screen.
3. Locate the transaction in the Transactions table.
4. In the SPID column, click the SPID of the process associated with your transaction. (The ID number is a link.)
The Processes screen appears, displaying information about your transaction's parent process.

See also

- *Process Statistics and Details* on page 74
- *Transaction Statistics and Details* on page 78

Transaction Statistics and Details

Interpret information about Adaptive Server transactions.

The Transactions table displays information about all active transactions on the selected Adaptive Server. Details include the name of the transaction, the login of the user who owns it, the application that launched the transaction, the process that initiated the transaction (SPID column), the transaction's start time, the name of the host where the transaction is running, and the database it is running against. (If a transaction affects more than one database, the table

does not display them all—it shows the transaction’s current database and the process that started the transaction.)

Note: In Adaptive Server cluster configurations, information in the Transactions table is grouped by cluster instances.

When you select a transaction in the table, the User Log Cache Usage tab at the bottom of the screen displays statistics about the user log cache for the selected transaction. Details include bytes written, number of flushes and full flushes, maximum cache usage (in bytes), and current usage (in bytes).

See also

- *Process Statistics and Details* on page 74
- *Identifying a Transaction’s Process* on page 78

Server Configuration Values

Monitor Adaptive Server configuration values.

Displaying Configuration Values

Display the values and descriptions of parameters in the Adaptive Server configuration file.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Server Configuration**.
3. (Optional - for Adaptive Server cluster configurations only) Select a cluster instance from **Show Configuration for Cluster Instance**. Select **All** for information on all cluster instances.
4. Select a configuration category from **Show Configuration Parameters for**. For example, select **Cache Manager**.
5. (Optional) Enter filter text in **Show Configuration Parameter Matching**. For example, enter `sql`.

The Server Configuration table displays parameters whose names match or include your filter text. The filter narrows the selection made with **Show Configuration Parameters for**, so if you select **All** and then filter on `sql`, you find more parameters than when you select **Monitoring** and then filter on `sql`.

6. From the Server Configuration table, select a parameter.
A description of the parameter appears at the bottom of the screen.

For more information on configuration parameters, see the chapter on setting configuration parameters in the Adaptive Server *System Administration Guide*, Volume 1.

See also

- *Modifying Server Configuration Parameters* on page 80

- *Server Configuration Statistics and Details* on page 80

Modifying Server Configuration Parameters

Use Sybase Control Center to configure server configuration parameters.

1. In the Adaptive Server monitor, select **Server Configuration**.
You see a table with fields parameter names, value, default value, maximum value, minimum value, and restart required. Editable columns are indicated by a "pencil" icon. To reset a value to the previously configured value, click the "Reset" icon that appears after you have edited a field.
2. Select the server configuration parameter to configure. For example, increase or decrease the size of the procedure cache by selecting the **procedure cache size** parameter, or of the statement cache by selecting **statement cache size** parameter.
 - a) Click **Configure Value** for the selected row.
 - b) Enter the new value for the configuration parameter.
If the value is invalid, you see an error message. As some parameters require a restart, the changed value is in the **Pending Value** column till the server is restarted.
 - c) Click **Save All** to update the server with the new values or **Reset All** to restore the original values for the resource.

See also

- *Displaying Configuration Values* on page 79
- *Server Configuration Statistics and Details* on page 80

Server Configuration Statistics and Details

Interpret Adaptive Server server configuration information.

The Server Configuration screen displays information about all configuration parameters for this Adaptive Server. You can modify the configuration parameters using either the **sp_configure** stored procedure, or by editing the configuration file. The parameters are stored in the Adaptive Server configuration file on each server.

You can also use the Sybase Control Center to change the current values of the server configuration parameters and change the size of the procedure cache and statement cache by editing the **Value** fields corresponding to the specific row.

For each parameter, the Server Configuration table gives the name; displays the current, default, minimum, and maximum values; and indicates whether you have to restart the server to make a change to this parameter take effect.

See also

- *Modifying Server Configuration Parameters* on page 80
- *Displaying Configuration Values* on page 79

SQL Activity

Monitor SQL queries on Adaptive Server.

Monitoring SQL Queries

Display details about recently executed SQL queries.

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **SQL Activity**.

The SQL Activity screen appears. It lists queries executed during the current trend period, along with details including each query's server process identifier (spid), the login account that executed the query, the kernel process identifier (KPID), batch identifier, and execution statistics.

Note: In Adaptive Server cluster configurations, information in the Active SQLs table is grouped by cluster instances.

3. Select a SQL query from the Active SQLs in Batch table.
The SQL statement appears at the bottom of the screen.

Note: The SQL Activity screen displays SQL text for the most recent collection interval only.

Replication Agents

Monitor RepAgent threads for replicate databases on the selected Adaptive Server.

In Adaptive Server, there is one RepAgent thread for each database from which data is replicated. The RepAgent thread reads the transaction log of a primary database. It sends the transaction as Log Transfer Language (LTL) commands for replicated tables and replicated stored procedures to the primary Replication Server, which converts the LTL commands into SQL and applies the SQL to the replicate database.

For more information about replication, see the *ASE Replicator Users Guide*.

Monitoring RepAgent Threads

Display the status and transaction log details of the RepAgent threads running in Adaptive Server.

Prerequisites

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

1. In the Perspective Resources view, right-click the server to monitor and select **Monitor**.
2. In the left pane, select **Replication Agent**.

Note: If Sybase Control Center does not detect a Replication Management Agent Plug-in, or detects one that is incompatible with your Adaptive Server version, the **Replication Agent** option is greyed out.

3. Select a RepAgent thread in the table.
You see the **Log Size** and **Activity** graphs.

For more information about replication, see the *ASE Replicator Users Guide*.

See also

- *Setting Replication Parameters* on page 82
- *Replication Agent Statistics and Details* on page 82

Setting Replication Parameters

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

Replication Agent Statistics and Details

Interpret the Replication Agent information for Adaptive Server.

The Replication Agent screen displays the name, status, and controlling Replication Server of the RepAgent thread for the current Adaptive Server. The Transaction Log Details tab displays the log data of the selected RepAgent thread.

Transaction Log Details tab

Log Size	Displays a graph of the log size, in megabytes.
Activity	Displays a log count activity graph of the of the scanned and processed log records.

For more information about replication, see the *ASE Replicator Users Guide*.

See also

- *Monitoring RepAgent Threads* on page 81
- *Setting Replication Parameters* on page 82

Settings

Learn about Adaptive Server monitoring controls on the Settings screen.

Table 21. Controls on the Settings screen

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

Statistics

Availability and performance statistics in Sybase Control Center for Adaptive Server can help you identify if your system is running as efficiently as possible.

Availability statistics are concerned with present conditions; they help you determine whether a system component you are monitoring (a server or a storage device, for example) is running and functioning properly. Performance statistics are concerned with behavior of the same components over time. You can use them to spot trends, identify problems or potential problems, and make plans.

Sybase Control Center includes predefined key performance indicators (KPIs) for Adaptive Server that are grouped into collections. KPIs such as Server Status, which serves as an availability statistic when it is fresh, also has long-term value as historical performance statistics.

Availability statistics appear on the heat chart and on the screens of the Adaptive Server monitor for each Adaptive Server. The Heat Chart includes these KPIs:

- Server Status (up or down)
- Server CPU Utilization
- Number of Blocked Processes
- Number of Suspended Processes

These KPIs are part of the default collection: `collection_ase_availability`.

Performance statistics appear on the Statistics Chart and on the screens of the Adaptive Server monitor for each Adaptive Server. In the Statistics Chart, you expand folders in the Statistics tab to drill down to specific KPIs. See the KPI values as tables or graphs, and compare them by displaying several KPIs together. These are the folders that contain KPIs you can use to display data:

- Server Overview
- Devices
- Engines
- Segments
- DataCaches

Adaptive Server cluster configurations have these additional folders on the Statistics Chart:

- Cluster Instances
- Cluster Workload
- Logical Clusters
- tempDBs Activity

To make specific KPIs available to the Statistics Chart and to the Adaptive Server monitor screens that use them, in addition to the availability statistics scheduled by default, set up collection jobs in the scheduler for:

- `collection_ase_histmon` – The KPIs in this collection are only available in the Statistics Chart, and for setting alerts.
- `collection_ase_all_client_kpis` – This collection is necessary to gather statistics for historical charts in the Adaptive Server Monitor. The statistics are also available in the Statistic Chart, and for setting alerts.

Several configuration options affect the collection and appearance of Adaptive Server data in Sybase Control Center:

- Collection repeat interval – the frequency of data collection. Set the repeat interval on the collection job in the scheduler. This option is set when the collection is scheduled, but may be modified later.

- Screen refresh interval – the period between screen refreshes. Refreshing the screen redraws it with the latest available data. This option can be configured from the Settings window.
- Chart trend period – the period over which data appears in historical charts. This option can be configured from the Settings window.

Interpreting Statistics

Understand the scope and freshness of Adaptive Server data in Sybase Control Center.

Each Adaptive Server statistic presented in Sybase Control Center has a scope:

- Delta – the number of occurrences since the last screen refresh. For example, the user log cache statistics on the Transactions screen are delta values.
- Rate – the number of occurrences over the given period of time. Device I/O is given as a rate.
- Percentage or ratio – an amount, number, or rate stated as a proportion to a whole. Percentage statistics include CPU utilization, space usage on devices and segments, and cache hit rates. Ratios include cache volatility.
- Count – a simple value; for example, the size of a database or cache in megabytes, or the number of partitions in a cache.
- Cumulative – the number of occurrences since Adaptive Server started, or since the counter wrapped. On the Processes screen, you can set the Top 5 User Processes chart to display the five processes that use the most CPU, disk I/O, or network resources as rates or as cumulative values.

Note: When a server has been running for a long time, its statistical counters can wrap, which means they restart from zero. This most often affects cumulative statistics. Information about when or how many times a counter has wrapped is not available.

Most Adaptive Server statistics not otherwise labeled are presented as deltas since the last screen refresh; rates, percentages, and cumulative numbers are labeled as such.

Sybase Control Center displays statistics promptly. However, there are several factors that affect the freshness of the data on the screens:

- The screen refresh interval, which you can set on the Settings screen
- The collection repeat interval, which you can set in the scheduler for each Adaptive Server collection
- Network latency

For more information about interpreting Adaptive Server statistics, see:

- *Adaptive Server Performance and Tuning Series: Improving Performance with Statistical Analysis*
- *Adaptive Server Performance and Tuning Series: Monitoring Adaptive Server with sp_sysmon*

Manage and Monitor

- *Adaptive Server Performance and Tuning Series: Basics*

Troubleshoot Sybase Control Center for Adaptive Server

Solve problems that occur in Sybase Control Center for Adaptive Server.

Error: Unable to Format the Date String

Problem: While using Adaptive Server 15.5, the error log contains error messages about date string format.

If you use Adaptive Server 15.5, you see error messages in the Sybase Control Center server log such as this:

```
2009-11-17 09:13:14,493 ERROR [RMI TCP
Connection(12)-10.33.55.77] Unable to format the date string
00000 08:39:50.23 using the format yyyy/MM/dd HH:mm:ss.SSS
```

Solution: Follow this procedure:

1. Shut down the Unified Agent managing the Adaptive Server.
2. Rename `UAF-2_5/plugins/com.sybase.ase/lib/ASEAgentPlugin.jar`.
3. Copy `ASEAgentPlugin.jar` from `SCC-3_1/plugins/ASEMAP/ASEAgentPlugin.jar` to `UAF-2_5/plugins/com.sybase.ase/lib/ASEAgentPlugin.jar`.
4. Start the Unified Agent that is managing the Adaptive Server.

KPI is Not Updated

Problem: Cannot update KPI: **Number of Transactions**

Solution: The **Number of Transactions** KPI is populated only if the monitored Adaptive Server version is 15.0.3 or later.

If your server version is 15.0.3 or later and the **Number of Transactions** KPI is not being updated, verify that you have installed the latest version of the Adaptive Server `installmaster` script on your server.

Invalid Connection Profile

Problem: During the login authentication step, a security error warns that the connection profile for this Adaptive Server is invalid.

Solution: Check the connection information stored by Sybase Control Center for this Adaptive Server. You can make sure the information is valid by using it with iSQL or Sybase Central to connect to the server.

See also

- *Registering an Adaptive Server* on page 7

Cannot Monitor Adaptive Server or Display Statistics Chart

Problem: In the Perspective Resources view, the Monitor and Statistics Chart context menu items for a monitored Adaptive Server are grayed out.

Solution: Make sure your user account is authenticated on the Adaptive Server. To monitor the Adaptive Server, you must also make sure that your account has been granted **mon_role** on the Adaptive Server.

See also

- *Role Assignment in Sybase Control Center for Adaptive Server* on page 11

Data on Screens or Charts Is Missing

Problem: Data on screens or charts is missing.

Solution 1: Schedule collections

Schedule collection jobs for Adaptive Server statistics collections. If collections are scheduled, make sure the start and stop times are set correctly.

- In the Adaptive Server monitor, the Overview, Devices, Engines, and Segments screens use data from collection_ase_all_client_kpis.
- Also in the Adaptive Server monitor, the Replication Agent screen uses data from collection_ase_rat.
- The Statistics Chart uses data from collection_ase_histmon and collection_ase_all_client_kpis.
- The Heat Chart uses data from collection_ase_availability, which is the default collection.

Solution 2: Check revalidation frequency

If appropriate collections are scheduled but data is missing on caches, devices, engines, or segments that were recently added to a registered Adaptive Server, check the value of the `revalidation_frequency` parameter. If the repository has not refreshed (revalidated) its list of monitored resources since the new resources were added, data on the new resources is not collected or displayed. Wait for the next revalidation to see the data.

Solution 3: Reset system clocks

If appropriate collections are scheduled but data is missing, truncated, or incomplete, compare the clock settings on the machines where Sybase Control Center and the client browser are running. Clocks that are out of sync by a few minutes cause screens and charts to display incomplete data. If the clocks' time difference exceeds the value of the chart trend period, displays that use the chart trend period contain no data. (This problem occurs when the time on one or both system clocks is incorrect. It is not caused by time zone differences—servers and clients can operate successfully in different time zones.)

Solution 4: Reset collection interval or trend period

If appropriate collections are scheduled but lines on graphs are missing, truncated, or incomplete, the values of the collection repeat interval and the chart trend period might be too close together. When these options are set to similar values, graphs sometimes contain only a single data point. Because Sybase Control Center needs at least two data points to draw a curve on a graph and a single data point is not displayed, graphs are empty. For example, if the collection interval is 12 minutes and the trend period is 15 minutes, graphs will display two data points for only a few minutes at a time, so they will appear to be blank more often than not.

To resolve the problem, decrease the collection repeat interval (set on the collection job in the scheduler) or increase the chart trend period (set on the Settings screen in the Adaptive Server monitor) so that multiple collection intervals occur within the trend period.

See also

- *Setting Adaptive Server Parameters in the Configuration File* on page 21
- *Setting Up Statistics Collection* on page 12

Adaptive Server Is Responding Slowly

Problem: A monitored Adaptive Server is responding slowly. How do you tell whether the problem lies in the network or the server?

Solution: On the Adaptive Server monitor for the Adaptive Server in question, select **Engines**. On the Engines screen, select an engine from the Engines table and check the **Engine CPU Utilization** graph. If the graph shows high activity for the period of slow response, the engine

might be overloaded. If the graphs for all engines on this server show low activity, a network problem is more likely.

See also

- *Displaying Engine CPU Utilization* on page 69

Error: No Result Set for this Query

Problem: The agent log contains one or more instances of the error “No result set for this query.”

Solution: This error occurs when queries executed by Sybase Control Center for Adaptive Server cannot be completed. To solve the problem, try increasing the space available to the tempdb on Adaptive Server. This example uses the **alter database** command to increase the size of tempdb by 20MB:

```
alter database tempdb on tempdb_dev=20
```

For more information on increasing the size of the tempdb, see:

- The **alter database** command in the Adaptive Server *Reference Manual: Commands*
- The chapter on temporary databases in the Adaptive Server *Performance and Tuning Series: Physical Database Tuning*

If the problems persists, contact Sybase technical support.

Collection Job for Adaptive Server Fails

Problem: A collection job for Adaptive Server may fail when the number of open databases is too low.

Solution: Modify the value of number of open databases by using either the Server Configuration screen of the Adaptive Server monitor, or these steps:

1. Log in to the Adaptive Server.
`isql -S<server_name> -U<sa user name> -P<sa password>`
2. Run this command to display the current configuration value:
`sp_configure 'number of open databases'`
3. Run this command to change the current configuration value:
`sp_configure 'number of open databases', <number>`
Add 10 to the current configuration value and substitute this number for <number>.

See also

- *Modifying Server Configuration Parameters* on page 80

Cannot Authenticate Server Configured with a Multibyte Character Set

Problem: If the Adaptive Server is configured to use a language that requires a multibyte character set such as Chinese, an attempt to authenticate the server fails if the correct character set is not specified in the connection profile for the server.

Solution: The character set for the connection profile can be specified in either of these ways:

- On the Resource Registration screen while registering the resource.
- On the Connection page of the Properties dialog for the Adaptive Server resource.

For example, if your Adaptive Server is using the Chinese language, it may be using character set **gb18030**. In this case, specify **gb18030** as the character set.

See also

- *Registering an Adaptive Server* on page 7

Glossary: Sybase Control Center for Adaptive Server

Glossary of Sybase Control Center terms related to Adaptive Server.

See the Adaptive Server glossary for a complete list of Adaptive Server terms.

- **Adaptive Server** – a server in the Sybase client/server architecture that manages multiple databases and multiple users, keeps track of the actual location of data on disks, maintains mapping of logical data description to physical data storage, and maintains data and procedure caches in memory. Sybase Control Center can manage multiple Adaptive Servers.
- **alert** – a mechanism for notifying administrators when a managed resource experiences a status change, or when a performance metric passes a user-specified threshold.
- **alert instance** – a copy of an alert type that has been configured to generate a particular kind of alert for a specific managed resource.
- **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 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** – (1) A template that defines the resource type and the key performance indicator associated with an alert instance. (2) The basis on which an alert fires: state or threshold. Some alerts are triggered by the state of their key performance indicator (for example, running or stopped), while other alerts are triggered when their KPI's numerical value passes a specified threshold.
- **availability** – indicates whether a resource is accessible and responsive.
- **blocking** – waiting for a lock; a task that needs to acquire a lock on a row, page, or table must wait, or block, if another process holds an incompatible lock on its target object.
- **cache** – See data cache, procedure cache, or statement cache.
- **chart trend period** – the period, in minutes, over which data is displayed in historical charts. Set the chart trend period on the Settings screen of the Adaptive Server monitor. Contrast with screen refresh interval.
- **collection** – a named, predefined set of key performance indicators for which values are collected from monitored servers at the same time. Collections supply the performance and availability data shown on Sybase Control Center screens and charts. Use the scheduler to view a list of collections and to control which collections run, how often they run, and the length of time for which they run.

- **collection repeat interval** – the period, in seconds, between successive repetitions of a statistics collection job. The collection repeat interval determines how often new data on historical monitoring screens is available to be refreshed. Set the collection repeat interval in the scheduler. See also screen refresh interval.
- **data cache** – also called buffer cache and named cache. An area of memory within Adaptive Server that contains the images of database pages and the data structures required to manage the pages. Each cache is given a unique name that is used for configuration purposes. By default, Adaptive Server has a single cache named “default data cache.” Caches configured by users are called user-defined caches.
- **device** – in Adaptive Server, any piece of a disk or file in the file system used to store databases and their component objects.
- **engine** – an instance of the Adaptive Server executable that can communicate with other Adaptive Server engines in shared memory. An Adaptive Server running on a uniprocessor machine always has one engine, engine 0. An Adaptive Server running on a multiprocessor machine can have one or more engines.
- **event** – an activity in the system, such as a user logging in, a service starting or stopping, or a condition changing. Use the alerts feature to detect and notify you about system events.
- **heat chart** – a graphical view of resource availability in the current perspective.
- **index** – a database object that consists of key values from data tables and pointers to the pages that contain those values. Indexes speed up access to data rows by pointing Adaptive Server to the location of a table column’s data on disk.
- **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 Server Availability and Number of Blocked Processes.
- **lock** – a concurrency control mechanism that protects the integrity of data and transaction results in a multiuser environment. Adaptive Server applies table, page, and row locks to:
 - Prevent two or more users from changing the same data at the same time
 - Prevent processes from reading data that is in the process of being changed
- **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 ways that make sense in your environment—for example by location, department, or project.
- **procedure cache** – memory used for stored procedures, batch query plans, triggers, the statement cache, datachange tracking, query compilation, and other objects used during query execution.

- **query plan** – the ordered set of steps required to carry out a SQL query, complete with the access methods chosen for each table. Query plans are chosen by the Adaptive Server optimizer.
- **repository** – a database that stores information related to managed resources, along with user preference data, operational data, and performance statistics.
- **resource** – a unique Sybase product component (such as a server) or a subcomponent.
- **SCC-enabled login account** – a user account that has been granted privileges in Sybase Control Center by mapping appropriate Sybase Control Center roles. (Roles are typically mapped to a group to which the account belongs rather than to the account itself.) The user account and group can be native to Sybase Control Center or created in the operating system or the LDAP directory service to which Sybase Control Center authentication is delegated. You must use an SCC-enabled account to log in to Sybase Control Center.
- **schedule** –
the definition of a task (such as the collection of a set of statistics) and the time interval at which Sybase Control Center executes the task.
- **screen refresh interval** – the period in seconds between refreshes of screens in the Adaptive Server component of Sybase Control Center. Refreshing a screen redraws it with the most recent available data. Set the screen refresh interval on the Settings screen of the Adaptive Server monitor. See also collection repeat interval.
- **segment** – space allocated on one or more database devices. Segments can be used to control the placement of tables and indexes on specific database devices.
- **semaphore** – a simple internal locking mechanism that prevents a second task from accessing the data structure currently in use. Adaptive Server uses semaphores to protect transaction logs, user log caches, and I/O devices. A semaphore is relevant only in symmetric multiprocessing (SMP) environments.
- **statement cache** – memory used to store computed query plans. The statement cache is part of the procedure cache.
- **transaction** – a set of related SQL statements that are treated as a single unit of work. To ensure consistency, if all the statements in the set cannot be executed, the changes made by the query are rolled back. The tables queried during the transaction are locked until a transaction is completed.
- **Transact-SQL** – the SQL dialect used in Sybase Adaptive Server.
- **trend period** – See chart trend period.
- **view** – a window in a perspective that displays information about one or more managed resources. Some views also let you interact with managed resources or with Sybase Control Center itself. For example, the Perspective Resources view lists all the resources managed by the current perspective. Other views allow you to configure alerts, view the topology of a replication environment, and graph performance statistics.
- **wait event** – a condition that causes an Adaptive Server process to pause and wait for another event. Common wait events are waiting for disk I/O to complete, waiting on the Adaptive Server scheduler runnable queue for a CPU to become available, and waiting for another process's lock on a table to be released.

Index

A

Adaptive Server

- configuring to be monitored 6
- criteria for declaring down 21, 22
- defined 93
- enabling performance statistics collection 14, 20
- performance overview screen 53
- versions supported 1

administer, Adaptive Server

- authenticating Unified Agent 39
- error log 38
- server shutdown 38
- server start 38
- start Adaptive Server 40
- stop Adaptive Server 41
- Unified Agent tasks 38
- view error log 41

adminster, Adaptive Server

- registering Unified Agent 33, 38

alert instances

- defined 93

alert list size

- Adaptive Server, setting 37, 83

alert notifications

- defined 93

alert storm

- defined 93

alert storm suppression

- defined 93

alert templates

- Adaptive Server 26

alert type

- defined 93

alert-triggered scripts

- example 25

alerts

- configuring duplicate alerts 24
- configuring escalations 24
- configuring storm suppression 24
- configuring subscriptions 24
- configuring to execute scripts 24
- creating 23
- defined 93

- scripts executed by 30
- setting triggering states and thresholds 23
- state values in Adaptive Server 30
- substitution parameters for scripts 31
- triggering scripts, about 30

attempts_reopen_con parameter, Adaptive Server

- setting 21, 22

authenticate Adaptive Server, troubleshooting 91

availability

- defined 93

availability statistics

- Adaptive Server 83

B

blocking

- defined 93

blocking processes

- displaying SQL for 72
- displaying wait events for 72
- identifying 72

buffer cache

- defined 94

C

cache

- user log, for Adaptive Server 78
- See also data cache

chart trend period

- Adaptive Server, setting 37, 83
- defined 93
- interaction with collection repeat interval in displaying Adaptive Server graphs 88

clocks

- synchronizing to correct missing data problems in Adaptive Server 88

collection job

- for Adaptive Server in Replication, troubleshooting 90

collection repeat interval

- defined 94

- interaction with chart trend period in
 - displaying graphs 88
- collection_ase_all_client_kpis
 - Adaptive Server statistics collection 14, 20
- collection_ase_availability
 - Adaptive Server statistics collection 14
- collection_ase_histmon
 - Adaptive Server statistics collection 14, 20
- collection_ase_rat
 - Adaptive Server statistics collection 14, 20
- collections
 - defined 93
- component log
 - Adaptive Server 38
- configuration
 - optional 34
- configuration file
 - Adaptive Server 6
- configuration parameters, Replication
 - configuring 82
 - setting up 82
- configuration values
 - displaying 79
- configuring
 - replication parameters 82
- connection profile
 - troubleshooting for Adaptive Server 88
- counts 85
- CPU utilization, Adaptive Server
 - displaying for an engine 69
- cumulative values 85

D

- data caches, Adaptive Server
 - buffer pools 57
 - defined 94
 - interpreting display 57
 - modifying size of 56
 - monitoring 55
- data collection jobs
 - adding 12
 - creating 12
- data collection schedules
 - adding 12
- data missing on Adaptive Server monitoring screens
 - 88
- database data, Adaptive Server
 - in cache 57

- databases, Adaptive Server
 - backup status of 62
 - displaying devices used 63
 - displaying disk usage 63
 - displaying processes for 63
 - displaying resources used 63
 - displaying segments used 63
 - displaying segments used by 65
 - displaying unused indexes 63
 - in-memory 61
 - interpreting display 64
 - managing size 64
 - monitoring 62
- delta values 85
- device semaphore 67
- devices, Adaptive Server 67
 - defined 94
 - determining I/O response time and I/O per
 - second 66
 - displaying for a database 63
 - interpreting display 67
- disk usage
 - displaying for an Adaptive Server database 63
- display and copy options 1
- displaying
 - Adaptive Server log 38

E

- engine CPU utilization, Adaptive Server
 - displaying 69
- engines, Adaptive Server
 - defined 94
 - interpreting display 69
- error messages
 - troubleshooting for Adaptive Server 87
- events
 - defined 94

F

- frequency of revalidation in Adaptive Server
 - setting 21, 22

G

- glossaries
 - Adaptive Server terms 93

- graphing statistics 36
 - troubleshooting for Adaptive Server 88
- groups, Adaptive Server
 - assigning monitoring and administration roles 11

H

- heat chart 35
 - customizing columns 35
 - defined 94
 - displaying 35
 - filtering resources displayed 35
- historical performance monitoring 36

I

- I/O rate for Adaptive Server 66
- I/O response time for Adaptive Server 66
- icons
 - on Adaptive Server Processes screen 74
- in-memory storage, Adaptive Server
 - interpreting display 61
 - monitoring 61
- indexes
 - defined 94
 - unused, displaying for an Adaptive Server database 63
- interfaces files, importing resources from 9

J

- JDBC
 - controlling reconnection attempts for Adaptive Server 21, 22
- jobs
 - defined 94

K

- key performance indicators
 - Adaptive Server 15
 - defined 94
 - state values in Adaptive Server 30
- KPIs
 - defined 94

L

- locks
 - defined 94

- login accounts
 - Adaptive Server, configuring monitoring role 6
 - authenticating 12
- login accounts, Adaptive Server
 - assigning monitoring and administration roles 11
- logs
 - agent log, no result set for this query error 90
- LRU buffers, Adaptive Server 57

M

- managing, Adaptive Server
 - registering 7
- mon_role, granting 6
- monitored objects in Adaptive Server
 - refreshing SCC's list of 21, 22
- monitoring
 - performance 36
 - unavailable for Adaptive Server, troubleshooting 88
- monitoring, Adaptive Server 1
 - cluster instances 44, 45
 - cluster interconnect, cluster interprocess communication, CIPC Links, CIPC Mesh, CIPC EndPoints 46, 47
 - cluster overview, displaying 43
 - cluster workload management 49, 50
 - cluster workload management, load profile 49, 50
 - cluster workload management, logical clusters 49, 50
 - cluster workload management, routes 49, 50
 - cluster workload management, workloads 49, 50
 - configuring for 6
 - overview, displaying 52
 - overview, interpreting display 53
- MRU buffers, Adaptive Server 57
- multibyte character sets in Adaptive Server, troubleshooting 91

N

- named cache
 - defined 94
- no result set for this query error 90

P

- parameters
 - displaying 79
- parameters for scripts 31
- percentage or ratio values 85
- performance statistics
 - Adaptive Server 83
- Perspective Heat Chart view 35
- perspectives
 - adding resources 10
 - creating 10
 - defined 94
- pools in caches, Adaptive Server 57
- procedure cache, Adaptive Server
 - defined 94
 - interpreting display 59
 - monitoring 58
- processes, Adaptive Server
 - blocked, identifying 71
 - blocking, identifying 71, 72
 - blocking, terminating 72
 - displaying for a database 63
 - displaying SQL and query plans for 73
 - displaying wait events for 74
 - finding 70
 - interpreting display 74
 - monitoring 70

Q

- queries, SQL
 - in Adaptive Server statement cache 59
- query plan
 - defined 95
- query plan for an Adaptive Server process,
 - displaying 73

R

- rate values 85
- registering, Adaptive Server 7
- RepAgent threads
 - monitoring in Adaptive Server 81
- Replication Agent screen, Adaptive Server
 - interpreting display 82
- Replication Agent threads
 - monitoring in Adaptive Server 81
- Replication Agents
 - running in Adaptive Servers 81

- replication parameters
 - configuring 82
 - setting up 82
- repository
 - defined 95
- resources
 - adding to a perspective 10
 - authenticating 12
 - defined 95
 - displaying availability 35
 - importing in batch 9
- response time for Adaptive Server, troubleshooting 89
- revalidation_frequency parameter, Adaptive Server
 - and missing data 88
 - setting 21, 22
- roles, Adaptive Server
 - assigning to users and groups 11
- RSSD user name, using to authenticate 12

S

- SCC-enabled login account
 - defined 95
- schedules
 - creating for a data collection job 12
 - defined 95
- scope of Adaptive Server statistics 85
- screen refresh interval
 - Adaptive Server, setting 37, 83
 - defined 95
 - setting 14, 20
- scripts
 - example 25
 - substitution parameters 31
 - triggered by alerts 30
- segments
 - finding reserved table space for 76
- segments, Adaptive Server
 - defined 95
 - displaying for a database 63, 65
 - extending 76
 - interpreting display 77
- semaphore
 - defined 95
- server configuration, Adaptive Server 80
 - interpreting display 80
- servers
 - authenticating 12

- displaying availability 35
 - importing in batch 9
 - setting up
 - replication parameters 82
 - Settings screen
 - for Adaptive Server monitoring 37, 83
 - sp_configure stored procedure 6, 80
 - sp_role stored procedure 6
 - SQL
 - displaying cached queries in Adaptive Server 59
 - displaying for Adaptive Server processes 73
 - monitoring queries 81
 - sql.ini files, importing resources from 9
 - statement cache, Adaptive Server
 - defined 95
 - interpreting display 60
 - monitoring 59
 - states
 - values for alerts in Adaptive Server 30
 - values for KPIs in Adaptive Server 30
 - statistics
 - about 14
 - Adaptive Server, about 83
 - Adaptive Server, interpreting 85
 - Adaptive Server, scope of 85
 - availability 14
 - for Adaptive Server, enabling collection 14, 20
 - performance 14
 - statistics chart
 - graphing performance counters 36
 - Statistics Chart
 - unavailable for Adaptive Server, troubleshooting 88
 - stored procedures, Adaptive Server
 - in cache 59
 - storm suppression 24
 - substitution parameters for scripts 31
 - Sybase Control Center for Adaptive Server 1
 - Sybase Control Center, Adaptive Server
 - managing 37
 - monitoring 37
 - system resources, Adaptive Server
 - identifying processes that use 70
- T**
- tables, in Adaptive Server
 - columns, in Adaptive Server
 - choosing for display 1
 - copying 1
 - nodes, cluster, in Adaptive Server
 - expanding and collapsing 1
 - rows, in Adaptive Server
 - copying 1
 - tempdb, Adaptive Server
 - no result set for this query error 90
 - terms
 - Adaptive Server 93
 - time_between_reattempts parameter, Adaptive Server
 - setting 21, 22
 - Transact-SQL
 - defined 95
 - transactions, Adaptive Server
 - defined 95
 - identifying associated processes 78
 - interpreting display 78
 - monitoring 78
 - user log cache usage 78
 - trend period
 - setting 14, 20
 - troubleshooting
 - Adaptive Server 87
- U**
- update KPI
 - troubleshooting for Adaptive Server 87
 - user interface, Adaptive Server
 - choose columns 1
 - collapse all nodes 1
 - copy selected row 1
 - copy table 1
 - expand all nodes 1
 - options 1
 - user log cache usage for Adaptive Server
 - transactions 78
- V**
- viewing
 - Adaptive Server log 38
 - views
 - defined 95

W

wait event
 defined 95

wait events
 displaying for Adaptive Server processes 74
wait events, Adaptive Server 74