



Sybase® Adaptive Server™ Enterprise
Managing and Monitoring Sybase Adaptive Server Enterprise

Adaptive Server™

Document ID: 35138-01-1150

September 1997

Copyright Information

Copyright © 1989–1997 by Sybase, Inc. All rights reserved.

Sybase, Inc., 6475 Christie Avenue, Emeryville, CA 94608.

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. 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.

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, the Sybase logo, APT-FORMS, Certified SYBASE Professional, Data Workbench, First Impression, InfoMaker, PowerBuilder, Powersoft, Replication Server, S-Designer, SQL Advantage, SQL Debug, SQL SMART, SQL Solutions, Transact-SQL, VisualWriter, and VQL are registered trademarks of Sybase, Inc. Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Monitor, ADA Workbench, AnswerBase, Application Manager, AppModeler, APT-Build, APT-Edit, APT-Execute, APT-Library, APT-Translator, APT Workbench, Backup Server, BayCam, Bit-Wise, ClearConnect, Client-Library, Client Services, CodeBank, Column Design, Connection Manager, DataArchitect, Database Analyzer, DataExpress, Data Pipeline, DataWindow, DB-Library, dbQ, Developers Workbench, DirectConnect, Distribution Agent, Distribution Director, Dynamo, Embedded SQL, EMS, Enterprise Client/Server, Enterprise Connect, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, EWA, Formula One, Gateway Manager, GeoPoint, ImpactNow, InformationConnect, InstaHelp, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Logical Memory Manager, MainframeConnect, Maintenance Express, MAP, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, MethodSet, Net-Gateway, NetImpact, Net-Library, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, PB-Gen, PC APT-Execute, PC DB-Net, PC Net Library, Power++, Power AMC, PowerBuilt, PowerBuilt with PowerBuilder, PowerDesigner, Power J, PowerScript, PowerSite, PowerSocket, Powersoft Portfolio, Power Through Knowledge, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Quickstart Datamart, Replication Agent, Replication Driver, Replication Server Manager, Report-Execute, Report Workbench, Resource Manager, RW-DisplayLib, RW-Library, SAFE, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, smart.partners, smart.parts, smart.script, SQL Anywhere, SQL Central, SQL Code Checker, SQL Edit, SQL Edit/TPU, SQL Modeler, SQL Remote, SQL Server, SQL Server/CFT, SQL Server/DBM, SQL Server Manager, SQL Server SNMP SubAgent, SQL Station, SQL Toolset, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Gateways, Sybase IQ, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase Synergy Program, Sybase Virtual Server Architecture, Sybase User Workbench, SybaseWare, SyBooks, System 10, System 11, the System XI logo, SystemTools, Tabular Data Stream, The Architecture for Change, The Enterprise Client/Server Company, The Model for Client/Server Solutions, The Online Information Center, Translation Toolkit, Turning Imagination Into Reality, Unibom, Unilib, Uninull, Unisep, Unistring, Viewer, Visual Components, VisualSpeller, WarehouseArchitect, WarehouseNow, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, and XA-Server are trademarks of Sybase, Inc. 6/97

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

Table of Contents

About This Book

Audience	xxix
How to Use This Book	xxix
Release Compatibility	xxxix
Related Documents	xxxii
Other Sources of Information	xxxiii
Training	xxxiii
Conventions	xxxiii
Naming Conventions	xxxiii
Style Conventions	xxxiv
Shortcuts	xxxiv
Privileges	xxxv
Mouse Buttons	xxxv
Screen Images	xxxvi
If You Need Help	xxxvi
Using Online Help	xxxvi
Contacting Technical Support	xxxvii

1. Introduction

About Sybase Central	1-1
Performing System and Database Administration Tasks	1-2
Monitoring the Performance of Adaptive Server	1-2
How Sybase Central Makes Administration Easier	1-2

2. Sybase Central Windows and Dialog Boxes

What's in This Chapter	2-1
Starting Sybase Central	2-1
Stopping Sybase Central	2-1
The Sybase Central Window	2-1
Sybase Central Features	2-2
Window Panes	2-3
Folder and Object Icons	2-3
Selecting Objects	2-3
Moving Through the Sybase Central Object Tree	2-4
Customizing the Display	2-6
Selecting Display Formats	2-6

Turning the Toolbar and Status Bar On and Off	2-8
Using Drag-and-Drop Shortcuts	2-8
Valid Drag-and-Drop Operations	2-9
Using Keyboard Shortcuts in the Sybase Central Window	2-11
Menus and Toolbars	2-12
Standard Menus	2-12
Shortcut Menus	2-13
Toolbar and Status Bar	2-14
Toolbar	2-14
Refreshing Window and Dialog Box Displays	2-16
Using the Code Editor	2-16
Opening the Code Editor	2-17
Syntax Highlighting	2-18
Finding and Replacing Code in the Code Editor	2-18
Executing Code in the Code Editor	2-18
Printing Code	2-18
The Code Editor Toolbar	2-18
Using the Sybase Tools Console	2-20

3. Common Tasks

What's in This Chapter	3-1
Creating an Object	3-1
Viewing the Properties of an Object	3-2
Using Property Sheet Tabs	3-3
Changing the Properties of an Object	3-3
Closing a Property Sheet	3-3
Navigating to the Properties of a Related Object	3-4
Deleting an Object	3-4
Getting Help	3-5
Generating and Executing Adaptive Server DDL Scripts	3-5
Generating DDL Scripts	3-6
Executing DDL Scripts	3-6

4. Managing the Adaptive Server Operating Environment

What's in This Chapter	4-1
Connecting to and Disconnecting from Adaptive Server	4-1
Identifying Servers to Sybase Central	4-1
Connecting to Adaptive Server	4-2
Using Component Integration Services Passthrough Connections ..	4-4

Using Connection Profiles	4-4
Creating a Connection Profile	4-4
Connecting to Adaptive Server When Sybase Central Starts	4-5
Using Connection Profiles to Connect Manually	4-5
Editing a Connection Profile	4-6
Diagnosing an Unsuccessful Connection	4-6
Disconnecting from Adaptive Server	4-6
Starting and Stopping Adaptive Server	4-7
Starting Adaptive Server	4-7
Pausing Adaptive Server	4-7
Resuming Operation of A Paused Adaptive Server	4-8
Shutting Down Adaptive Server	4-8
Displaying Adaptive Server Properties	4-9
Configuring Adaptive Server	4-10
Who Can Set Configuration Parameters	4-10
Parameters Requiring Restart	4-10
Setting Configuration Parameters	4-11
Managing Adaptive Server Processes	4-12
Displaying the Properties of a Process	4-12
Process Parameters	4-13
Setting Execution Attributes for A Process	4-14
Returning to the Default Execution Attributes	4-14
Killing a Process	4-15

5. Using Monitors

What's in This Chapter	5-1
Introduction to Adaptive Server Enterprise Monitor	5-1
Adaptive Server Enterprise Monitor Architecture	5-2
Configuration Procedures	5-3
Starting Monitor Server	5-4
Configuring a Monitors Folder	5-5
Privileges Required for Monitoring	5-5
Opening and Closing Monitors	5-6
Opening Monitors	5-6
Closing Monitors	5-7
Monitor Window Description	5-8
Contents of Monitor Status Bar	5-9
Showing and Hiding the Status Bar	5-9
Monitor Menus	5-10
Setting up a Monitor Session	5-10

Setting the Statistic Type	5-11
Setting the Sample Interval	5-11
Notes about Setting Sample Interval	5-12
Pausing and Resuming Data Refresh	5-13
Setting Filters	5-13
Canceling A Filter	5-14
Setting Filters on Numeric Data	5-14
Setting Character String Filters without a Scroll List	5-16
Setting a Database Object Filter	5-17
Setting a Data Cache Filter	5-19
Setting a Lock Status Filter	5-20
Setting a Lock Type Filter	5-21
Setting a Process ID Filter	5-22
Setting a Process State Filter	5-23
Setting a Device Filter	5-24
Setting an Engine Filter	5-25
Setting up the Performance Trends Monitor	5-26
Troubleshooting	5-27
Unsuccessful Connections	5-27
Unexpected Adaptive Server Name	5-28
Cumulative Values Return to Zero	5-29
6. Managing Remote Servers	
What's in This Chapter	6-1
Configuring RPC Handling	6-1
Configuring Adaptive Server to Use Site Handler	6-2
Configuring Adaptive Server to Use CIS RPC Handling	6-3
Adding a Remote Server	6-4
Deleting a Remote Server	6-5
Displaying Remote Server Properties	6-5
Setting the Server Class of a Remote Server	6-6
Setting Security Options for Remote Servers	6-7
Specifying an RPC Security Model	6-8
Managing Remote Server Logins	6-10
Mapping Logins for Site Handler RPC Handling	6-10
Setting Default Login Mapping	6-12
Setting Explicit Login Mappings	6-13
Removing a Specific Login Mapping	6-14
Managing Logins for CIS RPC Handling	6-14
Removing a Login Mapping for CIS RPC Handling	6-16

Creating and Executing DDL Scripts	6-16
--	------

7. Managing Cache

What's in This Chapter	7-1
Creating a Named Data Cache	7-2
Deleting a Cache	7-3
Displaying Cache Properties	7-3
Changing Cache Size or Type	7-4
Defining Buffer Pools	7-5
Creating a Buffer Pool	7-7
Changing a Buffer Pool	7-8
Deleting a Buffer Pool	7-9
Binding Objects to a Cache	7-10
Rules for Binding	7-10
Creating a Binding	7-11
Deleting a Binding	7-13
Creating and Executing DDL Scripts	7-13

8. Managing Database Devices

What's in This Chapter	8-1
About Database Devices	8-2
Strategies for Recovery and Performance	8-3
Creating Database Devices	8-3
Deleting a Database Device	8-4
Displaying Device Properties	8-5
Examining Database Allocations on a Database Device	8-6
Examining Segment Mapping on a Database Device	8-6
Using Disk Mirroring	8-7
Adding Disk Mirroring	8-9
Disabling and Re-enabling Disk Mirroring	8-9
Disabling Disk Mirroring	8-10
Re-Enabling Disk Mirroring	8-11
Storage Management Considerations and Tips	8-12
Managing for Recovery	8-12
Changing Default Device Designation	8-12
Using Disk Mirroring	8-13
Separating Data and Transaction Log	8-13
Managing for Performance	8-14
Creating and Executing DDL Scripts	8-14

9. Managing Databases

What's in This Chapter	9-1
Considerations for Creating a New Database	9-1
Estimating Database Size	9-1
Deciding Where to Allocate Storage	9-2
Creating a Database	9-2
Deleting a Database	9-4
Deleting a Damaged Database	9-4
Displaying and Modifying Database Properties	9-4
Transferring Database Ownership	9-6
Modifying Database Storage Allocations	9-7
Examining Database Storage Allocations	9-7
Allocating Space on the Same Device	9-8
Allocating Space on a New Device	9-10
Changing Log Buffer Size	9-11
Changing Database Options	9-12
Database Options	9-13
Generating and Executing Database DDL Scripts	9-15
Generating DDL to Create a Database	9-15
Generating DDL for the Objects in A Database	9-16
Executing Database DDL Files	9-16
Issuing a Database Checkpoint	9-17
Checking Database Consistency	9-18
Consistency Check Options	9-19
Overall Database Option	9-19
Check Database Allocation	9-20
Check System Catalog	9-21
Repair and Drop Database	9-22
Setting a Database Online	9-22

10. Backing Up and Recovering Databases

What's in This Chapter	10-1
Preparing for Backup and Recovery	10-1
Developing a Backup and Recovery Plan	10-1
Identifying the Backup Server	10-2
Managing Dump Devices	10-2
Creating a Dump Device	10-2
Deleting a Dump Device	10-3
Examining Dump Device Properties	10-4
Backing Up a Database	10-5

Guidelines for Backing Up User Databases	10-5
Guidelines for Backing Up the <i>master</i> Database.....	10-5
Backup Procedure.....	10-6
To Proceed with the Backup	10-7
To Stop a Backup	10-7
To Get Information on Backup Server Messages.....	10-8
Restoring a Database	10-8
Deciding How to Restore.....	10-8
Recovery Considerations	10-8
Responding to Volume Change Prompts	10-10
Stopping A Restore	10-10
Getting Information About Backup Server Messages	10-10

11. Managing Segments

What's in This Chapter	11-1
Database Segments	11-1
Creating a Segment.....	11-2
Deleting a Segment	11-2
Displaying and Modifying Segment Properties	11-3
Extending a Segment	11-4
Displaying Objects Stored on a Segment	11-6
Using Segment Thresholds	11-6
Last-Chance Threshold	11-7
Hysteresis Value	11-7
Creating a Threshold	11-8
Modifying a Threshold	11-10
Deleting a Threshold	11-11
Creating and Executing DDL Scripts	11-11

12. Controlling Access to Adaptive Server

What's in This Chapter	12-1
Permissions and Sybase Central.....	12-2
About Access Privileges and Permissions	12-2
Access Based on Role.....	12-3
System Roles	12-3
Access Based on Special User Status.....	12-4
Access Based on Group Membership	12-5
Access Based on Command Permissions	12-6
Access Based on Object Permissions	12-6

Access Example	12-8
Summary of Access Privileges and Permissions	12-9
Summary of Transferable Privileges and Permissions	12-13

13. Managing Logins

What's in This Chapter	13-1
Creating Logins	13-1
Deleting A Login	13-2
Displaying Login Properties	13-2
Locking A Login	13-4
Changing a Login Password	13-6
Displaying Databases Owned by a Login	13-7
Assigning Roles to a Login	13-7
Setting Automatic Activation of A Role	13-9
Removing Automatic Activation of a Role	13-9
Sharing Logins by Using Aliases	13-9
Adding an Alias to a Login	13-10
Deleting an Alias from a Login	13-11
Creating and Executing DDL Scripts	13-11

14. Managing Users and Groups

What's in This Chapter	14-1
Managing Users	14-1
Creating Users	14-2
Deleting a User	14-3
Displaying User Properties	14-3
Changing a User's Group	14-4
Displaying Objects Owned by the User	14-5
Granting and Revoking Object Permissions for a User	14-5
Granting and Revoking Command Permissions for a User	14-8
Aliasing Users to Logins	14-9
Removing a Login Alias	14-11
Creating and Executing DDL Scripts for Users	14-11
Managing Groups	14-12
Creating a Group	14-12
Deleting a Group	14-12
Displaying Group Properties	14-13
Adding a User to a Group	14-14
Removing Users from a Group	14-15

Granting and Revoking Object Permissions for a Group	14-16
Granting and Revoking Command Permissions for a Group	14-18
Creating and Executing DDL Scripts for Groups	14-19

15. Managing Roles

What's in This Chapter	15-1
Creating a User-defined Role.	15-2
Deleting a User-defined Role.	15-2
Displaying Role Properties	15-3
Setting A Password For A Role	15-4
Assigning A Role to A Login	15-5
Removing A Role Assignment from A Login.	15-6
Activating and Deactivating A Role	15-7
Activating A Role By Default	15-8
Removing Automatic Role Activation.	15-8
Deactivating A Role	15-8
Specifying Mutual Exclusivity	15-9
Specifying A Role's Hierarchy	15-11
Example of Creating a Role Hierarchy	15-13
Managing Object Permissions for A Role.	15-15
Managing Command Permissions for A Role.	15-17
Creating and Executing DDL Scripts	15-18

16. Managing Defaults

What's in This Chapter	16-1
Creating a Default	16-1
Deleting a Default	16-3
Displaying Default Properties	16-3
Displaying Default Dependencies	16-4
Binding a Default.	16-5
Binding a Default to a Table Column.	16-5
Binding a Default to a User-defined Datatype	16-7
Unbinding A Default from A Table Column	16-8
Unbinding A Default from A User-defined Datatype.	16-8
Creating and Executing DDL Scripts	16-8

17. Managing Stored Procedures

What's in This Chapter	17-1
----------------------------------	------

Creating a Stored Procedure	17-1
Deleting a Stored Procedure	17-2
Viewing Stored Procedure Code	17-3
Displaying Stored Procedure Properties	17-4
Displaying Stored Procedure Dependencies	17-5
Setting Permissions for A Stored Procedure	17-6
Managing Extended Stored Procedures.	17-8
Creating An Extended Stored Procedure	17-8
Deleting an Extended Stored Procedure	17-9
Viewing Properties of an Extended Stored Procedure	17-10
Adding Parameters to an Extended Stored Procedure	17-11
Setting Permissions for an Extended Stored Procedure	17-11
Extended Stored Procedure Dependencies.	17-11
Configuring Adaptive Server to Use Extended Stored Procedures . .	17-12
Creating and Executing DDL Scripts	17-13

18. Managing Rules

What's in This Chapter	18-1
Creating a Rule	18-1
Deleting a Rule	18-3
Displaying Rule Properties	18-4
Displaying Rule Dependencies	18-5
Binding a Rule.	18-6
Binding a Rule to a Table Column	18-6
Binding a Rule to a User-defined Datatype	18-7
Unbinding A Rule	18-9
Unbinding A Rule From A Table Column	18-9
Unbinding A Rule From A User-defined Datatype.	18-9
Creating and Executing DDL Scripts	18-9

19. Managing Triggers

What's in This Chapter	19-1
Creating a Trigger	19-2
Deleting a Trigger	19-2
Editing A Trigger 's Code.	19-3
Displaying Trigger Properties	19-4
Displaying Trigger Dependencies	19-5
Creating and Executing DDL Scripts	19-6

20. Managing User-defined Datatypes

What's in This Chapter	20-1
Creating a User-defined Datatype	20-1
Deleting a User-defined Datatype	20-3
Displaying User-defined Datatype Properties	20-4
Binding Rules to User-defined Datatypes	20-5
Displaying User-defined Datatype Dependencies	20-5
Creating and Executing DDL Scripts	20-6

21. Managing Views

What's in This Chapter	21-1
Creating a View	21-1
Deleting a View	21-2
Displaying View Properties	21-3
Editing a View	21-4
Displaying View Dependencies	21-5
Granting Permissions on a View	21-6
Creating and Executing DDL Scripts	21-8

22. Managing User Tables

What's in This Chapter	22-1
Understanding Constraints	22-2
Implementing Rules and Check Constraints	22-2
Implementing Referential Constraints	22-3
Constraints Summary	22-3
Using the Table Editor	22-4
Opening the Table Editor	22-5
Creating a Table	22-6
Specifying Column Definitions	22-7
Specifying Column Check Constraints	22-8
Check Constraint Format	22-9
Specifying Primary Keys	22-9
Creating User Tables When Component Integration Services is Enabled ..	22-11
Deleting a Table	22-11
Displaying Table Properties	22-12
Adding Columns to a Table	22-13
Setting Cache Strategy	22-14

Displaying Table Dependencies	22-14
Setting Permissions for A Table	22-16
Placing a Table on a Segment	22-18
Creating Check Constraints on a Table	22-19
Creating Referential Constraints on a Table	22-21
Example of a Foreign Key	22-22
Displaying the Properties of A Foreign Key	22-23
Creating Unique Constraints	22-24
Viewing the Properties of a Unique Constraint	22-25
Recompiling Procedures and Triggers	22-26
Checking Table Consistency	22-26
Consistency Check Options	22-27
Overall Table Consistency	22-27
Table Allocations	22-28
Reindex	22-29
Fix Text	22-29
Managing Columns in the Columns Folder	22-30
Viewing the Properties of a Column	22-30
Binding a Rule to a Column	22-31
Unbinding a Rule from a Column	22-31
Binding a Default to a Column	22-32
Unbinding a Default from a Column	22-32
Deleting Check Constraints	22-32
Creating and Executing DDL Scripts	22-33

23. Managing Proxy Tables

What's in This Chapter	23-1
Understanding Proxy Tables	23-1
Component Integration Services Options	23-2
Creating Proxy Tables	23-3
Prerequisites for Creating Proxy Tables	23-3
Creating a Proxy Table	23-3
Considerations for Creating Proxy Tables	23-5
Creating Existing Tables Compared to Creating New Tables	23-6
Deleting a Proxy Table	23-6
Verifying the Status of a Proxy Table	23-6
Managing Proxy Tables	23-7
Creating and Executing DDL Scripts	23-7

24. Managing Indexes

What's in This Chapter	24-1
Creating an Index	24-2
Deleting an Index.....	24-6
Displaying Index Properties.....	24-6
Placing an Index on a Segment	24-7
Setting Maximum Rows Per Page	24-8
Checking Index Consistency.....	24-9
Updating Statistics	24-11
Creating and Executing DDL Scripts	24-11

25. Application Tuning

What's in This Chapter	25-1
Application Tuning Concepts	25-1
What is Execution Preference?	25-2
Execution Class Basics	25-2
Predefined Execution Classes.....	25-2
User-defined Execution Classes.....	25-3
Engine Affinity Basics	25-3
Scheduling with Execution Classes and Engine Groups	25-4
Scope and Precedence Rules	25-6
Precedence Rule	25-7
Scope Rule.....	25-8
Resolving Precedence Conflicts.....	25-8
Managing Engine Groups	25-9
Creating an Engine Group.....	25-9
Deleting an Engine Group.....	25-10
Displaying the Properties of an Engine Group	25-10
Changing the Engines in An Engine Group	25-12
Adding An Engine to An Engine Group	25-13
Deleting An Engine from An Engine Group	25-13
Creating and Executing DDL Scripts for Engine Groups	25-14
Managing Execution Classes.....	25-15
Creating an Execution Class	25-15
Deleting an Execution Class	25-15
Displaying the Properties of an Execution Class	25-16
Associating an Execution Class with an Engine Group	25-17
Binding an Execution Object to an Execution Class	25-17
Specifying the Scope of A Binding	25-18

Unbinding an Execution Object from an Execution Class	25-19
Creating and Executing DDL Scripts	25-19

26. Interpreting Monitor Data

Whats in This Chapter	26-1
Summary of Adaptive Server Monitors	26-2
Definition of Metrics	26-3
#/second for Sample	26-3
#/second for Cumulative	26-3
Rate Examples	26-4
Cumulative Counts	26-4
Cumulative Data for Processes	26-4
Resources for More Information	26-4
Information about CPU Usage and Engines	26-5
Information about Tuning Devices	26-5
Information about Memory Management	26-6
Information about Tuning Data Caches	26-6
Information about Tuning Tables and Indexes	26-6
Information about Tuning SQL Statements	26-7
Information about Processes and Task Management	26-7
Information about Tuning Transactions	26-7
Information about Locks and Lock Contention	26-8
Information about Networks	26-8
Application Activity Monitor	26-9
Cache Monitor	26-14
Data Cache Monitor	26-17
Device I/O Monitor	26-23
Engine Activity Monitor	26-27
Memory Utilization Monitor	26-30
Network Activity Monitor	26-31
Object Lock Status Monitor	26-37
Object Page I/O Monitor	26-42
Performance Summary Monitor	26-45
Performance Trends Monitor	26-52
Process Activity Monitor	26-54
Process Current SQL Statement Monitor	26-59
Stored Procedure Activity Monitor	26-64
Transaction Activity Monitor	26-67

27. Monitoring Techniques

What's in This Chapter	27-1
Creating Benchmarks	27-1
Designing a Benchmark.....	27-2
Statistics to Collect in a Benchmark.....	27-2
Statistics that Describe the Benchmark Environment.....	27-3
Running the Benchmark	27-6
Monitoring Strategies	27-6
Daily Monitoring	27-7
Statistics to Monitor.....	27-7
Recognizing a Problem	27-7
Researching a Stopped Process.....	27-8
Capacity Planning	27-9
Tuning Applications and Stored Procedures.....	27-9
Monitoring <i>tempdb</i> and Transaction Logs.....	27-10
Tuning Strategy.....	27-10
Other Useful Tools.....	27-11

A. Adaptive Server Configuration Parameters

What's in This Appendix	A-1
Adaptive Server Parameter List	A-1
Parameters by Functional Group	A-12
Adaptive Server Administration Parameters.....	A-12
Backup/Recovery Parameters	A-13
Cache Manager Parameters.....	A-14
Component Integration Services	A-14
Disk I/O Parameters	A-14
Error Log	A-14
Extended Stored Procedures.....	A-15
General Information Parameter	A-15
Languages Parameters.....	A-15
Lock Manager Parameters.....	A-15
Memory Use Parameter.....	A-15
Meta-data Caches	A-16
Network Communications Parameters.....	A-16
Operating System Resources Parameters	A-16
Parallel Query	A-16
Physical Memory Parameters.....	A-17
Processors Parameters	A-17
Rep Agent Thread Administration	A-17

Security Related	A-18
User Environment Parameters	A-18

Glossary

Index

List of Figures

Figure 1-1:	Sybase Central main window	1-1
Figure 1-2:	Opening a related property sheet.....	1-3
Figure 1-3:	Main window toolbar	1-3
Figure 2-1:	Sybase Central window	2-2
Figure 2-2:	Adaptive Server object tree	2-5
Figure 2-3:	Object display formats	2-7
Figure 2-4:	Drag-and-drop symbol	2-9
Figure 2-5:	Toolbar and status bar.....	2-14
Figure 2-6:	Sybase Central toolbar	2-15
Figure 2-7:	Code editor	2-17
Figure 2-8:	Code Editor toolbar	2-19
Figure 3-1:	Property sheet tabs	3-3
Figure 3-2:	Properties button on a property sheet	3-4
Figure 4-1:	Sybase Central window	4-2
Figure 4-2:	Logon dialog box	4-2
Figure 4-3:	Connection Profiles dialog box.....	4-4
Figure 4-4:	Create New Profile dialog box.....	4-5
Figure 4-5:	Server property sheet	4-9
Figure 4-6:	Server Configuration dialog box	4-11
Figure 4-7:	Process property sheet	4-12
Figure 4-8:	Process property sheet: Parameters tab.....	4-13
Figure 4-9:	Process property sheet: Execution Attributes tab	4-14
Figure 5-1:	Adaptive Server Enterprise Monitor architecture	5-2
Figure 5-2:	Sample monitor window	5-8
Figure 5-3:	Set Sample Interval dialog box	5-12
Figure 5-4:	Filter menu example.....	5-13
Figure 5-5:	Set <i><integer></i> Filter dialog box.....	5-15
Figure 5-6:	Set <i><string></i> Filter dialog box	5-16
Figure 5-7:	Set Database Object Filter dialog box	5-17
Figure 5-8:	Set Data Cache Filter dialog box.....	5-19
Figure 5-9:	Set Lock Status Filter dialog box	5-20
Figure 5-10:	Set Lock Type Filter dialog box.....	5-21
Figure 5-11:	Set Process Filter dialog box.....	5-22
Figure 5-12:	Set Process State Filter dialog box	5-23
Figure 5-13:	Set Device Filter dialog box.....	5-24
Figure 5-14:	Set Engine Filter dialog box	5-25
Figure 5-15:	Select Data Items dialog box	5-26
Figure 6-1:	Example of configuring site handler RPC handling.....	6-3

Figure 6-2:	Example of configuring CIS RPC handling	6-4
Figure 6-3:	Remote Server property sheet.....	6-5
Figure 6-4:	Remote server property sheet: Security tab	6-8
Figure 6-5:	Remote login mapping options	6-11
Figure 6-6:	Remote Server: Default Login Mapping tab.....	6-12
Figure 6-7:	Remote Server dialog box: Specific Login Mapping tab	6-13
Figure 6-8:	Add Specific Mapping dialog box.....	6-14
Figure 6-9:	Remote server property sheet: CIS Login Mapping tab.....	6-15
Figure 6-10:	Add CIS Mapping dialog box.....	6-15
Figure 7-1:	Cache Properties	7-3
Figure 7-2:	Cache property sheet: Configuration tab	7-4
Figure 7-3:	Default cache and two named data caches	7-6
Figure 7-4:	Cache Property Sheet: Buffer Pool tab.....	7-7
Figure 7-5:	Add/Change Memory Pool dialog box	7-8
Figure 7-6:	Cache property sheet: Cache Bindings tab	7-11
Figure 7-7:	New Cache Binding dialog box.....	7-12
Figure 8-1:	Logical database devices on a physical disk device	8-2
Figure 8-2:	Database Device property sheet	8-5
Figure 8-3:	Database Device property sheet: Databases tab.....	8-6
Figure 8-4:	Database Device property sheet: Segments tab.....	8-7
Figure 8-5:	Database device on B mirrored on D	8-8
Figure 8-6:	Database Device Properties: Mirror tab	8-9
Figure 8-7:	Database Device Properties: Mirror tab	8-10
Figure 8-8:	Database Device property sheet: Parameters tab.....	8-13
Figure 9-1:	Database property sheet	9-5
Figure 9-2:	New Database Owner dialog box	9-6
Figure 9-3:	Expanding a database	9-7
Figure 9-4:	Database property sheet: Usage tab.....	9-8
Figure 9-5:	Database property sheet: Devices tab.....	9-9
Figure 9-6:	Device Size dialog box	9-9
Figure 9-7:	Available Devices dialog box	9-10
Figure 9-8:	Database property sheet: Transaction Log tab	9-11
Figure 9-9:	Database property sheet: Options tab	9-12
Figure 9-10:	Check Overall Database output.....	9-20
Figure 9-11:	Check Database Allocation output	9-21
Figure 9-12:	Check System Catalog output.....	9-21
Figure 10-1:	Dump Device property sheet.....	10-4
Figure 10-2:	Backup Server Confirmation dialog box	10-7
Figure 11-1:	Segment property sheet	11-3
Figure 11-2:	Segment property sheet: Devices tab.....	11-4
Figure 11-3:	Add Database Device to Segment dialog box	11-5

Figure 11-4:	Segment property sheet: Contains tab.....	11-6
Figure 11-5:	Segment property sheet: Thresholds tab.....	11-8
Figure 11-6:	Add New Threshold dialog box.....	11-9
Figure 11-7:	Edit Threshold dialog box.....	11-10
Figure 13-1:	Login property sheet.....	13-3
Figure 13-2:	Login property sheet: Parameters tab.....	13-5
Figure 13-3:	Change Password dialog box.....	13-6
Figure 13-4:	Login property sheet: Databases Owned tab.....	13-7
Figure 13-5:	Login property sheet: Roles tab.....	13-8
Figure 13-6:	Add Roles dialog box.....	13-8
Figure 13-7:	Login property sheet: Users tab.....	13-10
Figure 13-8:	Add Alias dialog box.....	13-10
Figure 14-1:	User property sheet.....	14-4
Figure 14-2:	User property sheet: Objects Owned tab.....	14-5
Figure 14-3:	User property sheet: Object Permissions tab.....	14-6
Figure 14-4:	User property sheet: Command Permissions tab.....	14-8
Figure 14-5:	user property sheet: Logins tab.....	14-10
Figure 14-6:	Add Login Alias dialog box.....	14-10
Figure 14-7:	Group property sheet.....	14-13
Figure 14-8:	Group property sheet: Users tab.....	14-14
Figure 14-9:	Add User to Group dialog box.....	14-15
Figure 14-10:	Group property sheet: Object Permissions tab.....	14-16
Figure 14-11:	Group property sheet: Command Permissions tab.....	14-18
Figure 15-1:	Role property sheet.....	15-3
Figure 15-2:	Change Password for a Role dialog box.....	15-4
Figure 15-3:	Role property sheet: Logins tab.....	15-5
Figure 15-4:	Add Role to Login dialog box.....	15-6
Figure 15-5:	Role Activation dialog box.....	15-7
Figure 15-6:	Role property sheet: Exclusivity tab.....	15-10
Figure 15-7:	Add Mutually Exclusive Roles dialog box.....	15-10
Figure 15-8:	Role property sheet: Hierarchy tab.....	15-12
Figure 15-9:	Add a Role Assignment dialog box.....	15-12
Figure 15-10:	Illustration of role hierarchy.....	15-14
Figure 15-11:	Role property sheet: Object Permissions tab.....	15-15
Figure 15-12:	Role property sheet: Command Permissions tab.....	15-17
Figure 16-1:	New Default Properties dialog box.....	16-2
Figure 16-2:	Default property sheet.....	16-3
Figure 16-3:	Dependencies dialog box for defaults.....	16-4
Figure 16-4:	Column Properties: Rules and Defaults tab.....	16-6
Figure 16-5:	User Datatype Properties: Advanced Options tab.....	16-7
Figure 17-1:	Code Editor for a stored procedure.....	17-3

Figure 17-2:	Stored procedure property sheet.....	17-4
Figure 17-3:	Dependencies dialog box for a stored procedure	17-5
Figure 17-4:	Stored procedure property sheet: Permissions tab.....	17-7
Figure 17-5:	Extended Stored Procedure property sheet.....	17-10
Figure 18-1:	New Rule Properties dialog box.....	18-2
Figure 18-2:	Select a Creator dialog box.....	18-2
Figure 18-3:	Rule property sheet	18-4
Figure 18-4:	Dependencies dialog box for a rule	18-5
Figure 18-5:	Column property sheet: Rules and Defaults tab.....	18-7
Figure 18-6:	User-defined Datatype property sheet: Advanced Options tab	18-8
Figure 19-1:	Code editor for a trigger	19-3
Figure 19-2:	Trigger property sheet.....	19-4
Figure 19-3:	Dependencies dialog box for a trigger.....	19-5
Figure 20-1:	New User-defined Datatype dialog box.....	20-2
Figure 20-2:	User-defined Datatype property sheet	20-4
Figure 20-3:	Dependencies dialog box for a user-defined datatype.....	20-5
Figure 21-1:	View property sheet	21-3
Figure 21-2:	Code editor for a view	21-4
Figure 21-3:	Dependencies dialog box for a view.....	21-5
Figure 21-4:	View property sheet: Permissions tab	21-7
Figure 22-1:	Table editor toolbar.....	22-4
Figure 22-2:	Table editor	22-6
Figure 22-3:	Edit Check Constraints dialog box.....	22-8
Figure 22-4:	Edit Primary-key Columns dialog box.....	22-10
Figure 22-5:	Table property sheet.....	22-12
Figure 22-6:	Table Editor.....	22-13
Figure 22-7:	Dependencies dialog box for a table.....	22-15
Figure 22-8:	Table property sheet: Permissions tab	22-16
Figure 22-9:	Table property sheet: Miscellaneous tab	22-18
Figure 22-10:	New Check Constraint Properties dialog box.....	22-20
Figure 22-11:	Primary key - foreign key relationship.....	22-23
Figure 22-12:	Foreign Key property sheet.....	22-23
Figure 22-13:	Unique Constraint property sheet.....	22-25
Figure 22-14:	Overall Table Consistency - sample output	22-27
Figure 22-15:	Table Allocations - sample output.....	22-28
Figure 22-16:	Column Properties: General tab	22-30
Figure 22-17:	Column Properties: Rules and Defaults tab.....	22-31
Figure 22-18:	Column Properties: Check Constraints tab.....	22-33
Figure 23-1:	Proxy table mapping options.....	23-2
Figure 24-1:	Index property sheet	24-6
Figure 24-2:	Index property sheet: Miscellaneous tab.....	24-8

Figure 24-3:	Output from Full option of index consistency check	24-10
Figure 25-1:	Using an engine group.....	25-4
Figure 25-2:	How execution class and engine affinity affect execution.....	25-5
Figure 25-3:	Precedence rule	25-7
Figure 25-4:	Example of precedence	25-8
Figure 25-5:	Engine Group property sheet	25-10
Figure 25-6:	Engine Group property sheet: Engines tab.....	25-12
Figure 25-7:	Available Engines dialog box.....	25-13
Figure 25-8:	Execution Class property sheet	25-16
Figure 25-9:	Execution class property sheet: Bindings tab	25-17
Figure 25-10:	New Execution Class Bindings dialog box for logins	25-18
Figure 26-1:	Application Activity Monitor	26-9
Figure 26-2:	Cache Monitor.....	26-14
Figure 26-3:	Data Cache Monitor	26-17
Figure 26-4:	Device I/O Monitor.....	26-24
Figure 26-5:	Engine Activity Monitor	26-28
Figure 26-6:	Memory Utilization Monitor.....	26-30
Figure 26-7:	Network Activity Monitor	26-31
Figure 26-8:	Object Lock Status Monitor.....	26-37
Figure 26-9:	Object Page I/O Monitor.....	26-42
Figure 26-10:	Performance Summary Monitor.....	26-45
Figure 26-11:	A Sample Performance Trends Monitor.....	26-53
Figure 26-12:	Process Activity Monitor	26-54
Figure 26-13:	Process Current SQL Statement Monitor	26-60
Figure 26-14:	Stored Procedure Activity Monitor.....	26-64
Figure 26-15:	Transaction Activity Monitor.....	26-67
Figure 26-16:	Counting transactions and row inserts.....	26-69

List of Tables

Table 1:	Style conventions.....	xxxiv
Table 2:	Shortcut symbols	xxxiv
Table 3:	Displaying online help information.....	xxxvi
Table 2-1:	Valid drag-and-drop operations.....	2-9
Table 2-2:	Activities available from standard menus	2-12
Table 2-3:	Toolbar controls and key to Figure 2-6	2-15
Table 5-1:	Menus and commands on monitor windows	5-10
Table 6-1:	Inputs to Add New Server wizard.....	6-4
Table 7-1:	Inputs to Create Cache wizard	7-2
Table 8-1:	Inputs to Create a New Database Device wizard	8-4
Table 9-1:	Inputs to Create a New Database wizard	9-3
Table 10-1:	Inputs to Create a New Dump Device wizard.....	10-3
Table 10-2:	Inputs to Create Backup Command wizard	10-6
Table 10-3:	Inputs to Create a Restore Command wizard	10-9
Table 11-1:	Inputs to Create New Segment wizard	11-2
Table 12-1:	Object permission property sheet tabs	12-7
Table 12-2:	Granting access to create a table.....	12-8
Table 12-3:	Privilege and permission summary	12-9
Table 12-4:	Access granted through tabs and dialog boxes.....	12-13
Table 13-1:	Inputs to Create a New Login wizard	13-1
Table 14-1:	Inputs to Add a New User wizard.....	14-2
Table 14-2:	Input to Create Group wizard	14-12
Table 15-1:	Inputs for Create a New Role wizard	15-2
Table 17-1:	Inputs to Create a New Procedure wizard.....	17-2
Table 17-2:	Inputs to Create A New Extended Stored Procedure wizard	17-9
Table 19-1:	Inputs to Create a New Trigger wizard.....	19-2
Table 21-1:	Inputs to the Create View wizard.....	21-1
Table 22-1:	Types of table constraints	22-3
Table 22-2:	Table editor toolbar icons and key to Figure 22-1	22-4
Table 22-3:	Inputs to Create a New Foreign Key wizard	22-22
Table 22-4:	Inputs to Create New Segment wizard	22-24
Table 23-1:	Inputs to the Create a New Proxy Table wizard	23-4
Table 24-1:	Inputs to the Create a New Index wizard.....	24-2
Table 25-1:	Inputs to Create a New Engine Group wizard.....	25-9
Table 25-2:	Inputs to Create a New Execution Class wizard.....	25-15
Table 26-1:	Adaptive Server monitors	26-2
Table 26-2:	Tuning using the Application Activity Monitor	26-13
Table 26-3:	Tuning using the Cache Monitor.....	26-16

Table 26-4:	Tuning using the Data Cache Monitor	26-20
Table 26-5:	Tuning using the Device I/O Monitor	26-25
Table 26-6:	Tuning using the Engine Activity Monitor	26-29
Table 26-7:	Tuning using the Network Activity Monitor	26-34
Table 26-8:	Tuning using the Object Lock Status Monitor	26-40
Table 26-9:	Tuning using the Object Page I/O Monitor	26-44
Table 26-10:	Tuning using the Performance Summary Monitor	26-50
Table 26-11:	Tuning using the Process Activity Monitor	26-58
Table 26-12:	Tuning using the Process Current SQL Statement Monitor	26-63
Table 26-13:	Tuning using the Stored Procedure Activity Monitor	26-66
Table 26-14:	Tuning using the Transaction Activity Monitor	26-71
Table 27-1:	Suggested Statistics for Benchmarks	27-2
Table 27-2:	Describing the Benchmark Environment	27-4

About This Book

Managing and Monitoring Sybase Adaptive Server Enterprise describes how to manage Sybase® Adaptive Server™ Enterprise and monitor its activities using the Adaptive Server Enterprise plug-in for Sybase Central™. Sybase Central is a system administration tool that helps you manage Sybase Adaptive Server Enterprise and its databases. This manual provides step-by-step instructions for performing each of the activities supported by Sybase Central.

Unless otherwise noted, when this manual describes using Sybase Central to complete an Adaptive Server Enterprise management task, it means that you are using the features provided by the Adaptive Server Enterprise plug-in. For information about the Sybase Central architecture and plug-ins, refer to the online help.

Audience

This manual is for anyone who is authorized to perform administrative tasks on Adaptive Server Enterprise, or who owns one or more databases. Generally, System Administrators, Database Administrators, and application designers perform the activities described in this manual.

Familiarity with Adaptive Server Enterprise administration and some knowledge of the Sybase architecture are assumed. This manual also assumes that you are familiar with the Microsoft Windows environment you are using.

How to Use This Book

To get started, read:

- Chapter 1, “Introduction,” which gives an overview of Sybase Central.
- Chapter 2, “Sybase Central Windows and Dialog Boxes,” which introduces you to the Sybase Central interface and explains how to start and stop Sybase Central.
- Chapter 3, “Common Tasks,” describes generic procedures such as creating objects, deleting objects, displaying property sheets, and generating object DDL.

You might find it helpful to have Sybase Central running while you read, so you can try out the features.

To learn how to use Sybase Central to manage Adaptive Server Enterprise, read Chapters 4-27, in any order.

- Chapter 4, “Managing the Adaptive Server Operating Environment,” describes how to connect to, configure, and control an Adaptive Server Enterprise installation.
- Chapter 5, “Using Monitors,” describes how to configure and use monitor windows.
- Chapter 6, “Managing Remote Servers,” describes how to add and configure remote servers.
- Chapter 7, “Managing Cache,” describes how to create named caches and bind objects to caches.
- Chapter 8, “Managing Database Devices,” describes how to create and modify database devices and how to set up device mirroring.
- Chapter 9, “Managing Databases,” describes how to create and modify databases, how to set up and monitor the distribution of a database across one or more segments, and how to check database consistency.
- Chapter 10, “Backing Up and Recovering Databases,” describes how to create dump devices and how to use them to backup and restore databases.
- Chapter 11, “Managing Segments,” describes the use of segments to manage data storage.
- Chapter 12, “Controlling Access to Adaptive Server,” describes the use of Adaptive Server Enterprise logins and roles in controlling access to Adaptive Server Enterprise and its databases and objects.
- Chapter 13, “Managing Logins,” describes how to manage logins.
- Chapter 14, “Managing Users and Groups,” describes how to manage database users and database groups
- Chapter 15, “Managing Roles,” describes how to create user-defined roles and how to manage role assignments, mutual exclusivity, and role hierarchy.
- Chapter 16, “Managing Defaults,” describes how to manage defaults in a database.

- Chapter 17, “Managing Stored Procedures,” describes how to manage stored procedures and extended stored procedures.
- Chapter 18, “Managing Rules,” describes how to manage rules in a database.
- Chapter 19, “Managing Triggers,” describes how to manage triggers in a table.
- Chapter 20, “Managing User-defined Datatypes,” describes how to manage user-defined datatypes.
- Chapter 21, “Managing Views,” describes how to manage views on tables.
- Chapter 22, “Managing User Tables,” describes how to create and manage user tables and columns.
- Chapter 23, “Managing Proxy Tables,” describes how to manage tables created using Component Integration Services.
- Chapter 24, “Managing Indexes,” describes how to manage indexes on a table.
- Chapter 25, “Application Tuning,” describes execution classes and engine groups and their use in tuning performance of applications, logins, and stored procedures.
- Chapter 26, “Interpreting Monitor Data,” provides data item descriptions and tuning hints for each monitor window.
- Chapter 27, “Monitoring Techniques,” describes how to use monitors to create benchmarks and identify performance problems.

To look up reference information, see:

- Appendix A, “Adaptive Server Configuration Parameters,” gives a brief definition of each Adaptive Server Enterprise configuration parameter.
- The Glossary contains definitions of system administration and Sybase Central terms.

Release Compatibility

This manual describes features of the latest release of Adaptive Server. If you are managing SQL Server 11.0.x, objects and features described in this manual may not be visible in the Sybase Central object tree or on property sheets and dialog boxes.

Related Documents

The documentation sets for Adaptive Server Enterprise and Adaptive Server Enterprise Monitor provide detailed conceptual information about the tasks you can complete in Sybase Central. In particular, the following manuals are useful:

- *Adaptive Server Enterprise Utility Programs* documents the Sybase utility programs, such as isql, that you can execute from the operating system command line.
- *Sybase Troubleshooting Guide* contains information that helps you prevent or quickly respond to trouble with Adaptive Server Enterprise. It includes topics such as how to determine appropriate settings of configuration variables, how to manage transaction logs, how to develop good recovery procedures, and how to use disk mirroring effectively.
- *System Administration Guide* describes how to administer and control Adaptive Server Enterprise databases independent of system administration tools such as Sybase Central. It includes the Transact-SQL commands and system procedures used to perform the functions for which Sybase Central provides a graphical interface.
- *Adaptive Server Enterprise Reference Manual* provides basic syntax and usage information for every command, function, system procedure, and catalog stored procedure.
- *Transact-SQL User's Guide* describes database objects and the Transact-SQL commands you use to create triggers, views, and stored procedure.
- *Security Administration Guide* describes logins, users, groups, and permissions.
- *Adaptive Server Enterprise Monitor Server User's Guide* describes Adaptive Server Enterprise Monitor Server, which is required to use the monitoring features of Sybase Central.

Other Sources of Information

Use the SyBooks™ and SyBooks-on-the-Web online resources to learn more about your product:

- SyBooks documentation is on the CD that comes with your software. The DynaText browser, also included on the CD, allows you to access technical information about your product in an easy-to-use format.

Refer to *Installing SyBooks* in your documentation package for instructions on installing and starting SyBooks.

- SyBooks-on-the-Web is an HTML version of SyBooks that you can access using a standard Web browser.

To use SyBooks-on-the-Web, go to <http://www.sybase.com>, and choose Documentation.

Training

Sybase offers a system and database administration class called “Adaptive Server Enterprise Administration.” For details, contact:

Mail: Education Registrars
Sybase Professional Services
77 South Bedford Street
Burlington, MA 01803

Phone: (800) 8-SYBASE or (617) 564-6970

Fax: (800) 792-2733 or (617) 564-6960

E-mail: registrars@sybase.com

Conventions

The following sections describe the conventions used in this manual.

Naming Conventions

In the interest of brevity, the term Adaptive Server is used to refer to Adaptive Server Enterprise release 11.5.

Style Conventions

In this manual, the following typefaces and fonts have special significance:

Table 1: Style conventions

Example	Used For
procedure	Command keywords, terms that are being defined, and emphasized text
<i>exec.dtl</i>	Filenames, variable names, and book titles
check (pub_id in "1389")	Values that you type exactly as shown

Shortcuts

You can perform any activity in Sybase Central by executing a menu command. For most activities, there are shortcuts that help you work more efficiently. For example, Sybase Central has toolbars that let you select a button as an alternative to choosing the corresponding command from a menu.

In the descriptions of how to perform management tasks, this manual first presents the procedure for executing the appropriate menu command. If an alternative method exists for starting an operation, the manual describes it in a section following the basic procedure, titled "Shortcuts."

Shortcuts are presented in the following format:

Table 2: Shortcut symbols





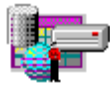
Graphic	Action
	Double-click an icon.
	Drag one of the icons shown and drop it on the other.
	Select the toolbar button shown.

Table 2: Shortcut symbols (continued)

Graphic	Action
	Execute a shortcut menu command. (Click the right mouse button.)
	This icon represents objects generically.

Privileges

Many activities in Adaptive Server Enterprise are restricted to users with specific roles or assignments, such as System Administrator, System Security Officer, System Operator, or database owner. When an activity has such restrictions, the manual indicates them at the beginning of the description. The following example illustrates a privilege description:

When you change ownership of a database, the new owner must have an Adaptive Server Enterprise login, but must not be a user of the database or have an alias in it.

Privileges	Only a System Administrator can transfer database ownership.
------------	--

► Note

Because a login with the System Administrator role can perform almost all Adaptive Server Enterprise activities, the “Privileges” notes include the System Administrator only when a System Administrator is the sole person who can perform the activity.

Mouse Buttons

To use Sybase Central, you need a pointing device such as a mouse. This manual refers to mouse buttons as the left and right buttons, assuming the mouse has two buttons and is configured for right-hand use.

Screen Images

Sybase Central runs on Windows NT 3.51 and 4.0, and Windows 95. In each environment, windows and dialog boxes have a somewhat different look, although their contents are identical. This manual displays screen images of Sybase Central running on Windows NT 4.0.

If You Need Help

Get help for managing Adaptive Server Enterprise from the printed documentation, online help, and the Technical Support Center.

Using Online Help

Online help is available. The following table shows how to access the various forms of help.

Table 3: Displaying online help information

To Display This Type of Online Help	Do This
Table of Contents for Sybase Central help. Gives access to all online help topics.	Choose Adaptive Server Enterprise Help from the Help menu on the menu bar of the main window. Select the Contents tab.
Help on utilities that come with Sybase Central.	Open the utility. Use the Help options in the utility.
Information on how to use the Windows help facility.	From the Start menu in Windows '95 or Windows NT 4.0, choose Help.
Search for a keyword.	Choose Adaptive Server Enterprise Help from the Help menu on the menu bar of the main window. Select the Index tab.
Search for text (full text search.)	Choose Adaptive Server Enterprise Help from the Help menu on the menu bar of the main window. Select the Find tab.
About dialog box for Sybase Central. Displays product version number.	Choose About Sybase Central from the Help menu on the menu bar of the main window.

Table 3: Displaying online help information (continued)

To Display This Type of Online Help	Do This
Brief description of a highlighted menu command or icon or of a pressed toolbar button.	See the status bar at bottom of the window.

Contacting Technical Support

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you cannot resolve a problem using the manuals or online help, please have the designated person contact Sybase Technical Support or the Sybase subsidiary in your area.

1

Introduction

About Sybase Central

Sybase Central is a graphical tool that provides a common interface for administering Sybase products. Products such as Adaptive Server Enterprise, Adaptive Server Anywhere, or NetImpact Dynamo include **plug-ins** that allow you to manage them through Sybase Central. This manual describes how to use Sybase Central to manage and monitor Adaptive Server Enterprise.

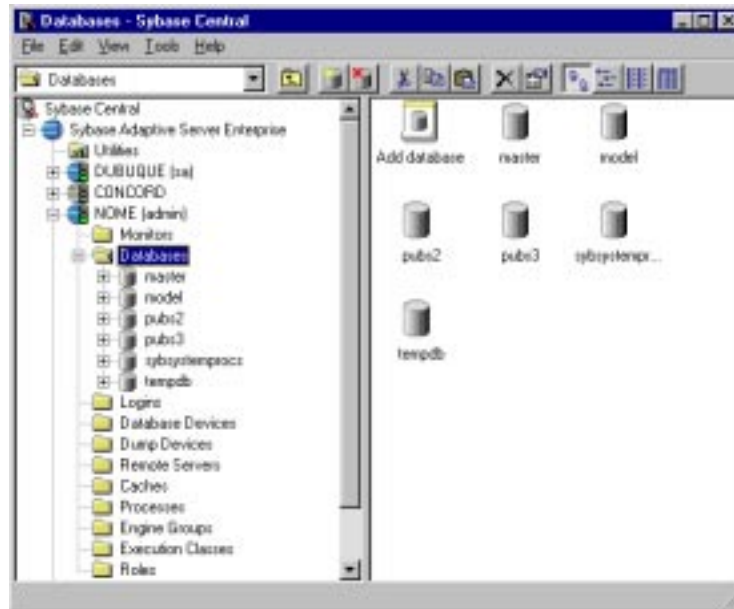


Figure 1-1: Sybase Central main window

Performing System and Database Administration Tasks

Using Sybase Central, you can manage most aspects of Adaptive Server, including the following:

- Adaptive Server configuration and processes
- Data caches
- Disk devices
- Databases and tables
- Access to Adaptive Server (logins, users, roles, and groups)
- Application queues (engine groups and execution classes)
- Component Integration Services connections

Monitoring the Performance of Adaptive Server

The Adaptive Server Enterprise plug-in is integrated with Adaptive Server Enterprise Monitor. You can monitor Adaptive Server performance from Sybase Central and, if necessary, make changes to your Adaptive Server configuration. Monitoring Adaptive Server is described in Chapter 5, “Using Monitors”, Chapter 26, “Interpreting Monitor Data”, and Chapter 27, “Monitoring Techniques.”

How Sybase Central Makes Administration Easier

Sybase Central helps you perform complex administration tasks without the need to remember the syntax of system stored procedures or Transact-SQL commands. Some special features that support this are:

- **Management of multiple servers from one console** – You can manage all the Adaptive Server installations listed in your interfaces file from the Sybase Central main window.
- **Generation of database definition language (DDL)** – You can generate DDL for the objects in Adaptive Server.
- **Visual representation of objects** – You can see the databases and logins in each Adaptive Server and the objects in each database. The display in the Sybase Central window expands and contracts to display information at the level of detail you need. (See Figure 1-1.)

- **Navigation between related objects** – To get more information about a database object related to the one whose property sheet you are displaying, you do not need to go back to the Sybase Central window to find and open the icon for the object. You can navigate directly through the displayed object's dialog box to the related object (Figure 1-2).



Figure 1-2: Opening a related property sheet

- **Toolbars** – As an alternative to executing commands from a menu, click a button in the toolbar (Figure 1-3) to perform such tasks as connecting to or disconnecting from Adaptive Server (A), deleting an object (B), or opening a property sheet (C).




Figure 1-3: Main window toolbar

2

Sybase Central Windows and Dialog Boxes

What's in This Chapter

This chapter explains how to start and stop Sybase Central and introduces the windows and dialog boxes of Sybase Central as they apply to managing Adaptive Server. For general information about monitor windows, see Chapter 5, “Using Monitors.”

Starting Sybase Central

Start Sybase Central using any of the standard methods for your operating system, such as:

- From the Start menu, choose Programs→Sybase→Sybase Central.
- In the Sybase program group or folder, double-click the Sybase Central icon.
- Create a shortcut on your desktop for Sybase Central.
- Double-click *scview.exe* in the *Sybttools\win32* directory.
- Add Sybase Central to your Startup program group.

Stopping Sybase Central

To stop Sybase Central, select the File menu and choose Exit.

The Sybase Central Window

The Sybase Central main window allows you to access Adaptive Server and Adaptive Server objects.

This section describes the Sybase Central main window and how to:

- Move through the Sybase Central hierarchy
- Customize the display
- Use drag-and-drop shortcuts

Sybase Central Features

Figure 2-1 illustrates the Sybase Central window and its main features.

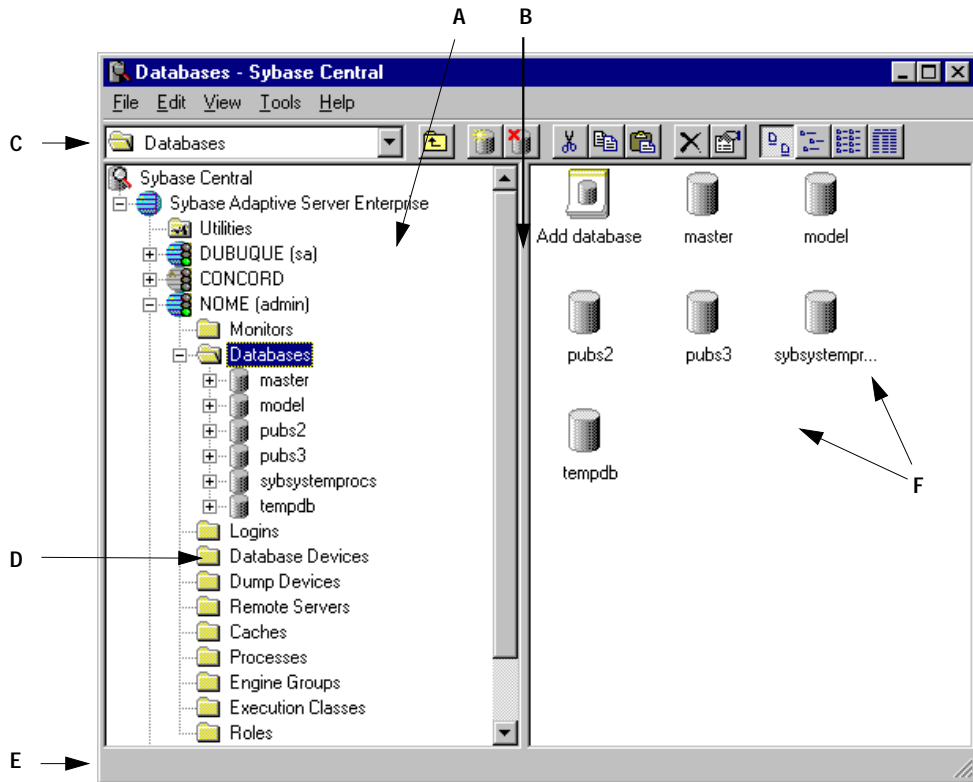


Figure 2-1: Sybase Central window

Key to Figure 2-1:	
A	Description
A	Object tree
B	Splitter bar
C	Toolbar
D	Folder icons
E	Status bar
F	Object icons

Window Panes

The window is split into left and right panes:

- The left pane contains a hierarchical list, or object tree, which shows:
 - The root level for the Adaptive Server plug-in
 - Adaptive Server installations and their objects
 - A folder that contains Adaptive Server utilities such as SQL Advantage
 - Management tools for other applications, such as Adaptive Server Anywhere
- The right pane displays the contents of the object selected in the left pane.

To adjust the size of the left and right panes, drag the splitter bar to the left or right with the mouse pointer.

Folder and Object Icons

The window includes folder icons and object icons:

- Each **folder icon** represents all objects of its type within Adaptive Server or a database. For example, the Databases folder in the Adaptive Server named NOME represents all databases in NOME. Folder icons can appear in either the left pane or the right pane of the window.
- Each **object icon** represents a single Adaptive Server or database object. Object icons, except databases, tables, and groups appear only in the right pane of the window. Icons for databases, tables, and groups are containers as well as objects and appear in the tree list in the left pane and in the object list in the right pane.
- Each monitor window icon represents a view of Adaptive Server.

Selecting Objects

For most activities, you must select an object before you do something to or with it. To select an object, click its icon.

Moving Through the Sybase Central Object Tree

To see different parts of the object tree, use the following techniques:

- To move vertically through the current display, use the scroll bar on the left or right pane.
- To expand or collapse the list to show different levels of detail, do one of the following:
 - Click plus or minus buttons. A plus button next to an icon indicates that the list of objects for that icon can be expanded. A minus button indicates that the list of objects for the icon is fully expanded.
 - Double-click a folder icon or its label.

Double-clicking a folder icon in the right pane also expands the list and changes the view in the right pane to a list of objects in the folder. For most objects, double-clicking an object icon in the right pane opens a property sheet that displays information about the object. Some objects open the code editor window when you double-click them.

Figure 2-2 illustrates the Adaptive Server object tree. The top level is the root of the Adaptive Server Enterprise plug-in.

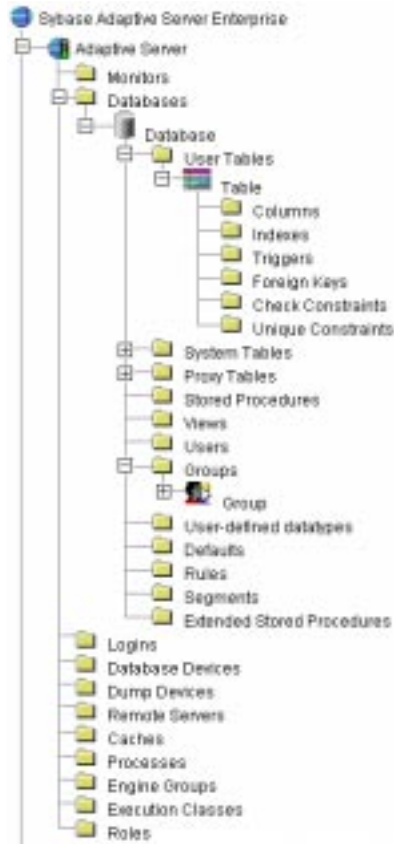


Figure 2-2: Adaptive Server object tree

Customizing the Display

To customize the Sybase Central window, you can:

- Choose the display format for the right pane
- Turn the toolbar and status bar on and off

Selecting Display Formats

To select the display format for the right pane, choose one of the following commands from the View menu:

- Large Icons – shows each object as a large icon with its label underneath.
- Small Icons – shows objects as small icons listed horizontally.
- List – shows objects as small icons listed vertically.
- Details – lists objects and displays detailed information about each one. The details vary with the type of object.

Figure 2-3 shows an example of each format.

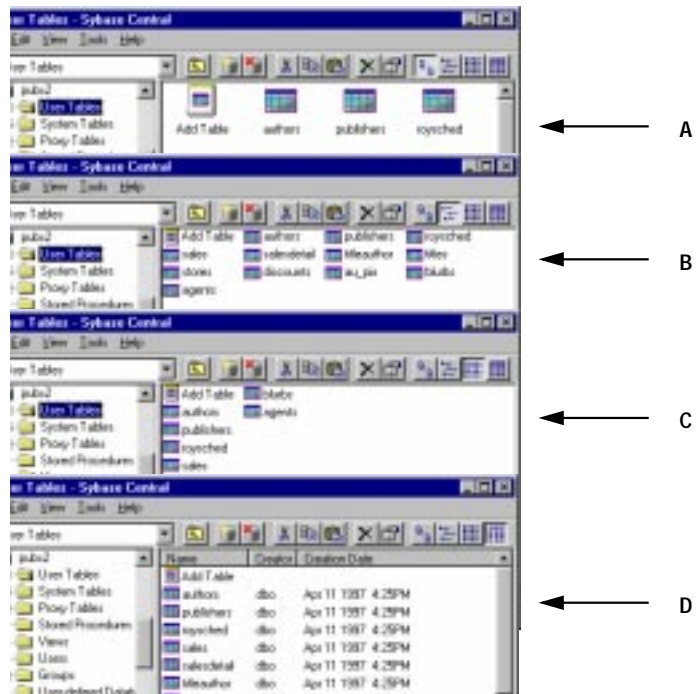


Figure 2-3: Object display formats

Key to Figure 2-3	Description
A	Large icons
B	Small icons
C	List
D	Details

Shortcuts



Choose the Large Icons, Small Icons, List, or Details button from the toolbar.



Click the right mouse button over an empty space in the right pane, and select Large Icon, Small Icon, List, or Details from the shortcut menu.

Turning the Toolbar and Status Bar On and Off

To hide the toolbar or status bar, choose Toolbar or Status Bar from the View menu. To redisplay the toolbar or status bar, repeat the procedure.

Using Drag-and-Drop Shortcuts

A quick way to perform operations involving the interaction of two Adaptive Server or database objects is to drag the icon for one object on top of an icon or folder for another object. When you drag and drop icons for supported operations, the application completes the operation if possible. Drag-and-drop is limited to situations where the source object is visible in either the left or right pane and the target object can be viewed (either simultaneously or by scrolling the window) in the other pane.

To drag and drop an icon:

1. Place the cursor on top of an icon.
2. Press and hold the left mouse button.
3. Continuing to hold the button down, move the mouse to drag the icon on top of another icon. As you drag the icon, a faint image of the icon travels across the screen, along with a symbol indicating that you cannot yet legally drop the icon. When the icon reaches its target location, the illegal drop symbol disappears, and you can release the mouse button (unless your intended target is illegal, in which case, you cannot complete the drag-and-drop operation).

► **Note**

To drag-and-drop an object within its folder (making a copy of it), hold down the Control key while you drag the object into empty space in the folder.

In Figure 2-4, the icon for the table named *roysched* is being dragged. The ghost image of the *roysched* has the illegal drop symbol on it because it is not in a place where it can do something.

4. Release the mouse button. Some tasks require confirmation, some do not.

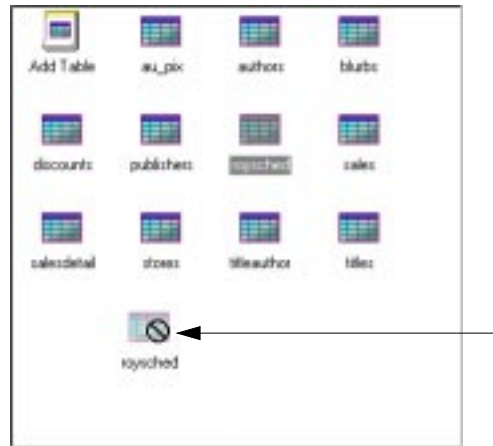


Figure 2-4: Drag-and-drop symbol

Valid Drag-and-Drop Operations

Table 2-1 summarizes the operations you can perform by dragging the icon in the middle column onto the icon in the right column.

Table 2-1: Valid drag-and-drop operations









Operation	Source	Target
Add a user (login) to a database (Users folder).		 User Folder
Add a user to a group.		
Grant a group all permissions on a procedure.		
Grant a group all permissions on a table.		

Table 2-1: Valid drag-and-drop operations (continued)



























Operation	Source	Target
Grant a group all permissions on a view.		
Grant a user all permissions on a table.		
Place a table onto a segment.		
Bind a database to a cache.		
Bind a table to a cache.		
Copy a table within a database or to another database		 Tables folder
Copy a table column within a table or to another table		 Columns folder
Copy a procedure within a database or to another database		 Stored procedures folder
Copy a default within a database or to another database		 Defaults folder

Table 2-1: Valid drag-and-drop operations (continued)

Operation	Source	Target
Copy a rule within a database or to another database		 Rules folder
Create a foreign key		 Foreign key table
Create a remote server		 Remote Servers folder
Connect Adaptive Server to SQL Advantage		 SQL Advantage

Using Keyboard Shortcuts in the Sybase Central Window

In addition to using a mouse, you can use keyboard shortcuts to choose menu commands and complete dialog boxes.

Every menu title and most commands on a menu have an underlined letter, called a mnemonic. To select a menu, press **ALT/mnemonic for the menu title**. To choose a menu command, press the mnemonic key for the command. Some commands can be run directly by pressing **Ctrl** plus another key or a function key. These shortcuts are listed on the menus.

To navigate to different boxes and buttons in a dialog box or property sheet, use the **Tab** key. To select different tabs in a property sheet, use the **Tab** key to move to the current tab, then use left and right arrow keys to select other tabs.

Menus and Toolbars

This section describes the menus and toolbar.

Standard Menus

The menu bar in the Sybase Central window contains the following menus:

- File
- Edit
- View
- Tools
- Help

The File menu is sensitive to the folder or container selected in the left pane. Menu commands change to meet the requirements of the object. When you highlight a menu command, a brief description of the command is displayed in the status bar at the bottom of the main window. Table 2-2 summarizes the menu commands.

Table 2-2: Activities available from standard menus

Menu	Activities
File	<ul style="list-style-type: none"> • Open a folder or object • Create a new object • Delete • Open the properties sheet for an object • Perform object specific activities such as generating DDL or running dbcc • Exit
Edit	<ul style="list-style-type: none"> • Cut, copy, or paste an object definition • Select a group of objects
View	<ul style="list-style-type: none"> • Hide or display the standard toolbar and status bar • Select the format for displaying object icons in the right pane of the Sybase Central window • Update the display with fresh data from Adaptive Server

Table 2-2: Activities available from standard menus (continued)

Menu	Activities
Tools	<ul style="list-style-type: none">• Connect to or disconnect from Adaptive Server• Create a connection profile• Manage plug-ins for Sybase Central
Help	<ul style="list-style-type: none">• Display Sybase Central online help• Display the About dialog box for Sybase Central

Shortcut Menus

Shortcut (pop-up) menus have the same commands as the File menu. To activate a shortcut menu, click the right mouse button over a folder icon or object icon representing an Adaptive Server or database object. A menu is displayed. From this menu, choose the appropriate menu command.

Toolbar and Status Bar

The Sybase Central toolbar gives you a quick alternative to executing commands from a menu. The status bar displays information about highlighted menu commands. Figure 2-5 illustrates the toolbar (A) and status bar (B).



Figure 2-5: Toolbar and status bar

Key to Figure 2-5:

A - toolbar

B - status bar showing message for Cut menu command

To display or hide the toolbar or status bar, toggle the Toolbar or Status Bar command on the View menu. For example, to display or hide the standard toolbar, choose Toolbar. When the toolbar command has a check mark (✓) next to it, the toolbar is visible in the window.

Toolbar

Figure 2-6 illustrates the main window toolbar. Table 2-3 describes the controls.



Figure 2-6: Sybase Central toolbar

Table 2-3: Toolbar controls and key to Figure 2-6

Key to Figure 2-6	Control	Function
A	Drop-down object tree	Displays the object tree of the currently selected object. Allows you to select a folder or object at a higher point on the tree without scrolling.
B	Up	Moves the Sybase Central display up one level in the object tree. For example, if the current focus is on a database, clicking this button moves the focus to the Adaptive Server in which the database is located.
C	Connect	Opens the New Connection dialog box so you can connect to another Adaptive Server or Sybase Central application.
D	Disconnect	Opens the Disconnect dialog box so you can choose an Adaptive Server installation to disconnect from.
E	Cut	Places the selected objects in the clipboard. If they are pasted elsewhere, the original objects are removed.
F	Copy	Copies the selected objects to the clipboard.
G	Paste	Pastes the contents of the clipboard into the selected container.
H	Properties	Opens the property sheet for the selected object.
I	Delete	Deletes the selected object.
J	Large Icons	Displays objects in the right pane of the Sybase Central window using large-format icons.
K	Small Icons	Displays objects in the right pane of the Sybase Central window using small-format icons.
L	List	Displays objects in the right pane of the Sybase Central window as a list of object names, along with small-format icons.
M	Details	Displays objects in the right pane of the Sybase Central window as a list of object names, along with details about each object. The information shown varies by object type.

To help identify the buttons in the toolbar, Sybase Central displays a **Tooltip**. To display a ToolTip, place the cursor over a toolbar button for a second. A floating label appears with a brief description of the button. It is not necessary to click the button.

Refreshing Window and Dialog Box Displays

As you work, information displayed in the Sybase Central window can get out of synchronization with the status of Adaptive Server. This is especially true if someone is updating Adaptive Server from another site. To update the contents of the main window, choose Refresh Folder or Refresh All from the View menu.

Using the Code Editor

The code editor is a separate window for displaying, editing, and printing the code of triggers, procedures, and views.

In addition to standard text-editing functions, it provides:

- Automatic syntax highlighting
- Language-sensitive indenting
- Unlimited undo and redo
- Drag-and-drop editing (you can drag selected text to a new location in the code)
- A toolbar and status bar

You can:

- Set the display font
- Find and replace text
- Jump to specific line numbers
- Open and save to external files
- Print the code
- Execute DDL code against a database

Opening the Code Editor

To open the code editor:

1. Select a stored procedure, a trigger, or a view.
2. From the File menu, choose Open.

The code editor opens with the code for the object displayed. The code editor also opens after you complete the wizard used to create one of these objects. Each menu item has status bar help. Figure 2-7 illustrates the code editor.

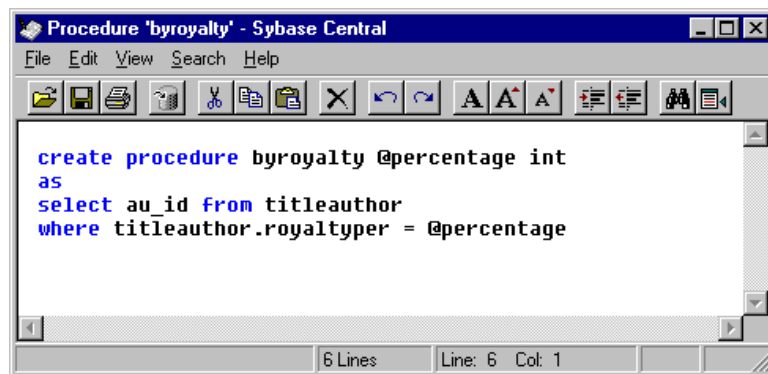



Figure 2-7: Code editor


► Note

If the object you open contains hidden (encrypted) text, when you open the code editor, it is empty.

Shortcuts



Double-click the object icon for triggers, views, or stored procedures.



Click the right mouse button over the object icon. Choose Open from the shortcut menu.

Syntax Highlighting

The code in the editor window is automatically highlighted as follows:

- Keywords and constants are blue
- Quoted strings are purple
- Partially quoted strings are red
- Comments are green

Finding and Replacing Code in the Code Editor

The code editor can find and replace text strings. The Find and Replace options are available on the Search menu.

Executing Code in the Code Editor

After you enter code in the editor, you can test it by choosing Execute Script from the File menu. If there are any errors in your SQL syntax, the editor alerts you and does not execute any of the code.

Executing code in the editor does not actually run the code against data - it only executes the `create <object>` or `alter <object>` command. For example, if you are creating a stored procedure, executing the code creates the stored procedure. It does not run the stored procedure against data.

Printing Code

To print the code, choose Print from the File menu. This opens the standard print dialog box.

The Code Editor Toolbar

The code editor has a toolbar that lets you perform tasks by clicking an icon rather than using a menu. Each icon has tooltip help for identification. Figure 2-8 illustrates the toolbar and identifies its components.



Figure 2-8: Code Editor toolbar

Key to Figure 2-8	Description
A	Open file
B	Save file
C	Print
D	Execute code
E	Cut
F	Copy
G	Paste
H	Clear (delete)
I	Undo
J	Redo
K	Change font
L	Larger font
M	Smaller font
N	Indent line
O	Move line left
P	Find
Q	Line number

Using the Sybase Tools Console

The Sybase Tools Console is a text editor window that Sybase Central uses to report the results of checking database consistency, generating DDL, verifying the accuracy of proxy tables, and providing other output from Adaptive Server.

You can save output in the Tools Console to a file and can open text files in the Tools Console.

The Tools Console has a set of menus and a toolbar that are similar to those in the Code Editor. If you have questions about the Tools Console, refer to descriptions of Code Editor components in this chapter, or view online Help.

3

Common Tasks

What's in This Chapter

This chapter describes how to do the following tasks common to most objects:

- Creating an object
- Viewing the properties of an object
- Changing the properties of an object
- Navigating to the properties of a related object
- Deleting an object
- Generating DDL scripts

Creating an Object

You create most objects using wizards. Wizards are a special form of dialog box that walks you through a procedure. Each page of a wizard asks you to enter information or make a choice. The wizard provides information to help you make a decision and prevents you from selecting mutually exclusive or incorrect options. Some tasks, such as creating a table, use special editor windows or dialog boxes.

To create an object:

1. Select the folder for the type of object you want to create.
2. From the File menu, choose New; then from the cascading menu, choose the object name. The create object wizard opens.

Shortcuts



In the right pane, double-click the wizard icon.



Place the cursor in an empty area in the right pane. Click the right mouse button. From the popup menu, choose New. From the cascading menu, choose the object name.


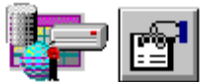
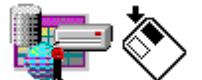
Viewing the Properties of an Object

To find out the properties of an object, you open a special dialog box called a property sheet. A property sheet contains information about the object and about how it relates to other objects in Adaptive Server or a database. The property sheet also provides a direct navigation path to its related objects.

To open the property sheet for an object:

1. Select the icon of the object you want to view.
2. From the File menu, choose Properties. The property sheet opens.

Shortcuts

	Double-click the object icon. (For tables, views, triggers, and stored procedures, hold down the ALT key while double-clicking the icon.)
	Select the object icon. Then, select the Properties toolbar button.
	Click the right mouse button over the object icon. Choose Properties from the shortcut menu.

Using Property Sheet Tabs

Information about Adaptive Server and database objects appears on multiple panels within property sheets. These panels are called tabs, referring to the row of tab-shaped controls across the top of the property sheet.



Figure 3-1: Property sheet tabs

The tabs vary according to the function of each object. To display a different tab, select that tab control.

Changing the Properties of an Object

Most objects have properties that can be changed. To change the properties of an object:

1. Open the property sheet for the object.
2. Select the tab with the property you want to change.
3. Change the property.
4. Click Apply to send the change to Adaptive Server and keep the property sheet open. Click OK to send the change to Adaptive Server and close the property sheet.

Refer to online help or this manual for instructions about changing specific properties.

Closing a Property Sheet

To close a property sheet and save any changes you have made, click OK. To close a property sheet without saving unapplied changes, click Cancel.

Navigating to the Properties of a Related Object

In many property sheets, you must select objects to create relationships to the object you are viewing. For example, in the Database property sheet, you can select a database device on which you want to allocate space for the database. In the Login property sheet, you can select users to map to the login.

In most cases, when you select an object in a list, you can click a Properties button that opens the property sheet for the selected object. This feature allows you to examine the object to be sure you want to create the relationship.




Figure 3-2: Properties button on a property sheet

Deleting an Object

To delete an object:

1. Select the icon of the object you want to delete.
2. From the File menu, choose Delete.
3. Confirm the deletion in the confirmation dialog box.

Shortcuts

	<p>Select the icon of the object you want to delete. Select the Delete toolbar icon.</p>
	<p>Click the right mouse button over the object you want to delete. Choose Delete from the shortcut menu.</p>

Getting Help

Each property sheet has a help button. Clicking the Help button opens up online help to a topic that generally describes the features of the property sheet and has links to related topics. For help with a specific control on a property sheet tab, use What's This? help. To get What's This? help, do one of the following:

1. Click the question mark button in the upper-right corner of the dialog box, then click the help cursor on the control.
2. Right-click the control, then click What's This? in the pop-up menu.
3. Make the control active (for example, by tabbing to it), then press F1.

This opens a pop-up window that describes the control.

You can also start help from the Help menu on the main window. Choose Adaptive Server Enterprise Help. It displays a table of contents. For more information about how to use online help, see the online help for your operating system.

Generating and Executing Adaptive Server DDL Scripts

Data Definition Language (DDL) consists of Transact-SQL commands combined into scripts that capture the schema definition of an Adaptive Server object. You can run DDL scripts to recreate an Adaptive Server configuration or database in its original location or on a different Adaptive Server installation.


Sybase Central allows you to generate DDL scripts. The DDL scripts are displayed in the Tools Console. You can save them to files for later editing or use.

Generating DDL Scripts

To generate DDL:

1. Open the folder of the object for which you want to generate DDL.
2. Select the objects for which you want to generate DDL.
3. From the File menu, choose Generate DDL. Sybase Central generates the DDL and displays it in the Tools Console. You can save the DDL to a file.

Shortcut



Click the right mouse button over the object icon.
Choose Generate DDL from the shortcut menu.

Executing DDL Scripts

To execute DDL scripts, choose one of the following options:

- Execute the DDL script from a utility such as SQL Advantage.
- Execute the DDL script from an isql session.

4

Managing the Adaptive Server Operating Environment

What's in This Chapter

This chapter describes the following Adaptive Server administration tasks:

- Connecting to and disconnecting from Adaptive Server
- Displaying information about Adaptive Server and the objects it controls
- Setting Adaptive Server configuration parameters
- Displaying user processes and killing processes
- Managing remote servers

Connecting to and Disconnecting from Adaptive Server

Before you can perform any Adaptive Server administration activities, you must be connected to the Adaptive Server you want to administer. You can connect to any Adaptive Server identified in the *sql.ini* file or through Directory Services, and you can maintain multiple Adaptive Server connections simultaneously.

Identifying Servers to Sybase Central

In order to manage Adaptive Server, Sybase Central must know the host name, port, and method of connection. This information is stored either in a *sql.ini* file or in Directory Services. The *sql.ini* file is the equivalent of the *interfaces* file used with Adaptive Server on UNIX platforms. You can update the *sql.ini* file using a text editor. You maintain Directory Services information with the Directory Services Editor (DSEDIT). For more information about the *sql.ini* file and DSEDIT see *Adaptive Server Enterprise Utility Programs for Windows and Windows NT*.

► **Note**

If you update the *sql.ini* file or Directory Services while Sybase Central is running, you must restart Sybase Central for the changes to take effect.

Connecting to Adaptive Server

When you start Sybase Central, the main window opens, displaying a list of icons representing the servers you have configured (Figure 4-1.)

► **Note**

Sybase Central displays an icon for each server regardless of its type, for example, Adaptive Server, Monitor Server, and OmniConnect. You can only connect to Adaptive Server or OmniConnect.



Figure 4-1: Sybase Central window

To connect to Adaptive Server:

1. From the Tools menu, choose Connect. The Logon dialog box opens.



Figure 4-2: Logon dialog box

2. In the User Id box, enter the login name you want to use to connect to Adaptive Server.
3. In the Password box enter the password for the login.

4. Select an Adaptive Server from the Server Name drop-down list.
5. Click OK. If the connection is successful, the icon changes to blue to indicate a successful connection. You can now manage that Adaptive Server Enterprise installation.

Shortcuts



Select the icon for the Adaptive Server you want to connect to. The server name is filled in and grayed out. You only need to enter the login and password.



Click the plus button for the Adaptive Server icon. The server name is filled in and grayed out. You only need to enter the login and password.



Click the Connect button in the toolbar.

Using Component Integration Services Passthrough Connections

If you are using Component Integration Services, you can specify that a login automatically connects to a remote server when it logs into a local server. If you have automatic passthrough specified, when you log into a server through Sybase Central, actions taken on the local server, such as creating objects, get sent to the remote server. There is no specific visible indication of this status in Sybase Central. You may notice that the objects you are trying to create on the local server are not being created. Therefore, if a login is set up with passthrough mode, that person must understand the ramifications of that status for using Sybase Central.

Using Connection Profiles

You can speed up the process of connecting to Adaptive Server by using connection profiles. A connection profile records connection information. You can use a connection profile to connect to Adaptive Server manually after Sybase Central opens or automatically when you start Sybase Central.

Creating a Connection Profile

To create a connection profile:

1. From the Tools menu, choose Connection Profiles. The Connection Profiles dialog box opens.

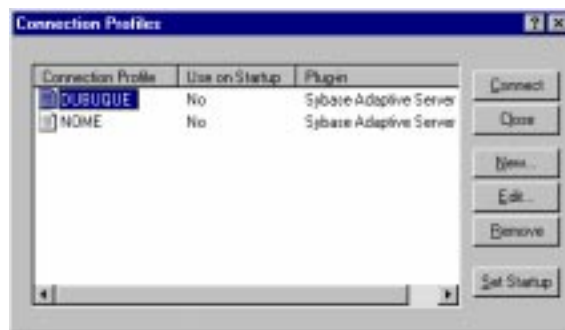


Figure 4-3: Connection Profiles dialog box

2. Click New.
3. In the Create New Profile dialog box, type a name for the profile.



Figure 4-4: Create New Profile dialog box

4. Click OK. The Logon dialog box opens.
5. Enter the User ID (login name) and password you want to use to connect to Adaptive Server.
6. Select the Adaptive Server installation you want to connect to from the drop-down list.
7. Click OK.

Connecting to Adaptive Server When Sybase Central Starts

You can configure Sybase Central to have connection profiles connect to Adaptive Server installations automatically when you start Sybase Central.

To connect at startup:

1. Open the Connection Profiles dialog box as described in “Creating a Connection Profile” on page 4-4.
2. Select the profile that you want to run when Sybase Central starts.
3. Click Set Startup.

Using Connection Profiles to Connect Manually

You can use a connection profile to connect to Adaptive Server instead of completing the Adaptive Server Enterprise Logon dialog box.

To use a connection profile manually:

1. Open the Connection Profiles dialog box as described in “Creating a Connection Profile” on page 4-4.
2. Select the connection profile for the Adaptive Server installation that you want to connect to.

3. Click Connect.

Editing a Connection Profile

To change the password or login in a connection profile:

1. Open the Connection Profiles dialog box as described in “Creating a Connection Profile” on page 4-4.
2. Select the connection profile that you want to change.
3. Click Edit.
4. In the Adaptive Server Enterprise Logon dialog box, change the entries.

Diagnosing an Unsuccessful Connection

If the connection is not successful, try these techniques:

- Check your login name and password to be sure they are correct.
- Try to ping Adaptive Server to see if it is running. Use any network utility that has this feature, for example DSEDIT.
- Make sure that the entries in Directory Services or the *sql.ini* file are correct.

Disconnecting from Adaptive Server

When you disconnect from Adaptive Server, the Adaptive Server icon changes to gray and you no longer have access to details about the objects in Adaptive Server.

1. Select the icon of the Adaptive Server from which you want to disconnect.
2. From the File menu, choose Disconnect.

Shortcuts



Right click the icon for the Adaptive Server you want to disconnect from. From the shortcut menu, choose Disconnect.



Select the Adaptive Server you want to disconnect from. Click the Disconnect button in the toolbar.

Starting and Stopping Adaptive Server

Sybase Central can start, stop, and pause Adaptive Server on a Windows NT machine on which they are both installed. It cannot start Adaptive Server on other machines. To restart on another machine, you must run the startserver utility on the host machine. For information on startserver, see the documentation for your platform.

Starting Adaptive Server

To start Adaptive Server:

1. In the left pane, click the icon for Adaptive Server. Sybase Central displays a message asking you if you want to start Adaptive Server.
2. Click Yes.

Shortcuts



Right click the icon for Adaptive Server. From the shortcut menu, choose Start.

Pausing Adaptive Server

Pausing Adaptive Server prevents users from connecting to Adaptive Server. It does not affect your connection. This allows you to manage Adaptive Server while preventing other server activity.

To pause Adaptive Server:

1. Select the icon for Adaptive Server.
2. From the File menu, choose Pause. The traffic light changes to yellow.

Shortcuts



Right click the icon for Adaptive Server. From the shortcut menu, choose Pause.

Resuming Operation of A Paused Adaptive Server

To resume normal operation of a paused Adaptive Server:

1. Select the icon for Adaptive Server.
2. From the File menu, choose Continue. The traffic light changes to green.

Shortcuts



Right click the icon for Adaptive Server. From the shortcut menu, choose Continue.

Shutting Down Adaptive Server

You can shut down an Adaptive Server that is running on the same machine as Sybase Central is running on. To shut down Adaptive Server:

1. Select the icon of Adaptive Server.
2. From the File menu, choose Stop.

When you shut down Adaptive Server, its icon changes to a red light.

Shortcuts



Right click the icon for Adaptive Server. From the shortcut menu, choose Stop.

Displaying Adaptive Server Properties

To display the properties of Adaptive Server:

1. Select the Adaptive Server icon.
2. From the File menu, choose Properties.



Figure 4-5: Server property sheet

The General tab lists information about Adaptive Server. The Defaults tab shows the default language and character set. The Monitor Server tab allows you to create a relationship with Monitor Server. For more information about this task, see Chapter 5, “Using Monitors.”

Shortcuts



Select the Adaptive Server icon. Click the Properties button in the toolbar.



Click the right mouse button over the Adaptive Server icon. Choose Properties from the shortcut menu.

Configuring Adaptive Server

When you install Adaptive Server, default configuration values take effect. Logins with the System Administrator role can reset most configuration parameters. Logins with the System Security Officer role can also reconfigure some of the parameters.

This section describes how to set configuration parameters.

Appendix A, “Adaptive Server Configuration Parameters” gives a short description of each configuration parameter. For more information, see *System Administration Guide* for the appropriate Adaptive Server release. For a discussion of configuration issues to consider when determining optimal settings, see *Performance and Tuning Guide*.

Who Can Set Configuration Parameters

The following rules govern who can set configuration parameters:

- Logins assigned the System Security Officer role can reset the following configuration parameters:
 - allow updates
 - audit queue size
- The default character set id parameter is set during Adaptive Server installation and cannot be reset from within Sybase Central.
- Logins assigned the System Administration role can reset all other parameters.

Parameters Requiring Restart

Some parameter values take effect as soon as you reset the value. Others do not change until you reset the value and then reboot Adaptive Server. In the Server Configuration dialog box, parameters requiring an Adaptive Server reboot are marked “Yes” in the Restart Required field.

Setting Configuration Parameters

Configuration parameters are organized by functional group, according to the areas of Adaptive Server behavior that they affect. When displaying parameters, you can select the functional group to display.

To set configuration parameters for Adaptive Server:

1. Select the Adaptive Server you want to configure.
2. From the File menu, choose Configure. The Server Configuration dialog box opens.



Figure 4-6: Server Configuration dialog box

3. In the Show Configuration Parameters For box, select the functional group to display, or select "All."
4. In the parameter list, select the parameter you want to update. For a brief description of the selected parameter, read the Explanation box.
5. Enter the new value in the Value column of the table.
6. Repeat steps 4 and 5 for each parameter you want to update.

7. Click OK. Sybase Central updates the configuration values in the dialog box as follows:
 - If the parameter you reset takes effect immediately, it is listed in the Value column.
 - If the parameter you reset requires you to reboot Adaptive Server, it is listed in the Pending Value column.

Managing Adaptive Server Processes

You can use Sybase Central to:

- View the status of user processes in Adaptive Server
- Set execution attributes of processes
- Kill processes

Displaying the Properties of a Process

To display user processes on Adaptive Server, select the Processes folder. Each process is represented by an icon. To display the properties of a process:

1. Select the process whose properties you want to display.
2. From the File menu, choose Properties.



Figure 4-7: Process property sheet

Shortcuts



Double-click the process icon.



Select the process icon. Click the Properties button in the toolbar.



Click the right mouse button over the process icon. Choose Properties from the shortcut menu.

Process Parameters

The Parameters tab of the Process property sheet lists additional information about the process.

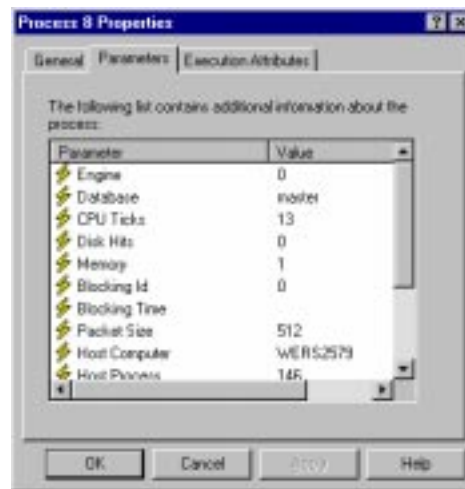


Figure 4-8: Process property sheet: Parameters tab

Setting Execution Attributes for A Process

You can change the priority of a process and you can assign it to an engine group. For a complete discussion of execution attributes and how they are used to tune performance, see Chapter 25, “Application Tuning.” To set execution attributes:

1. Open the property sheet for the process as described in “Displaying the Properties of a Process” on page 4-12.
2. Select the Execution Attributes tab.



Figure 4-9: Process property sheet: Execution Attributes tab

3. Change the priority or engine group, or both.
4. Click Apply or OK.

Returning to the Default Execution Attributes

To return to the default execution attributes:

1. Open the property sheet for the process as described in “Viewing the Properties of an Object” on page 3-2.
2. Select the Execution Attributes tab.
3. Check the Set System Attributes To Default Values box.
4. Click Apply or OK.

Killing a Process

To kill a process:

1. Select the process.
2. From the File menu, choose Delete.

Shortcuts



Select the icon of the process you want to kill. Then, select the Delete toolbar icon.



Click the right mouse button over the process you want to kill. Then, choose Delete from the shortcut menu.

5

Using Monitors

What's in This Chapter

This chapter describes how to configure and use the monitors in Sybase Central. The topics in this chapter are:

- Introduction to Adaptive Server Enterprise Monitor
- Configuring monitors in Sybase Central
- Opening and closing monitors
- Monitor window descriptions
- Setting up a monitor session
- Setting up the Performance Trends monitor
- Troubleshooting

Chapter 26, “Interpreting Monitor Data” provides data descriptions and tuning hints for each monitor. Chapter 27, “Monitoring Techniques” describes how to use the monitors to create benchmarks and solve specific problems.

Introduction to Adaptive Server Enterprise Monitor

The Adaptive Server Enterprise monitoring capability is collectively known as Adaptive Server Enterprise Monitor. It consists of four components that gather or display Adaptive Server Enterprise performance data:

- **Monitor Server** – This server collects Adaptive Server Enterprise performance data in real time and makes the data available to the other Adaptive Server Enterprise Monitor components. Monitor Server is a Sybase Open Server™ application.
- **Historical Server** – This server obtains Adaptive Server Enterprise performance data from Monitor Server and saves the data in files for deferred analysis. Historical Server is a Sybase Open Server application.
- **Monitors in the Adaptive Server Enterprise Plug-in for Sybase Central (Monitor Viewer)** – The monitors obtain Adaptive Server performance data from Monitor Server and display the data in real time in tables and graphs.

- **Monitor Client Library** – This application programming interface (API) to Monitor Server is available to users for developing monitoring applications. Monitor Viewer and Historical Server are Monitor Client Library applications.

Adaptive Server Enterprise Monitor Architecture

The following figure shows the relationships between Adaptive Server Enterprise and the various components of Adaptive Server Enterprise Monitor.

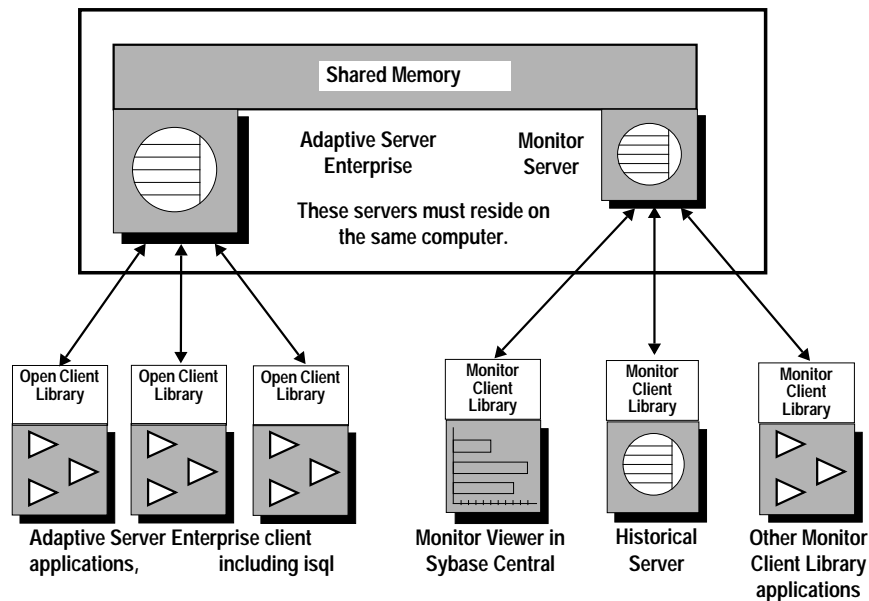


Figure 5-1: Adaptive Server Enterprise Monitor architecture

Adaptive Server Enterprise saves performance data in a shared memory area that Monitor Server reads. Because of this shared memory technique, Monitor Server must be installed and running on the same machine as the Adaptive Server Enterprise installation being monitored. A one-to-one relationship exists between an Adaptive Server Enterprise installation and a Monitor Server installation. For more information about Monitor Server, see *Sybase Adaptive Server Enterprise Monitor Server User's Guide*.

Monitor Client Library applications obtain Adaptive Server Enterprise performance statistics from Monitor Server. These applications are clients of Monitor Server. For performance reasons, Sybase recommends running Monitor Client Library applications on machines other than the ones where Adaptive Server Enterprise/Monitor Server pairs are running. See *Sybase Adaptive Server Enterprise Monitor Client Library Programmer's Guide* for more information.

Monitor Viewer in the Adaptive Server Enterprise Plug-in for Sybase Central includes a set of monitors showing different aspects of Adaptive Server resource usage at various levels of detail. Each open monitor is a separate application, with a unique client connection to Monitor Server. In Sybase Central, each Adaptive Server installation has its own Monitors folder containing the set of monitor objects. This chapter describes how to configure and use the monitors.

Historical Server collects performance information from Monitor Server and saves the information in files for deferred analysis. Two Historical Server interfaces let users specify the data to collect and the time period desired. They also include a historical data playback feature. The interfaces are:

- A command interface in `isql`. For more information, see *Sybase Adaptive Server Enterprise Monitor Historical Server User's Guide*.
- A programming interface using Monitor Client Library. For more information, see *Sybase Adaptive Server Enterprise Monitor Client Library Programmer's Guide* and *Sybase Adaptive Server Enterprise Monitor Historical Server User's Guide*.

Configuration Procedures

Before you can use a monitor:

- A Monitor Server must be configured on the same machine as the Adaptive Server you want to monitor. See *Sybase Adaptive Server Enterprise Monitor Server User's Guide* for Monitor Server configuration instructions.
- Monitor Server must be running. See "Starting Monitor Server" on page 5-4.
- An entry for the configured Monitor Server must exist in the `sql.ini` or Directory Services file on the machine where Sybase Central is running. See *Open Client/Server Configuration Guide for Desktop Platforms* for instructions on using `dsedit` to make these changes.

- The correct release of a Java virtual machine must be installed on your machine. For instructions, see the *Adaptive Server Enterprise Installation Guide* and the latest *Release Bulletin for Adaptive Server Enterprise*.
- In the Sybase Central tree, a Monitors folder must appear under the Adaptive Server icon. See “Configuring a Monitors Folder” on page 5-5.
- Your Adaptive Server login must have appropriate privileges.

Privileges	<p>Connect to Adaptive Server using either:</p> <ul style="list-style-type: none"> • a login with sa_role <p>or</p> <ul style="list-style-type: none"> • a login associated with a user in the <i>master</i> database that has permission to execute the <i>mon_rpc_connect</i> stored procedure. For this option to work, non-sa_role monitoring must be enabled. <p>For more information, see “Privileges Required for Monitoring” on page 5-5.</p>
------------	---

Starting Monitor Server

If Monitor Server is installed on the same machine as the Sybase Central installation, you can start it from the Sybase Central tree by right-clicking its icon and choosing the Start command from the pop-up menu. Monitor Server reads startup parameters from the NT Registry. The Monitor Server icon indicates one of the following server states:

- Green indicates running
- Yellow indicates paused
- Red indicates stopped

You might need to refresh the Sybase Central window to get an up-to-date status indication.

If Monitor Server is not installed on the Sybase Central machine, you do not have control over it from Sybase Central, Sybase Central does not know whether it is running, and the Monitor Server icon does not indicate the server state. See *Adaptive Server Enterprise Monitor Server User's Guide* for instructions on starting Monitor Server.

Configuring a Monitors Folder

To configure a Monitors folder:

1. In the Sybase Central tree, highlight the icon for the Adaptive Server you want to monitor.
2. From the right mouse menu, select Properties.
3. Select the Monitor Server tab.
4. In the drop-down box, select the Monitor Server that was configured for this Adaptive Server.

Be sure to select the correct Monitor Server name. If you choose the wrong Monitor Server, the monitors return data for the wrong Adaptive Server.

The list of Monitor Server names in the drop-down box on the Monitor Server tab lists the servers in the *sql.ini* or other Directory Service file on the Sybase Central machine. If the correct Monitor Server is not in the list, use the **dsedit** utility to add Monitor Server to the *sql.ini* or other Directory Service file.

Each Adaptive Server installation has a unique Monitor Server configured for it. If you do not know the correct Monitor Server name, refer to the Directory Service, *sql.ini* (Windows NT) or *interfaces* (UNIX) file on the Adaptive Server host machine.

5. Click Apply or OK.
6. These actions do not propagate to other Sybase Central installations. If you later want to monitor the same Adaptive Server from a Sybase Central installation on a different machine, you need to configure the Monitors folder in the other Sybase Central installation.

Privileges Required for Monitoring

To monitor Adaptive Server, you must connect to it using either:

- A login that has *sa_role*, or
- A login associated with a user in the *master* database that has permission to execute the *mon_rpc_connect* stored procedure. (The stored procedure is in the *master* database.) For this option to work, *non-sa_role* monitoring must be enabled.

To associate a login with a user, the system administrator can:

- Add a user to *master* for each login that wants to run monitors, or

- Add one user to *master* and associate multiple logins to it through aliases.

To add a user to *master*, see “Creating Users” on page 14-2. To set up login aliases, use the User property sheet.

To enable non-sa_role monitoring, the system administrator should:

1. Execute the **mon_authorize_non_sa** stored procedure in the Adaptive Server to be monitored. This stored procedure is in the *master* database.
2. If the **mon_authorize_non_sa** stored procedure is not in the *master* database or if it issues an error message indicating that two-phase commit needs to be installed, some steps were omitted during the Adaptive Server and Monitor Server installation process. See the latest *Release Bulletin for Adaptive Server Monitor* for specific instructions.

Opening and Closing Monitors

This section describes how to open and close monitors.

Opening Monitors

To open a monitor from the Sybase Central main window:

1. Click the Monitors folder under the Adaptive Server you want to monitor. If a Monitors folder is not in the tree, see “Configuring a Monitors Folder” on page 5-5.
2. In the right pane, double click the monitor you want to open.

To open a monitor from another monitor:

1. From the File menu on a monitor window, choose Monitors.
2. In the cascading menu, choose the monitor you want to open. The new monitor gathers statistics for the same Adaptive Server as the original monitor.

To open a Process Current SQL Statement monitor:

1. Open a Process Activity monitor.
2. Click on the process whose executing SQL statement you want to see.
3. From the File menu, choose Monitors. From the cascading menu, choose Process Current SQL Statement.

Opening a monitor establishes a connection to Monitor Server and a connection to Adaptive Server Enterprise. Each open monitor has a unique connection to each of these servers.

If the connection was not successful, see “Troubleshooting” on page 5-27.

Closing Monitors

To close a monitor:

1. Select the File menu on the monitor window
2. Choose Close.

Closing a monitor closes the connections to Monitor Server and Adaptive Server for the current monitor. It ends the current session.

If you reopen the same monitor, you are establishing a new session. All menu settings are reset to their default values and all activity counters and session values are reset to zero.

Monitor Window Description

Figure 5-2 shows a typical monitor window.

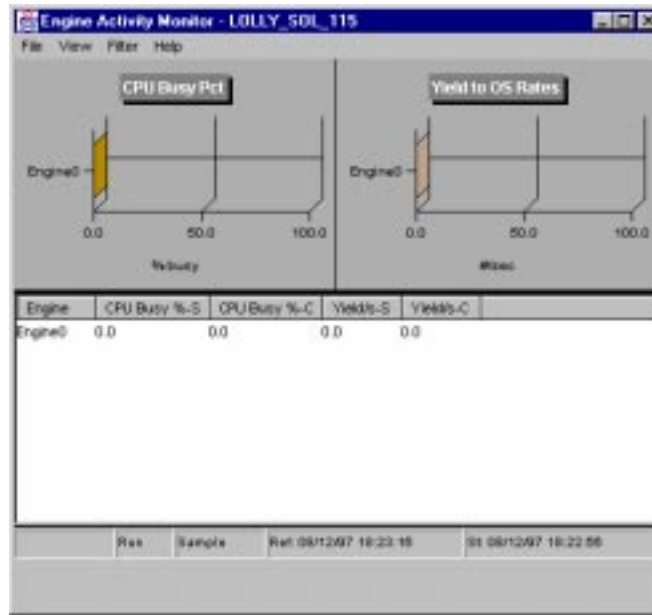


Figure 5-2: Sample monitor window

The title bar shows the name of the Adaptive Server being monitored. If this name is different from the Adaptive Server you selected in Sybase Central, see “Troubleshooting” on page 5-27.

The menu bar and menu commands are context sensitive based on the current monitor. Menus and commands that do not apply to the current monitor do not show up.

The monitor display contains scroll bars whenever they are needed to display the entire range of data.

You can resize a monitor window by grabbing the window edges. The fonts adjust accordingly (larger or smaller) to fit the new window size.

You can adjust column widths in monitor table displays by dragging the column header to the desired size.

Contents of Monitor Status Bar

The status bar at the bottom of a monitor shows the following information:

- Status of the current monitor.
 - Initializing – the monitor is in the process of connecting to Monitor Server and Adaptive Server, and has yet to collect the first sample.
 - Running – the monitor refreshes the data after each sample interval.
 - Paused – the data is fixed and will not be refreshed until you reverse the pause. During a pause, data continues to be collected in the background, maintaining accurate cumulative data.
 - Done – the monitor is finished gathering data. This status applies only to monitors that collect a one-time snapshot of data, such as the Process Current SQL Statement Monitor.
 - Error – an error condition is preventing the monitor from continuing to collect more data.
- Statistics setting – current setting of the Statistics option (sample, cumulative, or both).
- Refresh time – time the last sample was obtained by Monitor Server.
- Session start time – time this monitor was opened.

Showing and Hiding the Status Bar

The Show Status Bar command on the View menu makes the status bar visible at the bottom of the monitor window.

The Hide Status Bar command on the View menu makes the status bar disappear.

Monitor Menus

Table 5-1 lists the commands on the monitor menus.

Table 5-1: Menus and commands on monitor windows

Menu	Commands
File	<ul style="list-style-type: none"> Monitors - open another monitor Close - close the monitoring window
View	<ul style="list-style-type: none"> Statistics – change the statistic types displayed in the monitor tables (sample data, cumulative data, or both, where applicable) Sample Interval - change the sample interval Pause – stop refreshing the data display Resume – resume refreshing the data display Show Status Bar – show the status bar Hide Status Bar - hide the status bar
Filter	<ul style="list-style-type: none"> No Filter – delete all filters for the current monitor <i>specific filters ...</i> – choose the filter you want to set; the list of available filters is different for each monitor
Help	<ul style="list-style-type: none"> Adaptive Server Help - display help topics for the Adaptive Server Enterprise plug-in for Sybase Central Monitor Viewer Help - display help topics for Monitor Viewer About - displays version information for Monitor Viewer

Setting up a Monitor Session

This section explains how to set up a monitor to collect the type of data you need. Topics are:

- Setting the statistic type
- Setting the sample interval
- Pausing and refreshing data on the window
- Setting filters

Setting the Statistic Type

For most monitors, you have the choice of displaying:

- Sample data – performance activity that occurred during the last sample interval
- Cumulative data – performance activity since the session started (that is, since the monitor was opened)
- Both – both sample and cumulative data. In places where there are no column headings, the two values are displayed side by side, with the sample data being the first value and the cumulative data being the second value.

To determine which statistic type is currently being displayed:

1. Look at the status bar, if it is showing.
2. Otherwise, from the View menu, choose Statistics. A check mark indicates which statistic type is in effect.

To change the statistic type:

1. From the View menu, choose Statistics.
2. From the cascading menu, choose the desired statistic type.

► **Note**

In monitors that display some graphical data and some table data, the statistic type applies to the table data only. You cannot change the statistic type depicted in the graphs.

Setting the Sample Interval

The sample interval is the amount of time Monitor Server waits between obtaining samples on behalf of the monitor. You can set different sample intervals for each monitor.

The default sample interval is 10 seconds. The shortest allowable sample interval is 1 second. The longest allowable sample interval is one day.

To change the sample interval for a monitor:

1. From the View menu, select Sample Interval.



Figure 5-3: Set Sample Interval dialog box

2. Click in the appropriate box to change the hours, minutes, or seconds.
 - Hours – number of hours between samples.
 - Minutes – number of minutes, in addition to the specified hours, between samples.
 - Seconds – number of seconds, in addition to the specified hours and minutes, between samples.

Notes about Setting Sample Interval

Short sample intervals cause frequent refreshes of data in a monitor. When data refreshes so often that you cannot read the display, choose a longer interval or pause the refreshes.

The status bar indicates the time that the last sample was obtained by Monitor Server.

You might notice discrepancies in the rates reported by different monitors for related data items over the same sample interval. Discrepancies might occur because the rates are averaged over the period of the actual sample interval for each data point. The actual sample interval for several data points can vary with network traffic and server activity, even among sampling requests made at the same time. Differences in sample interval lengths can also cause discrepancies in data.

Very short sample intervals in monitors that obtain many event summaries can potentially have an effect on a busy Adaptive Server installation. For more information about event summaries, see the *Adaptive Server Enterprise Monitor Server User's Guide*.

You can change the sample interval in the middle of a session without affecting the cumulative values collected for the session.

Pausing and Resuming Data Refresh

The monitor window refreshes automatically at the end of each sample interval. You might need to stop the refreshes to examine data.

To pause refreshes:

1. From the View menu, choose Pause.

During the pause, Monitor Server continues to collect data in the background.

To resume refreshes:

1. From the View menu, choose Resume.

The next sample reflects all activity that occurred during the pause, plus the latest sample.

Setting Filters

Filters limit the data displayed in a monitor. Most data items in a monitor can be filtered, but only one filter can be active at a time in a monitor. The default when you first open a monitor is no filter.

The Filter menu lists the filters available for a monitor. A check mark beside a filter command indicates which filter is currently set. For example, Figure 5-4 shows an Engine Activity Monitor with no filters set.



Figure 5-4: Filter menu example

To set a filter:

1. From the Filter menu, choose the data item that you want to filter.
2. Complete the dialog box. The type of dialog box that opens depends on the filter you choose in step 1.

3. Click OK to activate the filter. The next refresh shows filtered data.

Canceling A Filter

There are two ways to cancel a filter:

1. Set another filter. Only one filter can be in effect at a time for a monitor.

OR

2. From the Filter menu, choose No Filter.

Setting Filters on Numeric Data

For data items whose values are numeric, you can choose between either of these filter types:

- Top *n* filter – displays the *n* rows with the highest values in the filtered column. For example, if you set a top 10 filter on the Logical Page Reads column in the Process Activity monitor, the monitor displays data for the 10 processes with the highest Logical Page Read values.
- Threshold value filter – displays rows with values in the filtered column that are equal to or exceed the specified threshold value. For example, if you set a threshold of 100 on the Logical Page Reads column in the Process Activity monitor, the monitor displays data for all processes with a Logical Page Read value equal to or greater than 100.

To set a top *n* filter:

1. For convenience, set the statistic type before setting the filter.
2. From the Filter menu, select the numerical data item you want to filter. The Set *<integer>* Filter dialog box opens. If the data item is a floating point number, a slightly different dialog box opens.




Figure 5-5: Set *<integer>* Filter dialog box

3. Select the Top N button.
4. In the text box next to the button, enter the number of rows you want displayed.
5. If the dialog box includes buttons for selecting the statistic type:
 - Select the Sample button to apply the filter to data from the most recent sample.
 - Select the Cumulative button to apply the filter to cumulative (session) data.

► **Note**

If the dialog box does not include these buttons, the filter applies to the statistic type currently in effect.

To set a threshold filter:

1. For convenience, set the statistic type before setting the filter.
2. From the Filter menu, select the numerical data item you want to filter. The Set *<integer>* Filter dialog box opens. If the data item is a floating point number, a slightly different dialog box opens.
3. Select the Threshold button.
4. In the text box next to the button, enter a threshold value. It must be a positive integer. The monitor displays values that are equal to and above the specified threshold value.

5. If the dialog box includes buttons for selecting the statistic type:
 - Select the Sample button to apply the filter to data from the most recent sample.
 - Select the Cumulative button to apply the filter to cumulative (session) data.

► **Note**

If the dialog box does not include these buttons, the filter applies to the statistic type currently in effect.

Setting Character String Filters without a Scroll List

For filters on character string data items when a scroll list of specific values is not possible, you specify the filter value in the filter dialog box. A filter on Application Name or Execution Class is specified in this way.

To specify filters without a scroll list of choices:

1. From the Filter menu, select a string data item:



Figure 5-6: Set <string> Filter dialog box

2. In the Value box, type the string value that you want to filter on. You can use a **single** instance of the wildcard character (%) anywhere in the string (at the beginning, end, or anywhere in the middle). For example, %account is a valid filter value; %account% is not valid.

3. Select one of the option buttons:

- The Include button defines a filter that includes statistics only when the value in the filtered data item matches the specified string.
- The Exclude button defines a filter that includes statistics **except** when the value in the filtered data item matches the specified string.

Setting a Database Object Filter

The database object filter includes or excludes entire databases or selected objects within databases. This section describes setting filters on databases, tables, or stored procedures.

To filter by database or database object:

1. From the Filter menu, choose Object Name or Stored Procedure Name.

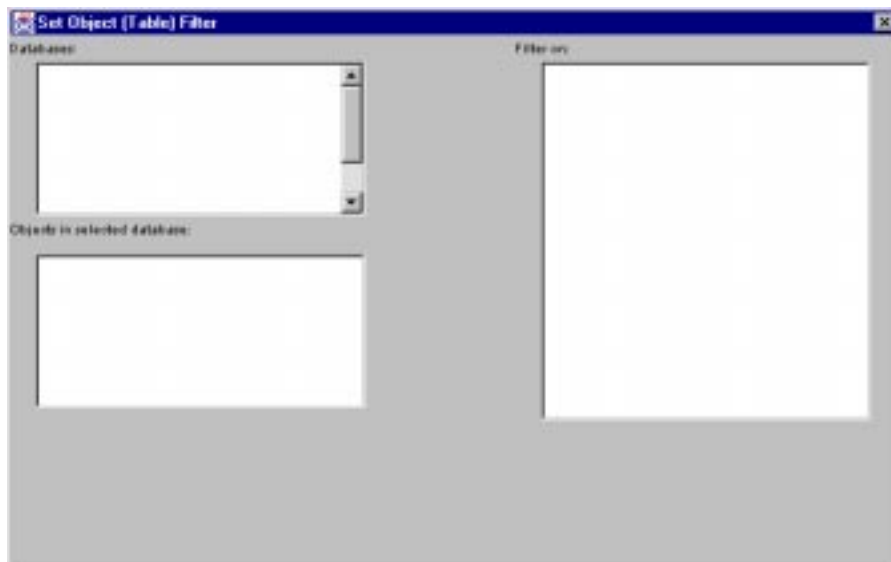


Figure 5-7: Set Database Object Filter dialog box

2. Select one of the option buttons:

- The Include button defines a filter that displays statistics only when the value in the filtered data item matches the specified string.
- The Exclude button defines a filter that displays statistics **except** when the value in the filtered data item matches the specified string.

To construct the Filter On list:

1. Select a database by clicking on it in the Databases scroll list. The list shows the databases defined in Adaptive Server.
2. To filter on all of the objects in the database, click the Add button while the database is selected in the top scroll list. This action causes the following entry in the Filter On scroll list:

```
databaseName.%
```

The % character is a wildcard that signifies any string of zero or more characters.

3. To filter on a subset of objects from the selected database:
 - a. Select one or more lines in the Objects in Selected Database scroll list. The list includes all tables or stored procedures defined in the selected database. To select an object, click on it. To select multiple objects, click on each one separately. To deselect an object, click on it again.
 - b. Click Add to move your selections into the Filter On scroll list. To remove a line from the Filter On list, select the line and click Remove.
 - c. Repeat until you complete your list of database objects for the filter.

Setting a Data Cache Filter

To filter by data cache name:

1. From the Filter menu, select Data Cache.



Figure 5-8: Set Data Cache Filter dialog box

2. Select one or more data caches in the scroll list. The list shows the default data cache and all named caches currently configured in Adaptive Server. Select a data cache by clicking on it. Select multiple data caches by clicking on each one individually. To deselect a data cache, click on it again.
3. Select one of the option buttons:
 - The Include button creates a filter that includes the data caches you selected in the scroll list.
 - The Exclude button creates a filter that reports on all data caches except those you selected in the scroll list.

Setting a Lock Status Filter

To filter by lock status:

1. From the Filter menu, select Lock Status.



Figure 5-9: Set Lock Status Filter dialog box

2. Select one or more lock status values in the scroll list. Select a lock status by clicking on it. Select multiple lock status values by clicking on each one individually. To deselect a lock status, click on it again.
3. Select one of the option buttons:
 - The Include button creates a filter that includes the locks whose status matches the ones you selected in the scroll list.
 - The Exclude button creates a filter that reports on all locks except those whose status matches the ones you selected in the scroll list.

Setting a Lock Type Filter

To filter by lock type:

1. From the Filter menu, select Lock Type.



Figure 5-10: Set Lock Type Filter dialog box

2. Select one or more lock types in the scroll list. Select a lock type by clicking on it. Select multiple lock types by clicking on each one individually. To deselect a lock type, click on it again.
3. Select one of the buttons:
 - The Include button creates a filter that includes locks whose type matches the ones you selected in the scroll list.
 - The Exclude button creates a filter that reports on all locks except those whose type matches the ones you selected in the scroll list.

Setting a Process ID Filter

To set a Process ID filter:

1. From the Filter menu, select Process ID.



Figure 5-11: Set Process Filter dialog box

2. Select one or more processes in the Process ID scroll list.

The list includes all currently running server process IDs (SPIDs) and their associated login names. To select a process, click on it. Select multiple processes by clicking on each one separately. To deselect a process, click on it again.

When a server process terminates, Adaptive Server can reuse its Process ID for a new process. Therefore, a SPID does not uniquely identify a process.

3. Select one of the option buttons:

- The Include button creates a filter that displays only the processes you selected in the scroll list. If a selected process terminates before the monitor session ends, the terminated process does not appear in the next refresh of the monitor display.

- The Exclude button creates a filter that displays all processes except those you selected in the scroll list.

Setting a Process State Filter

To filter by process state:

1. From the Filter menu, select Process State.



Figure 5-12: Set Process State Filter dialog box

2. Select one or more process states in the scroll list. Select a process state by clicking on it. Select multiple process states by clicking on each one individually. To deselect a process state, click on it again.
3. Select one of the option buttons:
 - The Include button creates a filter that includes the processes whose state matches the ones you selected in the scroll list.
 - The Exclude button creates a filter that reports on all processes except those whose state matches the ones you selected in the scroll list.

Setting a Device Filter

To filter by device name:

1. From the Filter menu, select Device.



Figure 5-13: Set Device Filter dialog box

2. Select one or more database device names in the scroll list. Select a device by clicking on it. Select multiple devices by clicking on each one individually. To deselect a device, click on it again.
3. Select one of the option buttons:
 - The Include button creates a filter that includes the devices you selected in the scroll list.
 - The Exclude button creates a filter that reports on all devices except those you selected in the scroll list.

Setting an Engine Filter

To filter by engine number:

1. From the Filter menu, select Engine.



Figure 5-14: Set Engine Filter dialog box

2. Select one or more engines in the scroll list. The list shows the engine numbers currently configured in Adaptive Server. Select an engine by clicking on it. Select multiple engines by clicking on each one individually. To deselect an engine, click on it again.
3. Select one of the option buttons:
 - The Include button creates a filter that includes the engines you selected in the scroll list.
 - The Exclude button creates a filter that reports on all engines except those you selected in the scroll list.

Setting up the Performance Trends Monitor

The Performance Trends Monitor displays, in a graphical format, values for the last 60 sample intervals for one or more data items. Using a dialog box, you choose which data items you want to see.

Each data item has a separate graph. If there is not enough room on your screen to display graphs for all of the items you choose, the monitor shows titles without the graphs.

To set up the Performance Trends Monitor:

1. Choose the Performance Trends Monitor from the Sybase Central window or from the File --> Monitors command in a monitor window. The Select Data Items dialog box opens.



Figure 5-15: Select Data Items dialog box

2. In the Select Data Items dialog box, choose one or more items that you want to see graphs for by clicking on them. Deselect an item by clicking on it again.

You can choose any number of items, but 4 to 6 is recommended. More than 6 graphs will not fit on most display terminals.

3. Click OK. The Performance Trends Monitor opens.
4. From the Performance Trends Monitor, choose View --> Sample Interval to set the sample interval length.

The monitor displays 60 sample intervals at a time. If you set the sample interval to 10 seconds, you can see the activity trends for 10 minutes at a time. If you set the sample interval to 30 minutes, you can see the activity trends for 30 hours at a time.

Troubleshooting

This section provides some monitor troubleshooting guidelines.

Unsuccessful Connections

If connections to Monitor Server are unsuccessful:

- Monitor Server might not be running.
- You might not have the necessary permissions.
- An error might exist in a *sql.ini*, *interfaces*, or other Directory Services file. Check the entries for Monitor Server on both the machine where Monitor Server is running and the Sybase Central machine. The entries should match. See the *Monitor Server User's Guide* for more information about configuring Monitor Server.
- The number of concurrent connections to Monitor Server is limited by a Monitor Server configuration setting. You might have too many monitors open, or other users, including users of Historical Server, might also hold connections to Monitor Server. See the *Monitor Server User's Guide* for more information about connection limits.
- If the error occurs only for the Transaction Activity Monitor or the Object Page I/O Monitor, and you are monitoring SQL Server release 11.0.x, you are connected to the wrong Monitor Server executable file. The Monitor Server executable file for monitoring SQL Server release 11.0.x is different from the one used for monitoring Adaptive Server release 11.5. See the *Release Bulletin for Adaptive Server Enterprise Monitor* for more information.

- The CLASSPATH system variable on the Sybase Central machine might not include all necessary paths. It must point to the following files:
 - *monclass.zip* and *3pclass.zip* in the Sybase Central installation directory. The default installation path names are:
disk:\Sybtools\ASEP\monclass.zip
disk:\Sybtools\ASEP\3pclass.zip
where *disk* is the installation disk.
 - *lib\classes.zip* in the JDK 1.1.3 installation directory.
- The PATH system variable on the Sybase Central machine might not include all necessary paths. It must point to the following directories:
 - *%SYBASE%\bin*, where *%SYBASE%* points to the Sybase installation directory.
 - *Sybtools\ASEP*, where *Sybtools* is the Sybase tools installation directory.
 - *Sybtools\win32*, where *Sybtools* is the Sybase tools installation directory.
 - *jdk_install\bin*, where *jdk_install* is the Java installation directory.

These variables are set appropriately during installation. See *Installing Sybase Products* for more information.

Unexpected Adaptive Server Name

If the Adaptive Server name in a monitor's title bar does not match the Adaptive Server icon you selected in Sybase Central:

- The wrong Monitor Server might be specified in the Monitor Server tab on the Adaptive Server property sheet. See “Configuring a Monitors Folder” on page 5-5.
- Aliases might be used for server names.
- The **-S** startup parameter to Monitor Server might specify the wrong Adaptive Server name.

In any of these cases, the monitor displays a warning dialog box when it first opens, and it displays performance statistics about the Adaptive Server whose name appears in the title bar.

Cumulative Values Return to Zero

If cumulative values return to zero in the middle of a session:

- A connection to Monitor Server is using the **sp_sysmon** command.

The **sp_sysmon** stored procedure collects server-level performance data. See the *Performance and Tuning Guide* for more information.

The **sp_sysmon** stored procedure and Monitor Server use the same internal counters for gathering statistics on Adaptive Server. Whenever **sp_sysmon** is invoked, it resets these counters to zero, which causes erroneous output for any Monitor Server connections in progress. Do not use **sp_sysmon** and Sybase Central monitors simultaneously.

6

Managing Remote Servers

What's in This Chapter

This chapter describes how to add and configure a remote server. It also describes how to set options that govern remote server connections.

You can configure the Adaptive Server installations you administer so that a user connected to Adaptive Server can request execution of a stored procedure on another Adaptive Server. The result of such a request, called a **remote procedure call (RPC)**, is returned to the calling process running on the Adaptive Server to which the user is connected.

If Component Integration Services (CIS) is enabled for Adaptive Server, you can also access data on a remote server as if it were on the local server.

Configuring RPC Handling

This section discusses the mechanisms that Adaptive Server can use to manage RPCs and how to configure them. Your choice of RPC handling method affects Adaptive Server configuration and login mapping for remote servers. The options are site handler and CIS RPC handling.

The default method for handling interaction between local and remote servers is through a **site handler**. A site handler creates a physical connection between the local server and remote server. Then it creates a logical connection for each RPC to the remote server. Adaptive Server creates a site handler for each remote server it connects to. Site handler is used only for connections between two Adaptive Server installations.

CIS RPC handling is always used for connections involving proxy tables. It creates connections using Client-Library functions. You can also enable it for use with all RPCs.

The principal difference between the two methods of handling RPCs is how the remote server views the RPC:

- If you use site handler, the remote Adaptive Server detects that the logical connection is made by another remote server and performs remote server verification through *sysremotelogins*.

- If you use CIS RPC handling, the remote server sees the RPC as an ordinary client connection. There is no verification using *sysremotelogins*. Therefore, connections must have a valid Adaptive Server login account established prior to the connection request. You cannot use trusted mode.

Use of CIS RPC handling allows you to include RPCs in a transaction. Work done by an RPC can be committed or rolled back along with the other work performed in the transaction.

Configuring Adaptive Server to Use Site Handler

Site handler is the default RPC handling method. To enable execution of RPCs using site handler:

1. Add each Adaptive Server of a remote server pair as remote servers to the other. See “Adding a Remote Server” on page 6-4.
2. For each Adaptive Server, add itself as a remote server with server class “local”. See “Adding a Remote Server” on page 6-4.
3. Make sure the *allow remote access* configuration parameter is set to 1 on each Adaptive Server. This is the default value. See “Setting Configuration Parameters” in Chapter 4, “Managing the Adaptive Server Operating Environment”.
4. On each Adaptive Server in the remote server pair, create a default or explicit login mapping. See “Setting Default Login Mapping” on page 6-12 and “Setting Explicit Login Mappings” on page 6-13.
5. Ensure that each Adaptive Server in the remote server pair lists the other in its **interfaces file** or its Directory Services.
6. Reboot both servers.

Figure 6-1 illustrates the process of adding remote servers. Assume that NOME is the Adaptive Server being managed (the local server) and DUBUQUE is the remote server. NOME adds remote server DUBUQUE and adds itself as a remote server with class local. DUBUQUE must also add NOME as a remote server and add itself as a remote server with class local. Note the bidirectional configuration.

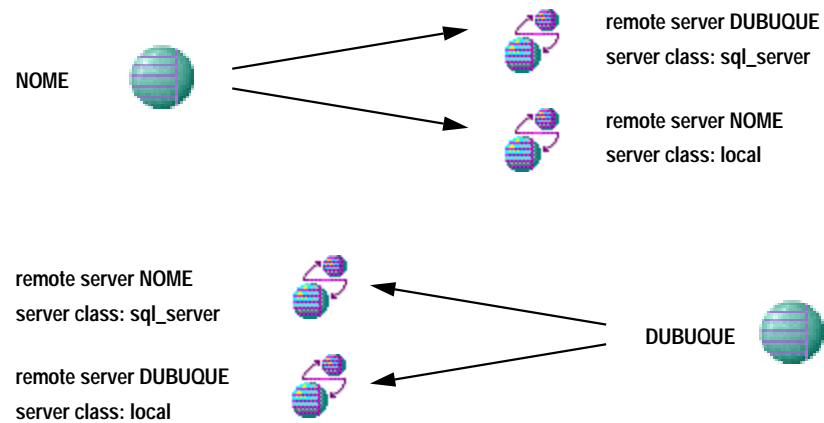


Figure 6-1: Example of configuring site handler RPC handling

Configuring Adaptive Server to Use CIS RPC Handling

This section explains how to enable CIS RPC handling when the remote server is Adaptive Server or a non-Sybase data store accessed using Component Integration Services. To enable CIS RPC handling:

1. On the local Adaptive Server add a remote server. See “Adding a Remote Server” on page 6-4.
2. Set the `cis rpc handling` parameter to 1. See “Setting Configuration Parameters” in Chapter 4, “Managing the Adaptive Server Operating Environment”.
3. By default, CIS RPC handling uses the local user’s login and password to connect to the remote server. The local login and password must be valid on the remote server or you must map it to a valid login using CIS Login Mapping. See “Managing Logins for CIS RPC Handling” on page 6-14. For connections to non-Sybase servers, see the server documentation for additional requirements for remote access.
4. Ensure that the local Adaptive Server lists the remote server in its **interfaces file** or its Directory Services.

Figure 6-2 illustrates configuration of a remote server for CIS RPC handling. NOME is the local server. CONCORD is the remote server. Compare it to the example in Figure 6-1. Note the unidirectional configuration of CIS RPC handling as compared to the bidirectional configuration of site handler.




Figure 6-2: Example of configuring CIS RPC handling

Adding a Remote Server

To gain access to a remote server, you must define the remote server on the local Adaptive Server.

Privileges Only a System Security Officer can create a remote server.

To add a remote server:

1. Select the Remote Servers folder.
2. From the File menu, choose New; then choose Remote Server.

The Add New Server wizard asks for the following information:

Table 6-1: Inputs to Add New Server wizard

Input	Description
Local Name	A local name for the remote server.
Remote Server	The network name of the remote server. (The list has the names of servers known to Sybase Central. The remote server does not need to be on this list. However, it must be in the interfaces file for the Adaptive Server on which you are creating this remote server.)
Server class	If Component Integration Services is enabled, specify the server class of the remote server. If Component Integration Services is not enabled, accept the default server class: sql_server.

Shortcuts



Double-click the Add Remote Server icon.

Deleting a Remote Server

To delete a remote server, follow the procedure for deleting objects described in “Deleting an Object” on page 3-4.

Displaying Remote Server Properties




To display information about a remote server:

1. Select the remote server icon.
2. From the File menu, choose Properties.



Figure 6-3: Remote Server property sheet

Shortcuts

	Double-click the remote server icon.
	Select the remote server icon. Then, select the Properties toolbar button.
	Click the right mouse button over the remote server icon. Choose Properties from the shortcut menu.

Setting the Server Class of a Remote Server

If you will access a remote server using Component Integration Services, you must specify the appropriate server class for the remote server. The server class tells Component Integration Services how to connect to the remote server. The server classes are as follows:

- *sql_server* -- indicates that the remote server is Sybase Adaptive Server, or Microsoft SQL Server.
- *local* -- the local Adaptive Server. There can only be one.
- *direct_connect* -- functionally the same as server class *access_server*. *access_server* is supported for backward compatibility.
- *db2* -- indicates that the remote server is a gateway to DB2 or DB2-compatible databases. The Sybase Net-Gateway and MDI Database Gateway for DB2, AS/400, and SQL/DS are in this category.
- *generic* -- indicates that the remote server is an Open Server application that conforms to the interface requirements of a Generic Access Module, as described in *OmniSQL Server Generic Access Module Reference Manual*. The OmniSQL Server Access Module for Rdb and OmniSQL Server Access Module for Informix are two examples of remote servers in this class.
- *sds* -- is functionally the same as *direct_connect*. It is listed for future implementation. If you select *sds*, the property sheet lists *direct_connect*.

To specify the server class:

1. Open the property sheet for the remote server as described in “Displaying Remote Server Properties” on page 6-5.
2. On the General tab, choose a server class from the drop-down list.

Setting Security Options for Remote Servers

You can set the following security options for a remote server:

- Disconnect on timeout
- Encrypt network password
- Readonly status (Component Integration Services only)
- RPC security model

Select Disconnect On Timeout if you want the remote server to drop the physical connection from your local server when there has been no activity for one minute.

Select Encrypt Network Password if you want the local server to encrypt its login passwords when it connects to the remote server. The local server use an encryption key provided by the remote server.

Select Read Only status to ensure that the local server cannot write to the data on the remote server. Read only status is enabled only if Component Integration Services is enabled.

Set an RPC Security Model if you want to control the level of security applied to the connections made between local and remote servers when a remote server executes an RPC.

To specify a security option:

1. Open the property sheet for the remote server as described in “Displaying Remote Server Properties” on page 6-5.
2. Select the Security tab.



Figure 6-4: Remote server property sheet: Security tab

3. Check the options you want to enable.

Specifying an RPC Security Model

The security model controls the level of security applied to the connections made between local and remote servers when a remote server executes an RPC. Depending on the security mechanism you use, RPC security provides mutual authentication, message confidentiality, and message integrity for remote procedure communications.

You can specify one of two RPC security models, A or B. Model A is equivalent to the level of security in Sybase SQL Server release 11.0.x and prior releases.

In security model B the local Adaptive Server gets a credential from the security mechanism and uses the credential to establish a secure physical connection with the remote Adaptive Server. You can use one or more of the following security services:

- Mutual authentication - The local server authenticates the remote server by retrieving the credential of the remote server and verifying it with the security mechanism. With this service, the credentials of both servers are authenticated and verified.

- **Message confidentiality via encryption** - Messages are encrypted when sent to the remote server, and results from the remote server are encrypted.
- **Message integrity** - Messages between the servers are checked for tampering.

Security model B implements security using one of the following security mechanisms:

- Distributed Computing Environment (DCE) (UNIX)
- CyberSAFE Kerberos (UNIX)
- NT LAN Manager (Windows NT)

You can configure RPC security in the Remote Server property sheet.

Managing Remote Server Logins

When a login requests access to a stored procedure on a remote server, the login assumes the identity and access privileges of a login on the remote server. The System Security Officer controls how this mapping between logins occurs and whether the password of a login making an RPC from another server is checked.

Management of logins for remote servers differs depending on the type of RPC handling in use.

Mapping Logins for Site Handler RPC Handling

When you use site handler, login mapping is managed from the perspective of the remote server. When you set up login mappings, you specify how logins from the calling, or remote, server map to logins on the local server. Figure 6-5 illustrates the options for remote login mapping. The options are as follows:

- Logins from the remote server keep their own names on the local server. For this scheme to work, you must set up a login on the local server for each login that will be issuing RPCs from the remote server.
- All logins from the remote server are mapped to a single login on the local server, taking on the identity and privileges of that single login.

► **Note**

Mapping more than one remote login to a single local login is not recommended. It reduces individual accountability on the local server. Audited actions can be traced only to the local login, not to the individual logins on the remote server.

- Each login from the remote server is mapped to a specific login on the local server. You can use this option in combination with either of the others. The explicit mappings set up for logins under this option override the mappings established under the other methods.

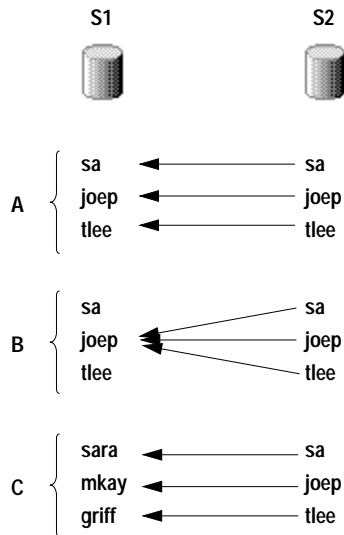


Figure 6-5: Remote login mapping options

Key to Figure 6-5:	Description
S1	Local server
S2	Remote server
A	Map remote login to local login with the same name
B	Map all remote logins to one local login
C	Map remote logins to specific local logins

► **Note**

Figure 6-5 does not apply to CIS RPC handling.

Setting Default Login Mapping

Privileges	Only a System Security Officer can establish remote server login mappings.
------------	--

To set up remote server login mapping:

1. Open the property sheet for the remote server as described in “Displaying Remote Server Properties” on page 6-5.
2. Select the Default Login Mapping tab.



Figure 6-6: Remote Server: Default Login Mapping tab

3. Select a login mapping method:
 - To specify that no automatic mapping is in place, select None. With this option, only logins explicitly mapped to logins on the local server are allowed access.
 - To allow each remote server login to keep the same login name on the local server, select Map to Local Logins with Same Name.
 - To map all remote server logins to a single local server login, select Map All to a Single Local Login. From the drop-down list below this option, select the local server login.

4. Optionally, set up explicit mappings between remote server logins and local server logins. See “Setting Explicit Login Mappings” on page 6-13 for this procedure.
5. When you are satisfied with all login mapping details, click Apply.

Setting Explicit Login Mappings

When you set up explicit mappings in addition to an automatic mapping scheme, the explicit mappings override the automatic ones. When you specify None as the default login mapping, only logins explicitly mapped to a local server login have access to the local server. To set up an explicit mapping for a remote server login:

1. Open the Remote Server property sheet as described in “Displaying Remote Server Properties” on page 6-5.
2. Select the Specific Login Mapping tab.

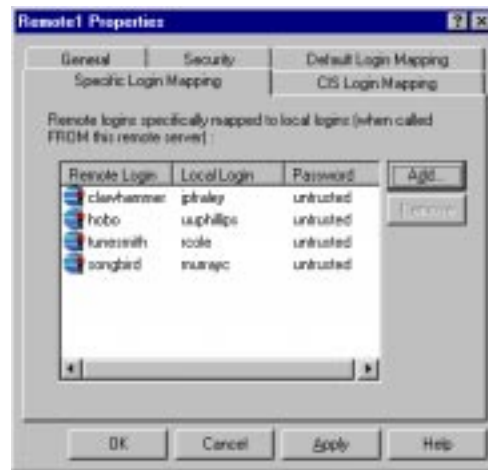


Figure 6-7: Remote Server dialog box: Specific Login Mapping tab

3. Click the Add button. The Add Specific Mapping dialog box opens.



Figure 6-8: Add Specific Mapping dialog box

4. Enter the name of the remote server login in the Remote Login box.
5. Select a local server login to map to in the Local Login drop-down list.
6. To specify that the remote server login should obtain access without password checking, choose Trusted Password.
7. Click OK.
8. Repeat steps 3 to 6 as often as necessary to set up the required explicit login mappings.

Removing a Specific Login Mapping

To remove a specific login mapping:

1. Open the Remote Server property sheet as described in “Displaying Remote Server Properties” on page 6-5.
2. Select the Specific Login Mapping tab.
3. Select the mapping you want to remove.
4. Click Remove.

Managing Logins for CIS RPC Handling

Logins for CIS RPC handling are mapped from the perspective of the local server. You identify a local login and the remote (or external) login that you want the local login to use when it connects to the remote server. The default is to use the local login of the user making the RPC call as the remote login.

To map logins for Component Integration Services:

1. Open the property sheet for the remote server as described in “Displaying Remote Server Properties” on page 6-5.
2. Select the CIS Login Mapping tab.




Figure 6-9: Remote server property sheet: CIS Login Mapping tab

3. Click the Add button. The Add CIS Mapping dialog box opens.



Figure 6-10: Add CIS Mapping dialog box

4. In the Local Login box, select a login from the drop-down list.
5. In the remote login box, type the remote login you want the local login to use to connect to the remote server.
6. Enter and confirm a password for the remote login.

Removing a Login Mapping for CIS RPC Handling

To remove a login mapping for CIS RPC Handling:

1. Open the property sheet for the remote server.
2. Select the CIS Login Mapping tab.
3. In the list, select the login mapping you want to remove.
4. Click the Remove button.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

7

Managing Cache

What's in This Chapter

Adaptive Server supports **cache** management features. Cache is a portion of Adaptive Server memory allocated when an Adaptive Server installation is created. Its purpose is to temporarily hold information currently in use by Adaptive Server and to hold recently used data pages. Because it is faster to read from memory than from disk, holding information in cache improves performance.

Adaptive Server cache is divided between memory reserved for data and memory reserved for stored procedure execution plans. You can partition the memory allocated for data cache (the default cache) into user-defined (named) data caches, and you can bind databases, indexes, or tables to a specific cache.

Within a named data cache, you can create buffer pools that are larger than the default size of 2K, enabling Adaptive Server to read larger amounts of data in a single I/O. Larger I/Os can enhance performance.

This chapter contains information about:

- Creating and deleting named data caches
- Displaying and modifying cache information
- Binding objects to a cache

Creating a Named Data Cache

This section describes how to create a named cache within the default data cache configured for Adaptive Server.

Privileges	Only a System Administrator can create, modify, and delete a cache.
------------	---

To create a named cache:

1. Select the Caches folder for Adaptive Server.
2. From the File menu, choose New; then choose Cache. The Create Cache wizard opens.

The Create Cache wizard asks for the following information:

Table 7-1: Inputs to Create Cache wizard

Input	Description
Name	Name for the new cache
Size	Size for the new cache. In pages, MB, KB, or GB.
Storage option	Store data and log pages or only log pages.

Notes

- The minimum cache size is .5MB.
- Because creating a cache involves some overhead, the actual amount of available memory will be slightly less than the amount you specify for size.
- Adaptive Server creates the cache with a status of “Activation Pending.” In this state, the cache is defined in Adaptive Server, but cannot be used for holding data and cannot be bound to a database, index, or table until you stop and restart Adaptive Server.

After you restart Adaptive Server, the cache status becomes “Active,” and the cache is ready for use.

Shortcut



Select the Caches folder. In the right pane double-click the Add Cache icon.

Deleting a Cache

To delete a cache, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Sybase Central changes the cache status to “Delete (pending).” In this state the cache is marked for deletion, but is not physically deleted until you stop and restart Adaptive Server.

Displaying Cache Properties




The Cache property sheet provides current information about a cache. To display cache properties:

1. Select the icon of the cache you want to display.
2. From the File menu, choose Properties.



Figure 7-1: Cache Properties

Shortcuts

	Double-click the cache icon.
	Select the cache icon. Then, click the Properties button in the toolbar.
	Click the right mouse button over the cache icon. Choose Properties from the shortcut menu.

Changing Cache Size or Type

To change the size or type of a named cache:

1. Display the Cache property sheet, as described in “Displaying Cache Properties” on page 7-3.
2. Select the Configuration tab.



Figure 7-2: Cache property sheet: Configuration tab

3. Select one of the options:
 - To modify a cache to hold both data and transaction log information, select “Stores Data and Log Pages.”
 - To modify a cache to hold only transaction log information, select “Only Stores Log Pages.” This option is not available if the cache has any databases, indexes, or tables bound to it. To change a cache to a log-only cache, first unbind any nonlog bound objects from it. You cannot change the default cache to a log-only cache.
4. To change the size, edit the Currently Configured text box. The new size is subject to the following constraints:
 - When you reduce the size, all space is removed from the default buffer pool. You cannot specify a size that reduces the default pool to less than the 512K minimum. For details, see “Defining Buffer Pools” on page 7-5.
 - When you increase the size, all new space is added in the default buffer pool. You cannot increase the size to an amount that exceeds the value in the Maximum Size box.

► **Note**

Because creating a cache involves some overhead, the actual amount of available memory is slightly less than the amount in the Maximum Size box.

5. Click OK or Apply. For the new specifications to take effect, you must stop and restart Adaptive Server.

Defining Buffer Pools

When you first create a cache, all space is assigned to the default memory (buffer) pool consisting of buffers that are the default I/O size for Adaptive Server (2K on most platforms). After creating a cache, you can create additional pools with buffers larger than the default. I/Os for these pools are the size of their buffers—for example, 4K I/Os for a pool with 4K buffers. In a database with a high level of I/O activity, defining pools with buffers larger than the default can make I/Os more efficient.

► **Note**

Moving space between buffer pools within a cache takes effect immediately. You do not need to restart Adaptive Server.

Each buffer pool you create is identified by the size of its buffers. The following sizes are allowed: 2K (the default on most Adaptive Servers), 4K, 8K, and 16K.

In any cache, you can have only one pool of each buffer size.

Figure 7-3 shows a data cache consisting of the default cache (A) and two named caches (B and C). The default cache has only the original 2K pool. Cache B has a 2K pool and a 16K pool. Cache C has a 2K pool, a 4K pool, and a 16K pool.




Figure 7-3: Default cache and two named data caches

When you create a buffer pool, you reassign space from the default pool. For example, if you create a data cache with 50MB of space, all of the space is initially assigned to the default pool. If you create a 30MB pool of 16K buffers, the default pool is reduced to 20MB. If you remove a 4K, 8K, or 16K buffer pool, the space returns to the default pool.

Creating a Buffer Pool

Privileges	Only a System Administrator can create, modify, and delete buffer pools.
------------	--

You can define buffer pools on named caches in Active status or on the default data cache.

To create a buffer pool:

1. Display the Cache property sheet as described in “Displaying Cache Properties” on page 7-3.
2. Select the Buffer Pool tab.



Figure 7-4: Cache Property Sheet: Buffer Pool tab

3. Click the Add button. The Add/Change Memory Pool dialog box opens.



Figure 7-5: Add/Change Memory Pool dialog box

4. In the I/O Size box, select a size from the drop-down list. Allowable sizes are 4K, 8K, and 16K. The list includes only those sizes that have not yet been created.
5. In the Amount in Pool box, enter the amount of space to reassign from the default buffer pool. The new pool must not reduce the 2K pool to less than 512K. As you change the Amount in Pool value, the space remaining in the 2K pool is reported in the Affected Pool group box.
6. Click OK.

► **Note**

When you create a buffer pool, you cannot set the wash size. Sybase Central sets a default wash size. After you create the buffer pool, you can edit the wash size.

Changing a Buffer Pool

To change a buffer pool:

1. Display the Cache property sheet as described in “Displaying Cache Properties” on page 7-3.
2. Select the Buffer Pool tab.
3. In the list box, select the buffer pool that you want to change.

4. Click the Change button. The Add/Change Memory Pool dialog box opens. You can change the size of the pool or the wash size.
5. To change the size of the pool, edit the number in the Amount In Pool text box.
6. To change the wash size, change the number in the Wash Size box. The **wash area** is a portion of each memory pool used to ensure that queries that need clean pages in a data cache can find them. When the number of dirty pages (pages that have been changed in cache) fills the rest of the buffers and begins to enter the wash area, Adaptive Server writes the data in the wash area to disk. When this write completes, the wash area is marked clean and is available for queries needing clean pages.

The minimum wash size is 10 times the I/O size for the pool. For example in a 4K pool, the minimum size is 40KB. The maximum wash size is 80% of the pool size.

7. Click OK.

Deleting a Buffer Pool

You can delete any buffer pool except the default pool. If you delete a 4K, 8K, or 16K buffer pool, the space returns to the default pool. To delete a buffer pool:

1. Display the Cache Properties tab, as described in “Displaying Cache Properties” on page 7-3.
2. Select the Buffer Pool tab.
3. In the Buffer Pools list, select the pool you want to delete, and click Remove.
4. Click Apply.

Binding Objects to a Cache

You can bind databases, indexes, or tables to a named data cache. Binding an object to a cache allows you to specify where recent reads from an object are held in memory and to take advantage of cache buffer pools that allow larger I/Os than the default. When you bind an object to a cache:

- Adaptive Server clears any pages for the object currently in memory.
- Adaptive Server reads pages for the object into the bound cache when subsequent queries need them.

Cache bindings take effect immediately. You do not need to restart Adaptive Server.

Privileges	Only a System Administrator can bind and unbind objects to a cache.
------------	---

Rules for Binding

The following rules govern cache bindings:

- You can bind objects only to an active cache.
- You can bind a database or database object to only one cache.
- You can bind a database to one cache and bind individual tables or indexes in the database to other caches.
- You can bind an index and the table it references to different caches.
- You cannot bind the *master* database, the system tables in *master*, or the indexes on the system tables in *master* to a named cache. However, you can bind nonsystem tables in *master* and their indexes.
- If you bind a database to a cache and do not bind its objects, Adaptive Server treats the objects as if they were bound to the database cache.
- If you bind neither a database nor its objects to a cache, Adaptive Server uses the default cache for their reads.
- When you delete a database, table, or index, Adaptive Server deletes all of the associated cache bindings.

Creating a Binding

To bind a database or object to a cache:

1. Display the Cache property sheet as described in “Displaying Cache Properties” on page 7-3.
2. Select the Cache Bindings tab. The tab displays the bound objects that meet the filtering criteria specified in the Show Object Bindings For drop-down list.



Figure 7-6: Cache property sheet: Cache Bindings tab


3. Filter the list so it includes the objects you want to bind. You can select databases, tables, or indexes.
4. Click the Bind button. The New Cache Binding dialog box is displayed.



Figure 7-7: New Cache Binding dialog box

5. For tables or views, select the database in which they are located from the Database drop-down list.
6. Select the objects you want to bind to this cache.
7. Click OK.

Shortcuts



Drag and drop the table or database icon from the left pane onto the cache icon in the right pane.

Deleting a Binding

To unbind a database or object from a cache:

1. Display the Cache property sheet as described in “Displaying Cache Properties” on page 7-3.
2. Select the Cache Bindings tab. The tab displays the bound objects that meet the filtering criteria specified in the Show Object Bindings For drop-down list.
3. Filter the list to find the object you want to unbind.
4. Select the object you want to unbind.
5. Click the Unbind button.
6. Click Apply.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

8

Managing Database Devices

What's in This Chapter

Adaptive Server has defaults for many aspects of storage management—where to place databases, tables, and indexes and how much space to allocate for each of them. However, as a System Administrator, you can override storage management defaults and customize the aspects of storage management that you require. This allows you to manage storage allocations for the level of performance and control you need in your environment.

Privileges	Only a System Administrator can perform the activities described in this chapter.
------------	---

This chapter describes how to use Sybase Central to:

- Create, modify, and delete database devices and to map them to physical resources in your environment
- Maintain a pool of default devices from which Adaptive Server allocates space when no specific device is requested
- Mirror database devices to ensure nonstop recovery in case of physical disk crashes

About Database Devices

A database device is a portion of a disk device that is used for storage of a database and its transaction log. Figure 8-1 illustrates a physical disk device (A). The physical device has two logical database devices (B and C).




Figure 8-1: Logical database devices on a physical disk device

An Adaptive Server installation includes the creation of a master device for the *master*, *model* and *tempdb* databases and another device for the storage of the *sybssystemprocs* database. The master device, *d_master*, is the initial default storage device for all user-defined databases. Because the master device stores the most critical Adaptive Server system tables, it is a good idea to:

- Designate other devices as default storage devices, and remove the default designation from the master device. See “Changing Default Device Designation” on page 8-12.
- For ease of recovery, reserve the master device for the system databases and create additional devices for your application databases.

Strategies for Recovery and Performance

To make good storage allocation decisions, you need to balance your requirements for recovery capability and performance. Adaptive Server supports the following strategies:

- Disk mirroring for nonstop recovery in the event of a device crash.
- Maintaining logs on a separate physical device. This enables you to recover by loading a database backup and applying transaction log records.

This chapter discusses these strategies in “Using Disk Mirroring” on page 8-7 and “Storage Management Considerations and Tips” on page 8-12.

Creating Database Devices

This section describes how to create logical database devices that map to your physical storage media. Before creating a new device, consider the following:

- When you create a database or dump device, Adaptive Server adds a row to the *sysdevices* table in the *master* database. Therefore, before creating a new device, you should back up the *master* database.
- Before creating a new database or dump device, make sure you have enough disk space on the physical device.
- On UNIX hosts, be sure that the “sybase” operating system account has write access on the physical device.
- After you create a new database or dump device, back up the *master* database again.

To create a database device:

1. Select the Database Devices folder.
2. From the File menu, choose New; then, choose Database Device from the cascading menu.

The Create a New Database Device wizard asks for the following information:

Table 8-1: Inputs to Create a New Database Device wizard

Input	Description
Name	Logical name for the new database device.
Pathname of device	Enter the full pathname of the physical device to which to map the logical name. On a UNIX host, this can be the name of a raw partition or operating system disk file.
Device number	Sybase Central provides the next available sequential number. You can enter any unused device number from 1 through 1 less than the value of the <i>devices</i> configuration variable.
Size	The device size in megabytes if your physical device is a file. If the physical device is a raw partition, the size of the partition.
Starting address	Default: 0
Mirror device	Specify if you want to mirror the device.
Pathname for mirror device	The physical name of the mirror device.

Shortcut



Select the Database Devices folder. In the right pane double-click the Add Database Device icon.

Deleting a Database Device

To delete a database device, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

If the device is mapped to an operating system disk file, delete the file to make the space available. This is not necessary if the device is mapped to a raw partition.

► Note

You cannot delete a device that has a database on it.

Displaying Device Properties

After you create a database device, you can examine or update its properties in the Database Device property sheet. You can also use the Database Device property sheet to get information about the databases allocated on the device.

To display the properties of a database device:

1. Select the icon of the database device.
2. From the File menu, choose Properties.



Figure 8-2: Database Device property sheet

Shortcuts



Double-click the database device icon.



Select the database device icon. Then, select the Properties toolbar button.



Click the right mouse button over the database device icon. Then, choose Properties from the shortcut menu.

Examining Database Allocations on a Database Device

The Database Device property sheet shows which databases have allocated storage space on the device and allows you to navigate to those databases. To display the allocation information:

1. Open the Database Device property sheet, as described in “Displaying Device Properties” on page 8-5.
2. Select the Databases tab.



Figure 8-3: Database Device property sheet: Databases tab

The Databases tab shows:

- The name of each database allocated to the device
- The amount of space allocated on the device for each database
- Creator of the device

Examining Segment Mapping on a Database Device

For a detailed look at the allocation of databases on a database device, display the Segments tab of the Database Device property sheet. The Segments tab shows the database segments mapped to the device. A **database segment** is a named portion of the database device storage allocated to a particular database. For information about database segments, see “Database Segments” on page 11-1.

To display segment mapping on a database device:

1. Open the Database Device property sheet, as described in “Displaying Device Properties” on page 8-5.
2. Select the Segments tab.



Figure 8-4: Database Device property sheet: Segments tab

Using Disk Mirroring

Mirroring is the duplication of the contents of an entire database device. When you set up a mirror device, all writes to the original device also go to the mirror device, and Adaptive Server reads from the original device or the mirror device based on efficiency. If one disk of a mirrored pair fails during a read or write, Adaptive Server automatically switches to the other disk and continues.

Figure 8-5 illustrates disk mirroring. Database device A, on disk B, is mirrored to mirror device C on disk D.




Figure 8-5: Database device on B mirrored on D

Disk mirroring provides a large benefit in recovery capability:

- It prevents downtime due to disk failures.
- It ensures full, nonstop recovery.
- It can improve read performance because either device is available as a source.

Disk mirroring also carries a cost:

- It consumes additional disk storage resources.
- It slows writes because they are duplicated.

Sybase highly recommends using mirror devices if you have the resources available.

► **Note**

Mirroring takes place for a database device, not for a database. To mirror the activities of a database that is allocated across several devices, you must mirror all of those devices.

Adding Disk Mirroring

To add disk mirroring to a database device:

1. Open the Database Device property sheet, as described in “Displaying Device Properties” on page 8-5.
2. Select the Mirror tab.

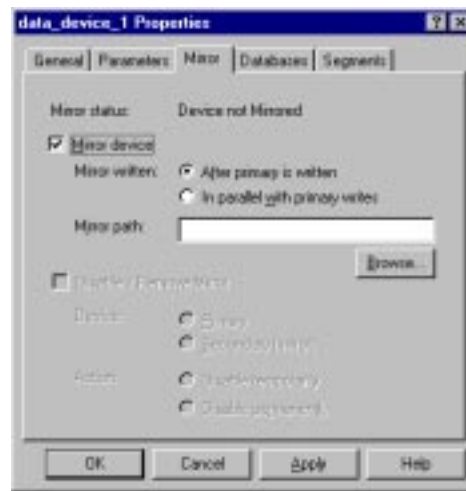


Figure 8-6: Database Device Properties: Mirror tab

3. Select the Mirror Device check box.
4. Specify the order in which the mirror device is written.
5. Specify the path of the mirror device.
6. Click Apply.

Disabling and Re-enabling Disk Mirroring

You can discontinue disk mirroring temporarily or permanently. If you temporarily disable mirroring, you can re-enable mirroring with the same mirror device.

Disabling Disk Mirroring

1. Open the Database Device property sheet, as described in “Displaying Device Properties” on page 8-5.
2. Select the Mirror tab.
3. Select the Disable/Remove Mirror check box.



Figure 8-7: Database Device Properties: Mirror tab

4. Specify which device to remove or disable.
5. Specify temporary or permanent removal of mirroring.
6. Click OK. Adaptive Server disables mirroring.

Re-Enabling Disk Mirroring

1. Open the Database Device dialog box, as described in “Displaying Device Properties” on page 8-5.
2. Select the Mirror tab.
3. Check the Mirror Device box.
4. Click Apply. Adaptive Server re-enables mirroring with the mirror device specified in the Mirror Path box.

► **Note**

If you permanently disabled a mirror device for which the physical device is an operating system file, you cannot re-enable mirroring using the same physical file name unless you first manually delete the operating system file for the disabled mirror device.

Storage Management Considerations and Tips

System Administrators must make many decisions regarding the physical allocation of space to Adaptive Server databases. Two major considerations that govern these decisions are recovery and performance. This section describes some strategies to consider.

Managing for Recovery

This section suggests strategies that can contribute to a clean recovery in the case of disk failure.

Changing Default Device Designation

When you create a database, you can specify the database device on which to create it. If you do not specify a device, Adaptive Server creates the database on a device that has been designated as a **default** device. If you have multiple default devices, Adaptive Server uses them in alphabetical order until each is full.

The master device created during an Adaptive Server installation holds the three system databases: *master*, *model*, and *tempdb*. This device is initially designated as a default device.

► **Note**

Sybase strongly recommends that you remove the default designation from the master device to prevent its getting cluttered and to make recovery easier in the case of a disk failure.

To change the default designation of a device:

1. Open the Database Device property sheet, as described in “Displaying Device Properties” on page 8-5.
2. Select the Parameters tab.



Figure 8-8: Database Device property sheet: Parameters tab

3. Select or clear the Default Device check box. Selecting the box makes the device a default device. Clearing the box removes the default designation.

Using Disk Mirroring

If you mirror each database device containing a database allocation, and each mirror device is on a different physical device from the database device, you can ensure database recovery without downtime. Any other scenario requires some amount of recovery from backups.

Separating Data and Transaction Log

If you do not mirror your database devices, another method to protect data is to keep data and transaction logs separate. If a database device fails, you can recover by loading from the most recent database backup and applying the transaction log backups from the time of the database backup until the time of the failure.

To separate data and transaction logs, create multiple database devices before creating a database. When you create the database, you can specify that the space you are allocating on one or more devices is for log only. Adaptive Server automatically makes any other devices allocated for the database into data-only devices. For details, see “Creating a Database” on page 9-2.

Managing for Performance

Separation is the key to maximizing performance through storage management. The following separations can improve performance:

- Placing a table on one physical disk and nonclustered indexes on another ensures that physical reads and writes are faster, since the work is split between two disk drives. For details, see “Placing an Index on a Segment” in Chapter 24, “Managing Indexes”.
- Splitting large tables across two physical disks can improve performance, especially for multi-user applications.
- Placing a database or object on one physical disk and its mirror on another reduces read time, since both devices are active. On the other hand, mirroring a device increases the time needed to write, because Adaptive Server writes transactions to both the database device and its mirror.

If you use Adaptive Server Enterprise Monitor at your installation, you may find it useful to run it with Adaptive Server to examine how performance is affected by changes in physical storage allocations. You can monitor Adaptive Server from Sybase Central using monitoring windows. See Chapter 5, “Using Monitors” and Chapter 26, “Interpreting Monitor Data.”

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

9

Managing Databases

What's in This Chapter

After creating database devices, you can create user databases and allocate them on the database devices. You can also create database segments, which enable you to control the placement of objects on the devices. This chapter describes how to:

- Create and delete a user database and allocate storage for it
- Transfer database ownership to a different user
- Generate and execute Data Definition Language (DDL) scripts
- Change database options
- Issue a manual database checkpoint
- Execute Database Consistency Checker (dbcc) commands

Considerations for Creating a New Database

Before creating a database, decide:

- What size to make the database
- Where to place the database and whether there is sufficient space
- Whether you will store the transaction log on a device separate from the data

Estimating Database Size

When estimating database size, keep in mind that it is easy to increase the size of a database but difficult to decrease it. To decrease the size of a database, you must delete the database and then recreate it.

To estimate the size of a database, consider:

- Anticipated size of its largest objects—tables and indexes
- Space for planned views, stored procedures, defaults, rules, and triggers
- Size of the transaction log
- Additional space for expansion based on anticipated activity

The Adaptive Server system procedure `sp_estspace` helps you estimate table and index space requirements based on the definition of a specific table. To use `sp_estspace`:

1. On a nonproduction database, create the tables and indexes you anticipate the new database will contain. For information about creating tables and indexes, see Chapter 22, “Managing User Tables” and Chapter 24, “Managing Indexes.”
2. Use a utility such as `isql` or SQL Advantage to open a session with Adaptive Server.
3. Run `sp_estspace` for each table, and add the results. For details on running `sp_estspace`, see the *Adaptive Server Reference Manual*.

The size of the transaction log depends on the type and quantity of transactions and the frequency of backups. As a starting point, allocate 10-25% of the overall size of the database for the transaction log.

Deciding Where to Allocate Storage

Consider the following storage options for a new database. You can:

- Store the database on a single database device or on multiple database devices.
- Specify the amount of storage space to allocate on each device.
- Store the transaction log on the same database device or on a different database device from the data. Storing the transaction log on a separate database device is highly recommended.

Creating a Database

Before creating a database, make sure enough space is available on the database devices you plan to use.

To create a database:

1. Select the Databases folder.
2. From the File menu, choose New; then choose Database in the cascading menu. The Create a New Database wizard opens.

The Create a New Database wizard asks for the following information:

Table 9-1: Inputs to Create a New Database wizard

Input	Description
Database name	Enter a name for the database
Database device	Specify the database device or devices on which to allocate the new database
Database device size	Specify a size for each database device
Data or log	Specify whether the database device will store data or the transaction log.
With override	Specify with override if you want to store data and log on the same device.
For load	If you are creating the database so you can restore it from a backup, check the For Load check box. This is the case only if you are recovering from media failure or if you are moving a database from one location to another.
Guest account	Specify whether to create a guest user in the database.

Notes

- If you do not enter a size, Adaptive Server allocates either the value of the `database size` configuration variable or the size of the `model` database, whichever is larger.
- If you have limited storage **and** must put the transaction log and the data on the same logical device, specifying With Override allows Adaptive Server to maintain the log on separate device fragments from the data.
- You cannot remove or change a database device after creating the database unless you first delete the database.

Shortcuts



Select the Databases folder. In the right pane double-click the Add Database icon.

Deleting a Database

Deleting a database deletes the database and all its objects.

Privileges	Only the owner of a database can delete it.
------------	---

To delete a database, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

► **Note**

It is a good practice to back up the *master* database after you delete a user database.

Deleting a Damaged Database

If a database is damaged, use the Repair and Drop option of the `dbcc` command to delete it. For information about running `dbcc`, see “Checking Database Consistency” on page 9-18.

Privileges	Only the owner of a damaged database can delete it.
------------	---

Displaying and Modifying Database Properties

After you create a database, you can examine or modify it in the Database property sheet. This section describes how to examine database details and how to make the following types of changes to a database that you can make in the Database property sheet:

- Extending storage allocations
- Modifying log cache properties
- Transferring database ownership

To view database properties:

1. Select the icon of the database you want to display.
2. From the File menu, choose Properties.



Figure 9-1: Database property sheet

Shortcuts



Select the database icon. Then select the Properties toolbar button.



Click the right mouse button over the database icon. Choose Properties from the shortcut menu.

Transferring Database Ownership

After creating a database, a System Administrator can transfer ownership to a different user. This allows the control of database creation and resource allocation to remain centralized in the hands of the System Administrator, while permitting database administration to rest with the principal database users.

When you change ownership of a database, the new owner must have an Adaptive Server login, but must not be a user of the database or have an alias in it.

Privileges	Only a System Administrator can transfer database ownership.
------------	--

To change database ownership:

1. Open the Database property sheet as described in “Displaying and Modifying Database Properties” on page 9-4.
2. Click the Change button that is next to the Owner box. The New Database Owner dialog box opens.



Figure 9-2: New Database Owner dialog box

3. Select the login for the new owner from the list.
4. To transfer aliases and permissions, check the Transfer Aliases and Their Permissions check box.
5. Click OK.

Modifying Database Storage Allocations

As a database grows, you may need to increase the storage space allocated to it. After allocating storage space, you cannot de-allocate it, you can only add to it.

Privileges	Only a System Administrator can change storage allocations.
------------	---

You can allocate additional space on the device where you made the initial allocation or on a different device. In Figure 9-3, database A is storing data on device C. It then allocates additional space on device C. Database B stores data on device D. It allocates additional space on device E.




Figure 9-3: Expanding a database

► **Note**

To reduce the size of a database, you must back up all the data, delete the database, recreate it with smaller storage allocations, and then load the data back in.

Examining Database Storage Allocations

To see the current storage allocation of a database, its transaction log, and any indexes it contains, display the Database property sheet.

The Devices tab shows how the total amount of space allocated to the database is distributed, in megabytes. The Usage tab shows how much space is currently used, how much is reserved but unused, and how much is free for database data, indexes, and the transaction log.

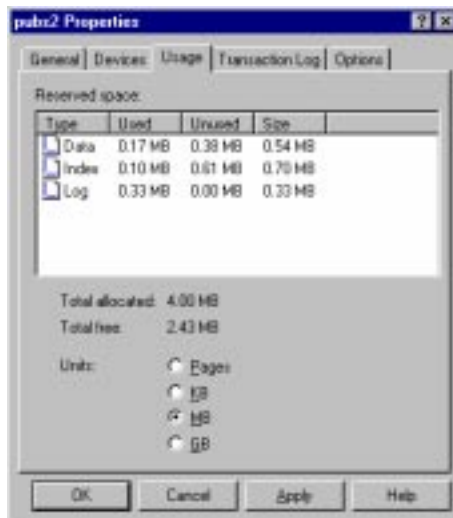


Figure 9-4: Database property sheet: Usage tab

Allocating Space on the Same Device

To expand a database by allocating space on the same device:

1. Open the Database property sheet as described in “Displaying and Modifying Database Properties” on page 9-4.
2. Select the Devices tab.



Figure 9-5: Database property sheet: Devices tab

- In the list of devices, select the device on which you want to allocate more space.
- Click the Edit button. The Device Size dialog box opens.



Figure 9-6: Device Size dialog box

- In the Add Space box, enter the amount of additional space you want to allocate from this device to the database.
- Click OK. Adaptive Server increases the allocation on the selected database device to the sum of the original allocation plus the additional allocation you specified.

Allocating Space on a New Device

To expand a database by allocating space on a new device:

1. Open the Database property sheet as described in “Displaying and Modifying Database Properties” on page 9-4.
2. Select the Devices tab.
3. Click the Add button. The Available Devices dialog box opens.



Figure 9-7: Available Devices dialog box

4. From the list of available devices, choose the device on which you want to allocate space to this database.
5. Specify whether the device is for Data or for Transaction Log.
6. In the Size box, specify the amount of space on the selected device that you want to allocate to this database.
7. Click OK.

When the original device is full, any additional database expansion takes place on the new device you specified.

Changing Log Buffer Size

The log buffer size determines the size of the I/Os that Adaptive Server can perform to the memory buffer in the cache designated for transaction log I/Os. (The cache used for transaction log I/Os is shown in the Log Cache box of the Database property sheet.) The default buffer size is 2K, indicating that Adaptive Server performs transaction log I/Os one data page at a time. You can change the log buffer size for the database's transaction log.

To change the log buffer size:

1. Open the Database property sheet as described in “Displaying and Modifying Database Properties” on page 9-4.
2. Select the Transaction Log tab.



Figure 9-8: Database property sheet: Transaction Log tab

3. Enter a value in the Log IO Size box. The value must 2, 4, 8, or 16.
4. Click OK.

Changing Database Options

Database options allow you to customize the behavior of user databases (you cannot change the option settings on the *master* database).

Privileges A database owner can change option settings.

A new user database takes its default option settings from the *model* database. To control the default settings of all new user databases, update the database option settings for the *model* database. Only the System Administrator can change the *model* database.

This section describes how to change database options and gives a brief description of each option. For more detailed option information, see the *System Administration Guide*.

To change database option settings for a database:

1. Open the Database property sheet as described in “Displaying and Modifying Database Properties” on page 9-4.
2. Select the Options tab.



Figure 9-9: Database property sheet: Options tab

3. Select or clear each option you want to change by clicking its check box. (Selecting an option sets its status to “ON” in Adaptive Server; clearing it sets the status to “OFF.”)
4. Click OK. Adaptive Server resets the options for the database.

Database Options

A description of each option follows:

- **Allow Select Into/Bulk Copy**—allows users to perform non-logged operations. Non-logged operations include `select into` for permanent tables, the bulk copy utility `bcp`, and the `writetext` utility.

You do not need to select this option to allow `select into` for temporary tables or to run `bcp` on a table with indexes, because inserts are logged.

Also, attempting to dump the transaction log in a database after unlogged changes have been made to the database with `select only` or bulk copy produces an error message instructing you to use `dump database` instead.

- **Truncate Log on Checkpoint**—truncates the transaction log (removes committed transactions) every time the database is checkpointed.

If you select this option, you cannot dump the transaction log. It may be useful to turn this option on during development work when backups of the transaction log are typically not needed.

► **Note**

If you select **Truncate Transaction Log on Checkpoint** for development purposes, clear it periodically and dump the transaction log. If you never dump the transaction log, it continues to grow, and eventually you run out of space in the database.

- **Checkpoint on Recovery**—Sets the database so that a checkpoint record is added to the database after it is recovered due to restarting Adaptive Server.

This checkpoint, which ensures that the recovery mechanism is not rerun unnecessarily, changes the sequence number on the database. If the sequence number on the secondary database has been changed, a subsequent dump of the transaction log from the primary database cannot be loaded into it.

Uncheck this option when an up-to-date copy of a database is kept. This prevents the secondary database from getting a checkpoint from the recovery process so that subsequent transaction log dumps from the primary database can be loaded into it.

- **Read Only**—prevents modification of any data in the database.
- **Only Usable By Owner**—restricts database access to the database owner.
- **Single User Mode**—allows only one user at a time to use the database.
- **Abort Transaction on Full Log**—decides how Adaptive Server treats active transactions when the database's log becomes critically low on space:
 - To cancel all user queries that need to write to the transaction log until space in the log has been freed, select this option.
 - To suspend transactions and awaken them when space has been freed, clear this option.
- **Allow Nulls by Default**—affects the ability of columns in newly created database tables to accept NULL values:
 - If you select this option, columns in newly-created tables allow null values unless the column definitions explicitly state "not null."
 - If this option is cleared, nulls are not allowed unless the column definitions explicitly permit them.
- **DDL in Transaction**—allows users to include DDL syntax within their transactions.

Generally, avoid using Data Definition Language commands inside transactions. For more information about this option, see *Adaptive Server Reference Manual*.

- **Free Space Accounting**—Determines whether the database enables free-space accounting and execution of threshold actions for the non-log segments.

Suppressing free-space accounting speeds recovery time because the free-space counts will not be recomputed for those segments. However, it disables updating the rows-per-page value stored for each table, so system procedures that estimate space usage may report inaccurate values.

Privileges	A System Security Officer can change the “No Free Space Accounting” option.
------------	---

- **Auto Identity**—automatically adds a 10-digit IDENTITY column in a new table when a user creates the table without specifying a primary key, a unique index, or an IDENTITY column.
- **Auto Identity Column in Non-Unique Indexes**— automatically includes an IDENTITY column in a table’s index keys, so that all indexes created on the table are unique.

Generating and Executing Database DDL Scripts

You can generate DDL to create a database or to create all its objects.

Privileges	Only a System Administrator can generate and execute database DDL scripts.
------------	--

Generating DDL to Create a Database

When you generate DDL for a database, you generate a script to create the database itself.

To generate database DDL:

1. Select the icon of the database for which you want to generate DDL.
2. From the File menu, choose Create Database DDL.

Sybase Central generates DDL to create the database and sends it to the tools console.

Shortcut



Click the right mouse button over the database icon. Choose Create Database DDL from the shortcut menu.

Generating DDL for the Objects in A Database

You can generate DDL for all the objects in a database (but not the database itself.) To generate DDL for database objects:

1. Select the icon of the database for which you want to generate object DDL.
2. From the File menu, choose All Objects DDL.

Sybase Central generates DDL for the database objects and sends it to the tools console.

Shortcut



Click the right mouse button over the database icon. Choose All Objects DDL from the shortcut menu.

Executing Database DDL Files

To execute database DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

Issuing a Database Checkpoint

A **checkpoint** is an automatic mechanism to guarantee that data pages changed by completed transactions are regularly written from the cache in memory to the database device. Each time it issues a checkpoint, Adaptive Server does the following:

- Freezes all current data modification transactions
- Writes pages that have been modified in memory, but not on disk, since the last checkpoint, to the database device
- Writes a checkpoint to the transaction log
- Unfreezes the current transactions

Through this mechanism, Adaptive Server regularly synchronizes the database and its transaction log, thereby shortening the recovery time in the event of a system failure.

Adaptive Server uses the *recovery interval* configuration variable to determine when to issue a checkpoint automatically. A System Administrator or database owner can also issue a checkpoint manually.

Privileges	A database owner can issue a manual checkpoint.
------------	---

To issue a manual checkpoint:

1. Select the icon of the database you want to checkpoint.
2. From the File menu, choose Checkpoint.

Shortcut



Click the right mouse button over the database icon. Choose Checkpoint from the shortcut menu.

Checking Database Consistency

Sybase Central lets you run the Database Consistency Checker (**dbcc**), a set of utility commands that check the logical and physical consistency of a database. Use the **dbcc** commands as follows:

- As part of regular database maintenance (periodic checks run by a System Administrator or database owner). These checks can detect, and often correct, errors before they affect a user's ability to use Adaptive Server.
- To determine the extent of possible damage after a system error has occurred.
- Before backing up a database.
- When you suspect that a database is damaged. For example, if using a particular table generates the message "Table corrupt," use **dbcc** to determine if other tables in the database are also damaged.

This section discusses the **dbcc** command options that deal with overall database consistency. It gives instructions for executing the commands and then describes each command.

Additional database command options enable you to check on the consistency of tables and indexes, and network I/O tasks.

Privileges	A System Administrator or database owner can execute database dbcc commands.
------------	---

To check database consistency:

1. Select the icon of the database you want to check.
2. From the File menu, choose Check Consistency. The Create DBCC command(s) wizard opens. Details of the various **dbcc** options are on the following pages.

Shortcuts



Click the right mouse button over the database icon and choose Check Consistency from the shortcut menu.

Consistency Check Options

The following `dbcc` command options are available for consistency checking at the database level:

- Check Overall Database
- Check Database Allocation
- Check System Catalog
- Repair and Drop Database

Overall Database Option

Overall Database corresponds to the `dbcc checkdb` command. This command checks each table in the selected database to see that:

- Index and data pages are correctly linked
- Indexes are in properly sorted order
- All pointers are consistent
- Information about internal allocation pages is synchronized with data rows

To skip checking nonclustered indexes on user tables, select the Ignore Non-clustered Indexes box. If you leave the box unselected, the `dbcc` command checks all indexes on all tables.

The report for each undamaged table shows the number of data pages used and the number of data rows for the table.

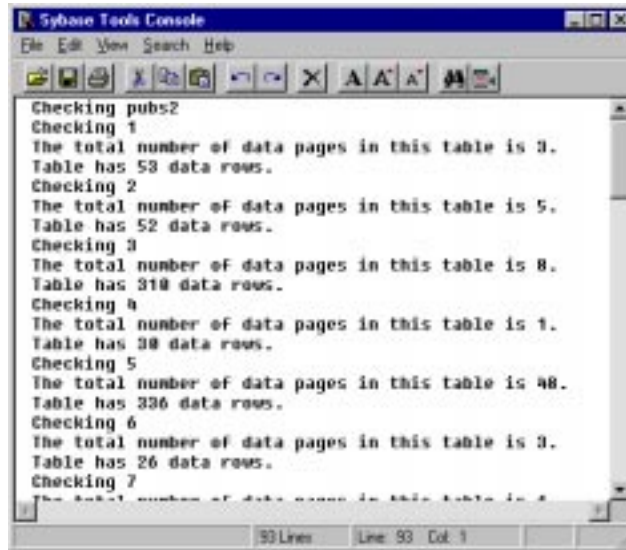


Figure 9-10: Check Overall Database output

Check Database Allocation

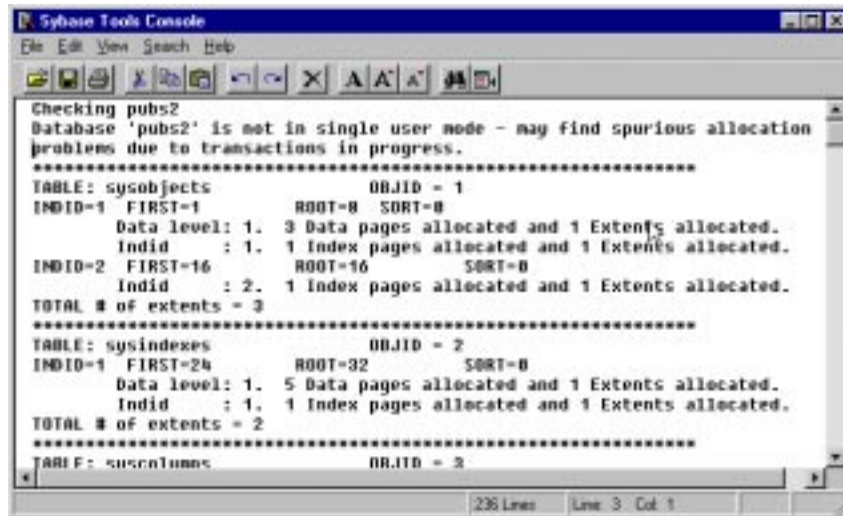
Check Database Allocation corresponds to the `dbcc checkalloc` command. This command checks the selected database to see that:

- All pages are correctly allocated
- No page is allocated that is not used
- No page is used that is not allocated

To fix allocation errors, check Fix Allocation Errors. Adaptive Server automatically places the database in single-user mode while executing `dbcc checkalloc` and then returns the database to multiple-user mode when processing is complete.

The report shows the amount of space allocated and used by each database table. The output consists of a block of data for each table, including the system tables, and the indexes on each table. For each table or index, it reports the number of pages and extents (8-page blocks of allocated space) used.

Because this `dbcc` option uses a large amount of overhead, run it in periods of low activity.



```

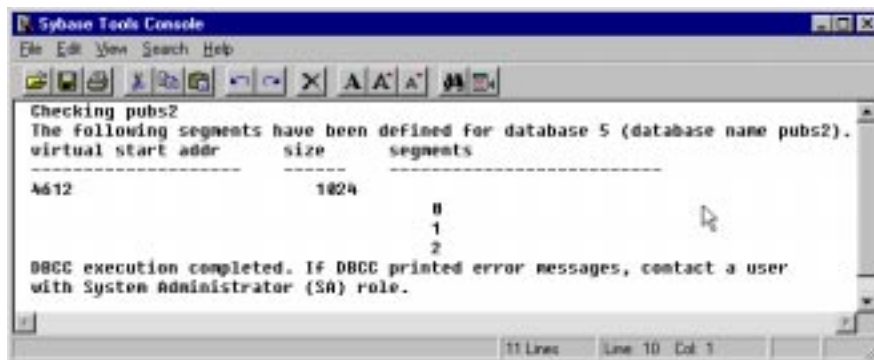
Checking pubs2
Database 'pubs2' is not in single user mode - may find spurious allocation
problems due to transactions in progress.
*****
TABLE: sysobjects          OBJID = 1
INDID=1  FIRST=1          ROOT=8  SORT=0
      Data level: 1.  3 Data pages allocated and 1 Extents allocated.
      Indid   : 1.  1 Index pages allocated and 1 Extents allocated.
INDID=2  FIRST=16        ROOT=16  SORT=0
      Indid   : 2.  1 Index pages allocated and 1 Extents allocated.
TOTAL # of extents = 3
*****
TABLE: sysindexes          OBJID = 2
INDID=1  FIRST=24        ROOT=32  SORT=0
      Data level: 1.  5 Data pages allocated and 1 Extents allocated.
      Indid   : 1.  1 Index pages allocated and 1 Extents allocated.
TOTAL # of extents = 2
*****
TABLE: syscolumns          OBJID = 3

```

Figure 9-11: Check Database Allocation output

Check System Catalog

Check System Catalog corresponds to the `dbcc checkcatalog` command. This command checks for consistency within and between the system tables found in a database. The report created by this option lists the segments defined for use by the database.



```

Checking pubs2
The following segments have been defined for database 5 (database name pubs2).
virtual start addr      size      segments
-----
4612                    1824
                        0
                        1
                        2
DBCC execution completed.  If DBCC printed error messages, contact a user
with System Administrator (SA) role.

```

Figure 9-12: Check System Catalog output

Repair and Drop Database

If a database is damaged, use Repair and Drop Database to repair the database and drop it. You cannot use this option to drop a database that is not marked suspect.

Setting a Database Online

Restoring a database or transaction log requires that the database in which it is performed be set off-line. Sybase Central automatically sets the database off-line. When the restore operation is complete, you must manually return the database to online status.

Similarly, it is possible that database administration activities performed outside Sybase Central can leave a database in the off-line state. If you receive an error message in Sybase Central indicating that the requested database is not online, you can set it back online by using a menu command.

To set a database online:

1. Select the icon of the database.
2. From the File menu, choose Set Online.

Shortcut



Click the right mouse button over the database icon. Choose Set Online from the shortcut menu.

10 Backing Up and Recovering Databases

What's in This Chapter

This chapter describes how to plan and perform database backup and recovery operations.

Preparing for Backup and Recovery

Backup and recovery (dump and load) are two of the most critical activities a system or database administrator performs. While Adaptive Server has automatic recovery procedures to protect you during power outages and computer failures, your best protection against media failure is regular and frequent database backups.

Before performing backup and recovery operations, you should do the following:

- Develop a backup and recovery plan.
- Make sure you can connect to the Backup Server from each Adaptive Server you administer.
- Decide on the backup media you will use and create dump devices that identify your physical backup media to Adaptive Server.

Developing a Backup and Recovery Plan

To develop a reliable backup and recovery plan, you must understand how Adaptive Server handles backup and recovery. To get started, read the chapter of the *System Administration Guide* on developing a backup and recovery plan. The *System Administration Guide* also has a detailed discussion of the commands and options used to perform backup and recovery operations.

The following are good practices to include in your backup and recovery procedures:

- Make frequent backups of the *master* database.
- Truncate the log of the *master* database frequently.
- Keep a current backup of the *model* database.

- Make frequent database and transaction log dumps for all databases.
- Keep statistics on how long it takes to back up and restore databases and how much space is required.

Identifying the Backup Server

Adaptive Server backups are performed by Backup Server, an Open Server™-based program that runs on the same host as Adaptive Server. Before backing up a database, make sure that the Backup Server for your Adaptive Server is up and running and that it is correctly identified in the interfaces file. Also, the login of the person who starts Backup Server (usually “sybase”) must have write permission for the physical dump device on which the backup is created.

Managing Dump Devices

Before backing up a database, you must have a **dump device** available to receive the backup. A dump device is a tape or disk device for storing a backup copy of a database or its transaction log.

This section describes how to create, delete, and examine the properties of a dump device.

Creating a Dump Device

Privileges	Only a System Administrator can create a dump device.
------------	---

Before you create a dump device, read “Creating Database Devices” in Chapter 8, “Managing Database Devices”.

To create a dump device:

1. Select the Dump Devices folder.
2. From the File menu, choose New; then choose Dump Device from the cascading menu.

The Create a New Dump Device wizard asks for the following information:

Table 10-1: Inputs to Create a New Dump Device wizard

Input	Description
Name	Logical name for the new dump device.
Path name of device	Enter the full path name of the physical device in your environment to which to map the logical name. On a UNIX host, this can be the name of a raw partition or operating system disk file.
Type	Specify a tape device or file device.
Capacity	For a tape device, specify the capacity.

Shortcuts



Select the Dump Device folder icon. Double-click the Add Dump Device icon.

Deleting a Dump Device

To delete a dump device, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Privileges	Only a System Administrator can delete a dump device.
------------	---

Examining Dump Device Properties

To display the properties of a Dump Device:

1. Select the icon of the dump device.
2. From the File menu, choose Properties. The Dump Device property sheet opens.



Figure 10-1: Dump Device property sheet

Shortcuts

	Double-click the dump device icon.
	Select the dump device icon. Then, select the Properties toolbar button.
	Click the right mouse button over the dump device icon. Choose Properties from the shortcut menu.

Backing Up a Database

Sybase Central helps you back up a database and its transaction log. The frequency of your database backups determines how little (or how much) work you will lose in the event of a media failure. This section presents some guidelines about when to schedule backups of user and system databases.

Guidelines for Backing Up User Databases

Back up a user database after the following events:

- After creating it, to develop a baseline backup. Thereafter, back up each database on a fixed schedule. Daily backups of the transaction log and weekly database backups are the minimum recommended. Many installations with large and active databases make transaction log dumps every half hour or hour and database dumps every day.
- After creating a new index. The transaction log records the create index transaction but does not log the filling of index pages with information.
- After executing the following unlogged transactions:
 - Unlogged writetext
 - select into on a permanent table
 - Bulk copy in “fast” mode (bcp into a table with no triggers or indexes)
- After truncating the transaction log without making a backup copy.

Guidelines for Backing Up the *master* Database

Back up the *master* database each time you change it. Each operation that affects disks, storage, databases, or segments makes changes to the system tables in the *master* database. For example, back up *master* after you:

- Create, modify, or delete a database, database device, database segment, dump device, or procedure
- Add or change disk mirroring specifications
- Change Adaptive Server configuration variables

Backup Procedure

Privileges	Operators can back up any database. Database Owners can back up databases they own.
------------	---

To back up a database:

1. Select the icon of the database you want to back up.
2. From the File menu, choose Backup. The Create a Backup Command wizard opens.

The Create a Backup Command wizard asks for the following information:

Table 10-2: Inputs to Create Backup Command wizard

Input	Description
Database name	Name of database or databases to back up.
Scope	Back up the entire database or just the transaction log.
Device name	Specify the dump device or devices for the backup.
Backup name	Optionally, specify a name for the backup.
Retention time	Optionally, override the default retention time.
Append	Specify whether or not to append to an existing backup.
Rewind	For tape devices, specify whether or not to rewind the tape after the backup.

Shortcuts



Click the right mouse button over the database icon. Choose Backup from the shortcut menu.

When a backup starts, it opens the tools console and sends messages from Backup Server to the console. For example, when a tape volume becomes full, Backup Server sends the following message to signal that a new tape volume must be mounted:

```
Mount the next volume to write.
```

When Backup Server needs operator assistance, Sybase Central displays the Backup Server Confirmation dialog box.



Figure 10-2: Backup Server Confirmation dialog box

The information to enter in the dialog box, and the command button to select, vary with the contents of the message. The following messages are the most common:

- If Backup Server reaches the end of the current tape volume, it sends the following message:
Mount the next volume to write.
- If you specified that the backup should be appended to an existing volume, and Backup Server cannot find the end-of-file mark on that volume, it sends the message:
Mount the next volume to search.

To Proceed with the Backup

1. Mount a new tape volume.
2. If the device on which the new tape volume is mounted differs from the device of the original tape volume, enter the new device name in the Dump Device box.
3. Click Proceed. Backup Server checks the tape and then continues the backup.

To Stop a Backup

If you want to stop the backup, click Abort. Backup Server stops the backup and returns control to Sybase Central.

To Get Information on Backup Server Messages

Sybase Central passes the information you supply in the Volume Change dialog box to Adaptive Server, which executes the `sp_volchanged` system procedure. For information about additional Backup Server messages and how to respond to them, see the description of `sp_volchanged` in the *Adaptive Server Reference Manual*.

Restoring a Database

Sybase Central helps you restore a database and its transaction log from database and transaction log backups. This section describes how to use the Database Restore dialog box to specify recovery options and execute a database recovery.

Deciding How to Restore

When your installation experiences media failure, the strategy for recovery depends on the state of your database and transaction log backups. Generally, you need to:

1. Restore the transaction log to capture the transactions that have occurred since the most recent transaction log backup.
2. Restore the database from the most recent complete database backup.
3. Apply each transaction log backup sequentially from the time of the most recent database backup through the time of media failure.

Recovery Considerations

In planning how to recover, consider the following:

- You can load a database backup into a pre-existing database, or you can create a new database with the `for load` option. For details on creating a database for load, see “Creating a Database” on page 9-2. With this method, you can move a database to a different Adaptive Server. However, you cannot load a database backup that was created on a different operating system or with an earlier release of Adaptive Server.

- To prevent users from making changes from the time you begin restoring a database until the time you finish applying the last transaction log backup, select the following database options in the Database Options dialog box:

- Single-User Mode
- Read Only
- Only Usable By Owner

Clear the Checkpoint On Recovery option.

When recovery is complete, clear these options to allow users to resume updating the database.

- For additional information about recovery, see *System Administration Guide* and *Sybase Troubleshooting Guide*.

Privileges	Operators can restore any database. Database Owners can restore databases they own.
------------	---

To restore a database:

1. Select the icon of the database you want to restore.
2. From the File menu, choose Restore. The Create a Restore Command wizard opens.

The Create a Restore Command wizard asks for the following information:

Table 10-3: Inputs to Create a Restore Command wizard

Input	Description
Database name	Name of database or databases to restore.
Scope	Restore the entire database or just the transaction log.
Device name	Specify the dump device or devices for the backup.
Backup name	Specify a backup name if the backup used one.
Rewind	For tape devices, specify whether or not to rewind the tape after the backup.

Shortcut



Click the right mouse button over the database icon, and choose Restore from the shortcut menu.

Responding to Volume Change Prompts

When you start a restore, Sybase Central opens the tools console. Messages from Backup Server are sent to the console. If Backup Server needs operator intervention, Sybase Central opens the Backup Server Confirmation dialog box (see Figure 10-2). For example, when Backup Server finishes reading the current tape volume, it sends the following message to signal that a new tape volume must be mounted:

```
Mount the next volume to read.
```

If you are prompted to change a tape volume:

1. Mount the new tape volume.
2. If you are using a different device, change the device name in the Device Name box.
3. Click Proceed.

Stopping A Restore

To stop a restore operation, click Abort in the Backup Server Confirmation dialog box. Backup Server stops the restore and returns control to Sybase Central.

Getting Information About Backup Server Messages

Sybase Central passes the information you supply in the Volume Change dialog box to Adaptive Server, which executes the `sp_volchanged` system procedure. For information about additional Backup Server messages and how to respond to them, see the description of `sp_volchanged` in the *Adaptive Server Reference Manual*.

11

Managing Segments

What's in This Chapter

After creating database storage devices, you can create user databases and allocate them on database devices. You can also create database segments, which enable you to control the placement of objects on the devices. This chapter describes how to create, extend, delete, and add thresholds to database segments.

Database Segments

Segments are named subsets of the database devices on which a particular database is stored. A segment is a label that points to one or more database devices. Within a particular database, you can define segments for the database devices already allocated to that database.

Each database can contain up to 32 segments. Every database contains segments called *system*, *logsegment*, and *default*. You can define additional segments, as needed.

Defining segments for a database lets you control placement of the objects that consume the most storage—tables and indexes. This gives you several performance and control advantages:

- By placing large tables on segments that span multiple physical devices, you can increase I/O throughput.
- By separating tables and their nonclustered indexes on different physical devices, you can also increase I/O throughput.
- By placing a table on a segment of a specific size, you can control space usage, since a table cannot grow larger than its segment allocation.

To monitor the use of space on a segment so that you can take action before a segment becomes full, you can define **threshold** values. Thresholds allow you to automatically trigger actions such as notification or dumping a transaction log when a segment is filled to the level of the threshold.

Creating a Segment

Before you create a segment, the database device on which you want to create it must exist and the database you create it for must have space allocated on the device.

Privileges	A database owner can create a segment.
------------	--

To create a segment:

1. Select the Segments folder in the database on which you want to create the segment.
2. From the File menu, choose New; then, choose Segment from the cascading menu. The Create New Segment wizard opens.

The Create New Segment wizard asks for the following information:

Table 11-1: Inputs to Create New Segment wizard

Input	Description
Name	Name for the new segment.
Database device	Device on which to create the segment.

Shortcuts



Select the Segments folder icon in the appropriate database. In the right pane, double-click the Add Segment icon.

Deleting a Segment

To delete a segment, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying and Modifying Segment Properties

Sybase Central lets you examine the mapping of database segments to database devices and the distribution of database objects on those segments. You can also extend an existing segment onto additional database devices. All of these activities are available through the Segment property sheet.




To display the properties of a segment:

1. Select the icon of the segment you want to display.
2. From the File menu, choose Properties.



Figure 11-1: Segment property sheet

Shortcuts

	Double-click the segment icon.
	Select the segment icon. Then select the Properties toolbar button.
	Click the right mouse button over the segment icon. Choose Properties from the shortcut menu.

Extending a Segment

Privileges A database owner can extend a segment.

To extend a segment onto another database device on which the database has been allocated:

1. Open the Segment property sheet as described in “Displaying and Modifying Segment Properties” on page 11-3.
2. Select the Devices tab.



Figure 11-2: Segment property sheet: Devices tab

3. Click the Add button. The Add Database Device to Segment dialog box opens.



Figure 11-3: Add Database Device to Segment dialog box

4. Select the devices you want to add to the segment.
5. Click OK.

Displaying Objects Stored on a Segment

You can specify the segment on which you want to place a table or index. The Segment property sheet lists the indexes and tables that use the segment. To display objects on a segment:

1. Open the Segment property sheet as described in “Displaying and Modifying Segment Properties” on page 11-3.
2. Select the Contains tab.

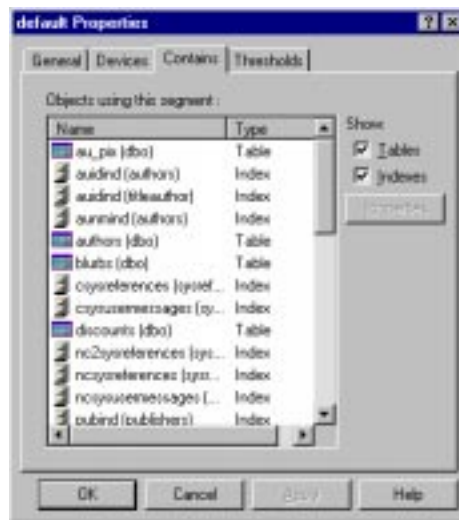


Figure 11-4: Segment property sheet: Contains tab

This tab lists the indexes and tables whose current growth is on the segment. Use the Indexes and Tables checkboxes under the Show: label to specify whether the list box displays indexes, tables, or both.

Using Segment Thresholds

Thresholds provide a way to monitor and preserve free space in a database. A threshold is a value specified as a number of pages of free space. Each threshold value is associated with a stored procedure. (For example, you can associate a threshold with a procedure that dumps the transaction log.) When free space on a segment falls below the threshold value, the stored procedure automatically executes.

You can define thresholds on segments that hold either data or transaction logs. A database can have as many as 256 thresholds. Using Sybase Central, you can create, delete, or modify a threshold, and you can create the stored procedure that executes when a threshold is reached.

Privileges	A System Administrator or database owner can create, delete, or modify a threshold.
------------	---

Last-Chance Threshold

If a transaction log is on its own segment (that is, if a segment is mapped to a device on which a database has allocated log-only space), Adaptive Server automatically creates a **last-chance threshold** on the segment, along with any other thresholds that you add. This threshold is an estimate of the number of free log pages that would be required to back up the transaction log. Adaptive Server automatically adjusts the last-chance threshold as you allocate more space to the log segment.

Hysteresis Value

To keep threshold procedures from firing continuously in response to minor fluctuations in space usage, Adaptive Server uses the Tolerance (Hysteresis) value. This server-wide value, stored in the system variable `@@thresh_hysteresis`, is specified in data pages. Before a threshold that has been activated and has triggered a stored procedure can be activated again, the space usage on the segment must decrease by at least the hysteresis value. Similarly, any two thresholds on a segment must be at least two hysteresis values apart. For example, if one threshold is set at .2MB and the hysteresis value for the Adaptive Server is 64 pages (.128MB), the closest you could place the next threshold would be .2+ (2*.128), or .456MB.

Creating a Threshold

Before you create a threshold, you must create the stored procedure that will run when the threshold is reached.

To create a threshold:

1. Select the segment on which you want to create the threshold.
2. Open the Segment property sheet as described in “Displaying and Modifying Segment Properties” on page 11-3.
3. Select the Thresholds tab.

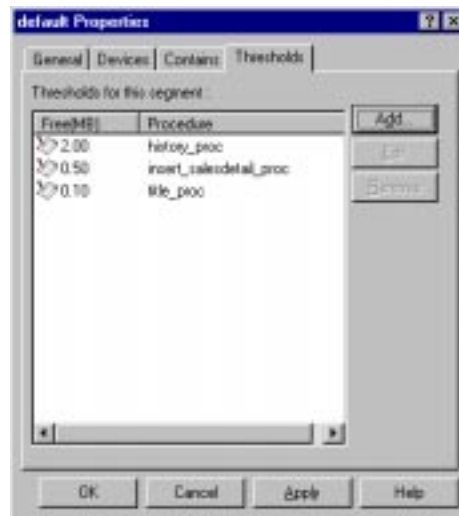


Figure 11-5: Segment property sheet: Thresholds tab

4. Click the Add button. The Add New Threshold dialog box is displayed.

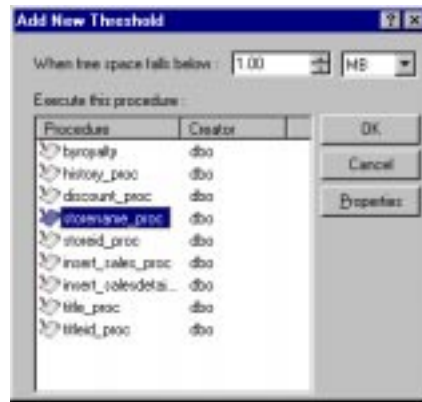


Figure 11-6: Add New Threshold dialog box

5. Set the free space value in the When Free Space Falls Below box.
6. Select a procedure from the list.
7. Click OK.

Modifying a Threshold

In thresholds other than the last-chance threshold, you can change the free space value and the stored procedure with which the threshold is associated. In a last-chance threshold, you can change only the stored procedure. Adaptive Server Enterprise automatically maintains the free space value. To modify a threshold:

1. Display the Thresholds tab of the Segment property sheet as described in “Creating a Threshold” on page 11-8.
2. Select the threshold you want to change.
3. Click the Edit button. The Edit Threshold dialog box is displayed.

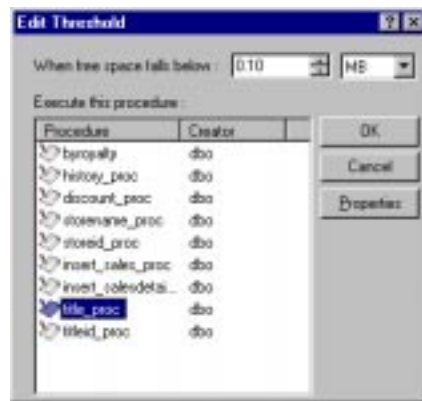


Figure 11-7: Edit Threshold dialog box

4. Change the threshold by editing the amount in the When Free Space Falls Below box or selecting a different procedure, or both.
5. Click OK to update the values in the Thresholds list.
6. Click Apply.

Deleting a Threshold

To delete a threshold:

1. Display the Thresholds tab of the Segment property sheet as described in “Creating a Threshold” on page 11-8.
2. Select the threshold you want to delete.

► **Note**

You cannot delete a last-chance threshold.

3. Click Remove.
4. Click Apply.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

12 Controlling Access to Adaptive Server

What's in This Chapter

Access to Adaptive Server, databases, and database objects is controlled through privileges that users receive automatically and through permissions that are assigned explicitly. Automatic privileges come with the roles associated with the user's login or come with database or object ownership. Explicit permissions are granted by authorized users.

For access to Adaptive Server, a Sybase Central user must have an Adaptive Server login. For access to a database or database objects, a Sybase Central user must be a user in the database. Unless a database user has automatic privileges in the database because of role assignment, group membership, or database or object ownership, the user must be granted explicit permissions. Therefore, to set up a new Sybase Central user to use a database, do the following:

- Create a new Adaptive Server login for the user and, optionally, assign roles or aliases to the login.
- Add the user to the database, and, optionally, add the user to a group.
- Grant permissions to the user to use commands and to use database objects.

► *Note*

Adaptive Server also permits access to a database to "guest" users. For information about guest users and for more information about controlling access to a database, see *System Administration Guide*.

This chapter describes how to control access through roles, shared logins (aliases), command permissions, and object permissions. Table 12-3 on page 12-9 summarizes permissions.

Permissions and Sybase Central

This chapter describes permissions and roles for standard command line access to Adaptive Server. Using Sybase Central to manage Adaptive Server may require permissions in addition to those listed in Table 12-3. This is because obtaining the information required to fill the fields in a property sheet can require permissions beyond those required for a particular task.

An indication that you have insufficient permission to manage an object is an error message indicating that select permissions have been denied on an object when you try to open a property sheet. If you receive this message, the property sheet can still open, but some information may be unavailable.

About Access Privileges and Permissions

Access privileges and permissions allow a Sybase Central user to create, view, and modify Adaptive Server and database objects. A user has access to commands and objects depending on the following:

- Adaptive Server role
- Status as a user (database owner or object owner)
- Group membership
- Specific command permissions granted
- Specific object permissions granted

Role assignment takes place at the Adaptive Server login level. Users have special user status when they become owners of a database or database object. Authorized Sybase Central users can grant command and object permissions to a database user directly, to a group to which the user belongs, to a role, or to another user whose login (alias) the user can use. This section describes access based on each of these categories of privilege or permission.

Access Based on Role

You can control user access to Adaptive Server through the use of system roles and user-defined roles. System roles are created by Adaptive Server and have a set of privileges that you cannot change. User-defined roles are created by the System Security Officer. Privileges are assigned explicitly to user-defined roles just as they are for users and groups.

You can create hierarchies of roles that inherit permissions from their members. You can also specify that roles are mutually exclusive. Mutual exclusivity is often used to prevent collaboration between job-related capabilities. For example, your business model might specify that separate transactions should be required to initiate a payment and to authorize a payment, and that no single individual should be capable of executing both transactions.

The next section describes the system roles. For more information about roles, see Chapter 15, “Managing Roles” and *System Administration Guide*.

System Roles

System roles help with the division of administrative responsibilities. A user can have multiple roles simultaneously. The roles are as follows:

- **System Administrator**—a user with the System Administrator role (*sa_role*) performs administrative tasks and has privileges for almost every Sybase Central function. Privileges acquired by a user with the System Administrator role are:
 - Managing disk space
 - Creating, deleting, and modifying logins
 - Granting or revoking the System Administrator role
 - Creating remote servers
 - Creating user databases and transferring ownership of them
 - Creating and managing the membership of groups
 - Granting permissions to Adaptive Server users
 - Executing the `dbcc` utility
 - Changing Adaptive Server configuration parameters

- Shutting down Adaptive Server or killing Adaptive Server processes
- Monitoring recovery

Additionally, Adaptive Server treats users with the System Administrator role as database owners in any database.

- **System Security Officer**—a user with the System Security Officer role (*sso_role*) performs security-related tasks. Privileges acquired by a user with the System Security Officer role are:
 - Creating and modifying logins and passwords
 - Setting the Adaptive Server password expiration interval
 - Setting the remote server option that controls whether a connection is maintained in the absence of login activity
 - Creating user-defined roles
 - Granting and revoking roles
 - Assigning aliases
- **Operator**—a user with the Operator role (*oper_role*) can back up and restore all databases and transaction logs.

To assign roles to a login, use the Roles tab of the Login property sheet. For details, see “Assigning Roles to a Login” in Chapter 13, “Managing Logins”.

Access Based on Special User Status

Some access privileges apply to users who assume a special status because they own a database or database object.

- **Database Owner**—users who own databases have authority over the databases they own. Database Owners have full permissions on objects inside their databases. Within their own databases, database owners can:
 - Issue a database checkpoint
 - Execute database-level `dbcc` commands
 - Delete a database
 - Create database users
 - Create database objects
 - Grant and revoke permission to create tables, views, defaults, procedures, and rules

- Grant and revoke permissions on database objects
- Back up and load the database and its transaction log

The creator of a database is its initial owner. Only a user with the System Administrator role can create a database. The System Administrator can then transfer database ownership to another user. See Chapter 9, “Managing Databases.”

- **Object owner**—a user who creates a database object owns the object and automatically has all object permissions on it. For a description of object permissions, see “Access Based on Object Permissions” on page 12-6. Users other than the object owner, including the owner of the database, must explicitly be granted each type of permission on the object. Along with object permissions, the owner of an object has the ability to:
 - Grant and revoke object permissions
 - Delete the object

The owner of a database table has the following additional permissions:

- Adding columns to the table
- Creating indexes and triggers on the table
- Executing the `update statistics` command

To become an object owner, a user:

- Must be the owner of a database, or
- Must be granted permission to create an object by a database owner or a user with the System Administrator role.

Access Based on Group Membership

Users who are members of a group take on the permissions explicitly granted to the group. Group permissions on an object override permissions granted to individual users on the same object.

Access Based on Command Permissions

To create databases or database objects, users must either have privileges based on role, special user status, or group membership, or be granted explicit permission to execute the corresponding create command.

A user with the System Administrator role can grant or revoke permission to create a database.

A database owner can grant to and revoke from other users or groups the permission to use the commands that create database tables, defaults, rules, procedures, and views.

In Sybase Central, use the following dialog boxes to grant or revoke command permissions:

- To grant permissions to or revoke permissions from a user, use the Command Permissions tab of the User dialog box. For details, see “Granting and Revoking Command Permissions for a User” on page 14-8.
- To grant permissions to or revoke permissions from a group, use the Command Permissions tab of the Group dialog box. For details, see “Granting and Revoking Command Permissions for a Group” on page 14-18.

Access Based on Object Permissions

To obtain access to database tables, views, and stored procedures, users must either have privileges based on role, special user status, or group membership, or be granted explicit permission for each type of access. Database and object owners can grant or revoke permission on objects they own.

For each permission granted, the grantor can specify that the recipient can grant the permission to another user. This form of granting permission is called granting **with grant**. When revoking permission, the revoker can specify that the permission be revoked from all users to whom the recipient granted it. In Sybase Central, this form of revoking permissions is called revoking **with cascade**.

You can use Sybase Central to grant or revoke the following object permissions:

- Select—select data from a table or view
- Insert—insert a row in a table or view
- Delete—delete a row in a table or view
- Update—update a row in a table or view
- Execute—execute a procedure
- Reference—create referential constraints on a table or view

When assigning permissions, you can update users, groups, or roles, or you can update the object to which permissions apply. For example, after you create a new user, you can assign all of the user's permissions from the user's property sheet. When creating a new table, you can assign permissions to each authorized user from the Table dialog box.

Table 12-1 shows which dialog boxes to use for granting or revoking object permissions.

Table 12-1: Object permission property sheet tabs

Method for assigning permissions	Property Sheet tab
Update a user	User property sheet Object Permissions tab
Update a group	Group property sheet Object Permissions tab
Update a role	Role property sheet Object Permissions tab
Update a table	Table property sheet Permissions tab
Update a view	View property sheet Permissions tab
Update a procedure	Procedure property sheet Permissions tab
Update an extended stored procedure	Extended stored procedure property sheet Permissions tab

Access Example

Table 12-2 illustrates several ways user *rcole* can be granted access to the *personnel* database to create a table.

Table 12-2: Granting access to create a table

Access test	Yes	No
Is <i>rcole</i> a user in <i>personnel</i> ?	Access granted	next test
Is <i>rcole</i> associated with a login that has the System Administrator role?	Access granted	next test
Is <i>rcole</i> the owner of <i>personnel</i> ?	Access granted	next test
Is <i>rcole</i> assigned a role that has been granted permission to create a table in <i>personnel</i> ?	Access granted	next test
Is <i>rcole</i> a member of a group that has been granted permission to create a table in <i>personnel</i> ?	Access granted	next test
Has <i>rcole</i> been granted permission to create a table in <i>personnel</i> ?	Access granted	No access

Summary of Access Privileges and Permissions

Table 12-3 summarizes the Adaptive Server protection system. The type of user listed as the default is the lowest level of user to which the permission is automatically granted. This user can grant the permission to other users or revoke it from other users, if it is transferable.

Table 12-3: Privilege and permission summary

Task	Default permissions	Can be granted or revoked?
Add an alias	System Administrator, Database Owner	no
Add columns to a table	Object owner	no
Assign a role to a login	System Security Officer	no
Bind a default to a table column or user datatype	Object owner	no
Bind an execution class to an execution object	System Administrator	
Bind a rule to a table column or user datatype	Object owner	no
Change database device default status	System Administrator	no
Change database options	System Administrator, Database Owner	no
Change database ownership	System Administrator, Database Owner	no
Change database storage allocations	System Administrator	yes (1)
Change group membership	System Administrator, Database Owner	no
Change login defaults or fullname	System Administrator	no
Change login password (5)	System Administrator, System Security Officer	no
Create a named cache	System Administrator	yes
Create a database	System Administrator	yes
Create a database device	System Administrator	no
Create a default	Database Owner	yes
Create a dump device	System Administrator	no
Create an engine group	System Administrator	no

Table 12-3: Privilege and permission summary

Task	Default permissions	Can be granted or revoked?
Create an extended stored procedure (create procedure)	System Administrator	yes
Create an execution class	System Administrator	no
Create a group	System Administrator, Database Owner	no
Create an index	Object owner	no
Create a login	System Security Officer	no
Create a procedure	Database Owner	yes
Create a remote server	System Administrator	no
Create a rule	Database Owner	yes
Create a segment	System Administrator, Database Owner	no
Create a table	System Security Officer (2), Database Owner	yes (2)
Create a proxy table	Database Owner	yes
Create a trigger	Object owner	no
Create a user	System Administrator, Database Owner	no
Create a user-defined datatype	Any user	
Create a user-defined role	System Security Officer	no
Create a view	Database Owner	yes
Define remote login mappings	System Security Officer	no
Delete an alias	System Administrator, Database Owner	no
Delete a database or dump device	System Administrator	no
Delete a database object	Object owner (3)	no
Delete an engine group	System Administrator	no
Delete an extended stored procedure	System Administrator	no
Delete an execution class	System Administrator	no
Delete a group	System Administrator, Database Owner	no

Table 12-3: Privilege and permission summary

Task	Default permissions	Can be granted or revoked?
Delete a login	System Administrator	no
Delete a remote server	System Administrator	no
Delete a segment	System Administrator	no
Delete a user	System Administrator, Database Owner	no
Delete a user-defined datatype	System Administrator, Database Owner	no
Delete a user-defined role	System Security Officer	
Disable disk mirroring	System Administrator	no
Dump a database	Operator, Database Owner	no
Dump a transaction log	Operator, Database Owner	no
Enable disk mirroring	System Administrator	no
Execute dbcc commands	Database Owner	no
Execute a procedure	Object owner (4)	yes
Extend a segment	System Administrator, Database Owner	no
Grant permission on a database object	Object owner	yes
Grant permission to create a database object	Database Owner	yes
Grant roles to logins	System Administrator, System Security Officer	no
Insert a row in a table	Object owner (3)	yes
Issue a database checkpoint	Database Owner	no
Issue an update statistics command	Object owner	no
Kill a user process	System Administrator	no
Load a database from a backup	Operator, Database Owner	no
Load a transaction log from a backup	Operator, Database Owner	no
Lock a login	System Administrator, System Security Officer	no
Move the transaction log to a different device	System Administrator, Operator	no
Place new table or view allocations on a segment	System Administrator, Database Owner, Object owner	no

Table 12-3: Privilege and permission summary

Task	Default permissions	Can be granted or revoked?
Reset Adaptive Server configuration options	System Administrator	no
Re-enable disk mirroring on an inactive mirror device	System Administrator	no
Revoke permission on an object	Object owner	no
Revoke permission to create an object	Database Owner	no
Shut down Adaptive Server	System Administrator	no
Unbind a default from a table column or user datatype	Object owner	no
Unbind a rule from a table column or user datatype	Object owner	no

(1) Transferred with database ownership.
(2) Public can create temporary tables, no permission required
(3) If a view, permission defaults to view owner
(4) Defaults to stored procedure owner
(5) All users can change their own passwords

Summary of Transferable Privileges and Permissions

The following table summarizes transferable access privileges, the administrative role or user who grants and revokes access, and the dialog box and tab to use for entering the information.

Table 12-4: Access granted through tabs and dialog boxes

Privilege	Grantor	Property sheet and tab where granted	
Change database storage allocations	System Administrator	Database (permission transfers with database ownership)	General
Create a default	Database Owner	Group	Commands
		User	Commands
Create a procedure	Database Owner	Group	Commands
		User	Commands
Create a rule	Database Owner	Group	Commands
		User	Commands
Create a table	Database Owner	Group	Commands
		User	Commands
Create a view	Database Owner	Group	Commands
		User	Commands
Delete a table or view	Object owner	Group	Object Permissions
		Table or View	Permissions
		User	Object Permissions
Define referential constraints on a table	Table owner	User	Object Permissions
		Group	Object Permissions
		Table	Permissions
Execute a procedure	Procedure owner	Group	Object Permissions
		Procedure	Permissions
		User	Object Permissions
Grant permission to transfer permission to other users	Any authorized grantor	Group	Object Permissions
		Table or View	Permissions
		User	Object Permissions

Table 12-4: Access granted through tabs and dialog boxes (continued)

Privilege	Grantor	Property sheet and tab where granted	
Insert rows in a table or view	Object owner	Group	Object Permissions
		Table or View	Permissions
		User	Object Permissions
Select rows in a table or view	Object owner	Group	Object Permissions
		Table or View	Permissions
		User	Object Permissions
Update rows in a table or view	Object owner	Group	Object Permissions
		User	Object Permissions
		Table or View	Permissions

13

Managing Logins

What's in This Chapter

Logins are the primary security mechanism for controlling access to Adaptive Server. This chapter describes how to create and manage logins. For more information about managing access to Adaptive Server, see Chapter 12, "Controlling Access to Adaptive Server."

Creating Logins

Creating an Adaptive Server login is the first step in giving a Sybase Central user access to Adaptive Server.

Privileges	Only a System Security Officer can create an Adaptive Server login.
------------	---

To create an Adaptive Server login:

1. Select the Logins folder.
2. From the File menu, choose New; then choose Login from the cascading menu. The Create a New Login wizard opens.

The Create a New Login wizard asks for the following information:

Table 13-1: Inputs to Create a New Login wizard

Input	Description
Name	Name of the new login.
Password	A password for the new login.
Default database	The database a user with this login is placed in upon logging in to Adaptive Server.
Default language	The language you want Adaptive Server to use to display messages.
Full name	The name of the person associated with this login.

Shortcuts



Select the Logins folder. In the right pane, double-click the Add Login icon.

Deleting A Login

Before you can delete a login, you must do the following for each user with which the login is associated:

- Revoke the user's command and object permissions with cascade. Revoking permissions with cascade revokes them from users to whom this user has granted them.
- Re-grant the revoked permissions to the other users, if appropriate.
- Delete objects owned by the user.
- Delete the user from all databases.

Because of all the preconditions for deleting a login, you may want to lock the login instead. See "Locking A Login" on page 13-4.

Privileges	Only a System Administrator can delete a login.
------------	---

To delete a login, follow the procedure for deleting objects as described in "Deleting an Object" on page 3-4.

Displaying Login Properties

Once a login is created, users with the System Security Officer role can lock or unlock the login or change the defaults, the password, and the user roles. System Security Officers can also assign additional users (aliases) to the login.

Privileges	A System Security Officer can modify an Adaptive Server login.
------------	--

To display the properties of a login:




1. Select the icon of the login you want to modify.
2. From the File menu, choose Properties.



Figure 13-1: Login property sheet

3. To change the login's name, enter the new name in the Full Name text box.
4. To change the database that the login points to automatically when connecting to Adaptive Server, select the new database name from the list in the Default Database text box.
5. To change the default language for the login on connecting to Adaptive Server, select the new language from the list in the Default Language box.
6. Click OK to change the login.

Shortcuts

	Double-click the login icon.
	Select the login icon. Then select the Properties toolbar button.
	Click the right mouse button over the login icon. Choose Properties from the shortcut menu.

Locking A Login

Locking an Adaptive Server login is a way to disable it without removing it from the *master* database. It may be easier to lock a login rather than delete it for the following reasons:

- You cannot delete the login if a user exists for that login in any database.
- You cannot delete a user from a database if the user owns any objects in the database or has granted permissions on objects to other users.
- Adaptive Server could compromise accountability by reusing a deleted login server user ID (suid) when the next login is created.
- You cannot delete the last remaining System Security Officer or System Administrator login.

Another common reason to lock a login is to implement roles. When you first install Adaptive Server, the default “sa” login has all privileges associated with every role. At many installations, after the “sa” login makes the initial role assignments, a user with the System Administrator or System Security Officer role locks the “sa” login, so that no single Sybase Central user has every available privilege and so that actions can be traced.

To lock or unlock a login:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Parameters tab.

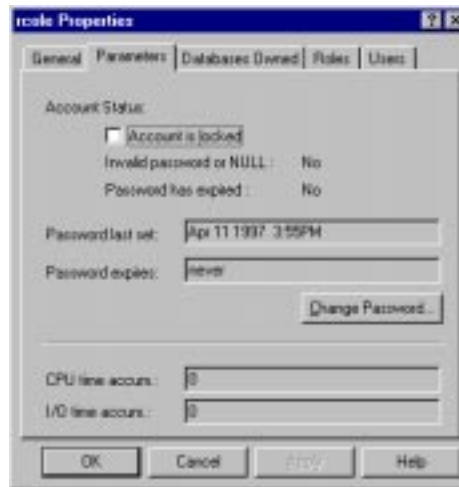


Figure 13-2: Login property sheet: Parameters tab

3. Check or clear the Account is Locked check box.
4. Click OK.

► **Note**

You can lock an account that is logged in—the user is not locked out until he or she logs out. You can lock an account that owns a database and objects in databases.

Changing a Login Password

To change the login password:



1. Select the login whose password you want to change.
2. From the File menu, choose Change Password. The Change Password dialog box opens.



Figure 13-3: Change Password dialog box

3. Enter your password in the Your Password box.
4. Enter a new password in the Login's New Password box.
5. Re-enter the password in the Confirm New Password box.
6. Click OK to change the password.

Shortcut

	Click the right mouse button over the login icon. Choose Change Password from the shortcut menu.
	Click the Change Password button on the Parameters tab of the property sheet.

Displaying Databases Owned by a Login

To display the databases a login owns:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Databases Owned tab.



Figure 13-4: Login property sheet: Databases Owned tab

Assigning Roles to a Login

For a complete discussion of roles, see “Access Based on Role” in Chapter 12, “Controlling Access to Adaptive Server” and Chapter 15, “Managing Roles.”

To assign a role to a login:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Roles tab (Figure 13-5.)



Figure 13-5: Login property sheet: Roles tab

3. Click the Add button. The Add Role to Login dialog box opens.



Figure 13-6: Add Roles dialog box

4. In the list, select the role you want to add to this login.
5. Click OK.
6. On the Roles tab, click Apply.

Setting Automatic Activation of A Role

If a role does not require a password, you can activate it automatically when a user logs in. If roles are mutually exclusive by activation, you can only activate one of them automatically.

To activate a role automatically at login:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Roles tab (Figure 13-5.)
3. In the list, select the role you want to activate automatically.
4. Click the Set Default button.
5. Click Apply.

Removing Automatic Activation of a Role

To remove automatic activation of a role:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Roles tab.
3. In the list, select the role for which you want to remove automatic activation.
4. Click the Unset Default button.
5. Click Apply.

Sharing Logins by Using Aliases

Using aliases enables more than one person to share identical privileges within a database. For example, to give several logins the automatic privileges associated with ownership of a database, you can associate those logins with the database owner, thus making each login an alias to the database owner. This section describes how to add aliases to a login and how to navigate to aliased users.

Privileges	A System Security Officer can update the aliases mapped to a login.
------------	---

Adding an Alias to a Login

To add an alias for a login:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Users tab.



Figure 13-7: Login property sheet: Users tab

3. Click the Add Alias button. The Add Alias dialog box opens.



Figure 13-8: Add Alias dialog box

4. Select the names of the users you want to alias to the login.

5. Click OK.

Deleting an Alias from a Login

To delete an aliased user from the login:

1. Display the Login property sheet, as described in “Displaying Login Properties” on page 13-2.
2. Select the Users tab.
3. Select the users whose aliases you want to delete.
4. Click the Remove button.
5. Apply the change.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

14 Managing Users and Groups

What's in This Chapter

This chapter explains how to create and manage users and groups. The information and procedures in this chapter are closely related to Chapter 12, “Controlling Access to Adaptive Server” and Chapter 15, “Managing Roles.”

Tasks involving users are:

- Creating and deleting users
- Modifying group membership and login aliases
- Displaying objects owned by a user
- Granting and revoking object permissions
- Granting and revoking command permissions

Tasks involving groups are:

- Creating and deleting groups
- Adding users to and deleting users from a group
- Granting and revoking object permissions
- Granting and revoking command permissions

Managing Users

The term user has a specific meaning in a database. To access a database, a login must be added as a user in the database. A common practice for setting users up in a database is to:

1. Add the user to the database.
2. Grant the user specific permissions on commands and objects.

Creating Users

Privileges	Database Owners can add and delete users in the databases they own.
------------	---

To create a user:

1. In a database, select the Users folder.
2. From the File menu, choose Create; then, choose User from the cascading menu. The Add a New User wizard opens.

The Add a New User wizard asks for the following information:

Table 14-1: Inputs to Add a New User wizard

Input	Description
Name	A name for the user. The name does not have to be the same as the login.
Login name	Login to which this user is assigned.
Group	Optionally, assign a group to the user. Default: public

► **Note**

A user can be a member of one assigned group or the default “public” group. The group permissions override the user’s permissions on the same objects.

Shortcuts



Select the Users folder. In the right pane, double-click the Add User icon.

Deleting a User

You cannot delete a user who owns objects. Since there is no command to transfer ownership of objects, you must delete objects owned by a user before you can delete the user. Also, you cannot delete a user who has granted permissions to other users without first revoking the permissions with cascade. If appropriate, re-grant the permissions to the other users.

Locking a login is a simple alternative to deleting a user. For details, see “Locking A Login” in Chapter 13, “Managing Logins”. If you decide to delete a user rather than lock a login, you must do the following before you can delete the user:

- Revoke the user’s command and object permissions with cascade.
- Re-grant the permissions to the other users, if appropriate.
- Delete the user’s objects.

To delete a user, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying User Properties

You can display and modify information about users and about their group membership, logins, permissions, and objects owned.

To display user properties:

1. Select the icon of the user.
2. From the File menu, choose Properties.



Figure 14-1: User property sheet

The General tab shows the user's name and login, as well as the group to which the user belongs.

Changing a User's Group

Privileges	Database Owners can modify a user's group and aliases in databases they own.
------------	--

To change a user's group membership:

1. Display the User property sheet, as described in "Displaying User Properties" on page 14-3.
2. On the General tab, select a group from the Group drop-down list.
3. Click Apply.

Displaying Objects Owned by the User

To see a list of the objects a user owns:

1. Display the User property sheet, as described in “Displaying User Properties” on page 14-3.
2. Select the Objects Owned tab.

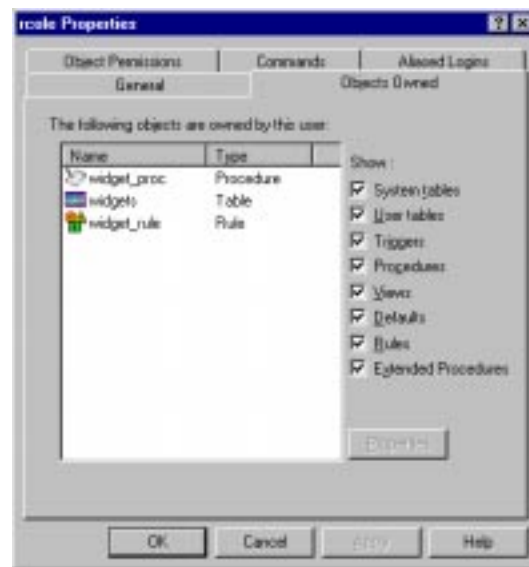


Figure 14-2: User property sheet: Objects Owned tab

3. To limit the list of displayed objects to specific object types, use the check boxes under the Show: label. By default all object types are selected. Clear the check box for a type of object to exclude it from the display.

Granting and Revoking Object Permissions for a User

Object permissions govern access to database objects. When assigning permissions to a user on an object, object owners can update either the user or the object. This section describes how to assign permissions by updating users.

Privileges	Object owners can assign privileges on the objects they own.
------------	--

To grant or revoke object permissions:

1. Display the User property sheet, as described in “Displaying User Properties” on page 14-3.
2. Select the Object Permissions tab.

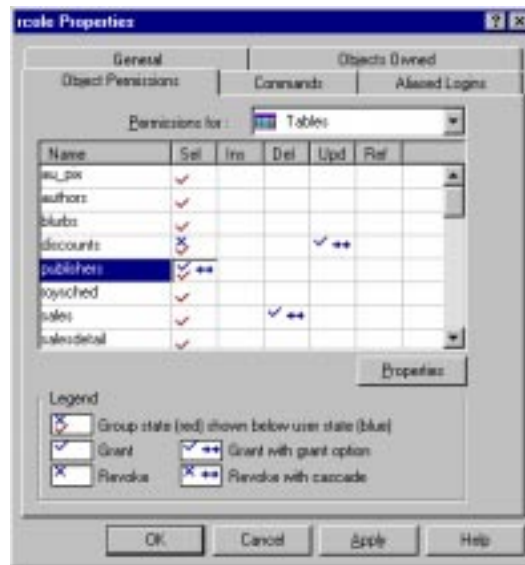


Figure 14-3: User property sheet: Object Permissions tab

3. Initially, the dialog box shows tables. You can also select procedures, extended stored procedures, or views from the Permissions For drop-down list.
4. To grant a permission on an object, click in the grid box for the object and type of permission. Continue clicking in the box to cycle through the options:
 - Grant
 - Grant with grant
 - Revoke
 - Revoke with cascade
 - No permissions

For stored procedures, you can grant execute permission. For tables and views, you can grant the following types of permissions:

- Select—permission to select a table or view
- Insert—permission to insert a row
- Update—permission to update a row
- Delete—permission to delete a row
- Reference—permission to create referential constraints

► **Note**

Sybase Central reports on permissions explicitly granted or revoked and those that users obtain by being members in a group. It does not show permissions associated with having an role, such as sa_role, associated with their logins.

5. Repeat Step 4 for each object on which you want to update permissions.

Shortcuts



Drag and drop the user icon onto a table to grant all permissions.



Right click a grid box. Select a permission from the popup menu.

Granting and Revoking Command Permissions for a User

Command permissions govern the ability to execute create commands for a database or database objects.

Privileges	System Administrators and database owners can grant command permissions. Only a System Administrator can grant permission to create a database. Database Owners can grant permission to create database objects in the databases they own.
------------	--

► **Note**

Permission to create a database applies only to users of the *master* database.

To manage command permissions:

1. Display the User property sheet, as described in “Displaying User Properties” on page 14-3.
2. Select the Commands tab.



Figure 14-4: User property sheet: Command Permissions tab

3. Assign or revoke a permission by clicking in the status box for the command you want to change. Each click changes the status to the next option in the cycle described in the Legend.
4. Click Apply to change the permissions in the database.

► **Note**

Sybase Central reports on permissions explicitly granted or revoked and those that users obtain by being members in a group. It does not show permissions associated with having an role, such as sa_role, associated with their logins.

Shortcuts



Right click a grid box. Select a permission from the popup menu.

Aliasing Users to Logins

Every user is created based on a login. However, you can alias a user to additional logins to share their privileges. To alias a user to a login:

1. Display the User property sheet, as described in “Displaying User Properties” on page 14-3.
2. Select the Logins tab.

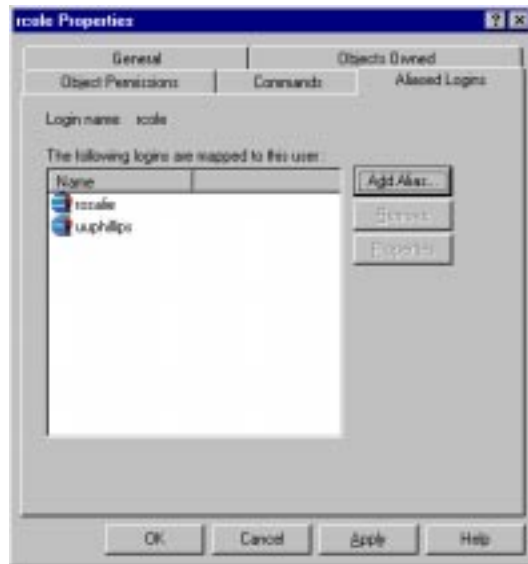


Figure 14-5: user property sheet: Logins tab

3. Click Add Alias. The Add Login Alias dialog box opens.

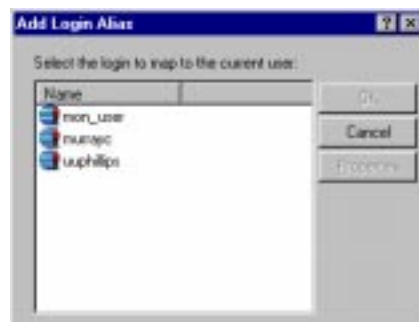


Figure 14-6: Add Login Alias dialog box

4. In the list, select the logins you want to alias to this user.
5. Click OK.
6. Click Apply.

Removing a Login Alias

To remove a login alias from a user:

1. Display the User property sheet, as described in “Displaying User Properties” on page 14-3.
2. Select the Logins tab.
3. In the list, select the logins you want to remove as aliases.
4. Click Remove.
5. Click Apply.

Creating and Executing DDL Scripts for Users

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

Managing Groups

Groups provide a convenient way to grant and revoke permissions to more than one user at once. Every user is a member of the “public” group and can be a member of one other group.

It may be convenient to create groups before adding users to a database so you can assign new users to existing groups.

Creating a Group

Privileges	Database Owners can create and delete groups in databases they own.
------------	---

To create a group:

1. In a database, select the Groups folder.
2. From the File menu, choose New Group. The Create Group wizard opens.

The Create Group wizard asks for the following information:

Table 14-2: Input to Create Group wizard

Input	Description
Name	Name of the new group.

Shortcuts



Select the Groups folder. In the right pane, double-click the Add Group icon.

Deleting a Group

You cannot delete a group that has members. Before deleting a group, you must first remove all of its users. See “Removing Users from a Group” on page 14-15.

To delete a group, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying Group Properties

You can display and modify information about groups, their membership, and permissions.

To display group properties:

1. Click the icon of the group you want to display.
2. From the File menu, choose Properties.



Figure 14-7: Group property sheet

Shortcuts



Select the group icon. Then select the Properties toolbar button.



Click the right mouse button over the group icon. Then, choose Properties from the shortcut menu.

Adding a User to a Group

Privileges	Database Owners can modify group membership in databases they own.
------------	--

To add a user to a group:

1. Display the Group property sheet as described in “Displaying Group Properties” on page 14-13.
2. Select the Users tab.



Figure 14-8: Group property sheet: Users tab

3. To add a user to the group, click Add. The Add User to Group dialog box is displayed.



Figure 14-9: Add User to Group dialog box

4. Select a user from the list.
5. Click OK.

Removing Users from a Group

Privileges	Database Owners can modify group membership in databases they own.
------------	--

To remove the users from a group:

1. Display the Group property sheet as described in “Displaying Group Properties” on page 14-13.
2. Select the Users tab.
3. In the list, select the users you want to remove from the group.
4. Click Remove.
5. Click Apply.

Granting and Revoking Object Permissions for a Group

Object permissions govern access to database objects. When assigning permissions to a group on an object, object owners can either update the group or the object. This section describes assigning permissions by updating database groups. For information about assigning permissions by updating database objects, see the description of each object in its respective chapter.

Privileges	Object owners can assign object permissions on groups they own.
------------	---

To grant or revoke object permissions:

1. Display the Group property sheet, as described in “Displaying Group Properties” on page 14-13.
2. Select the Object Permissions tab.

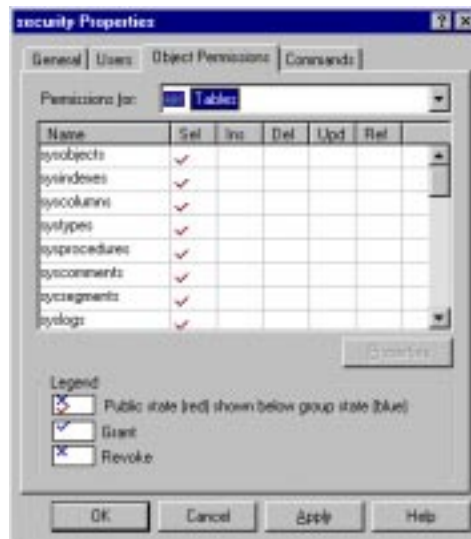


Figure 14-10: Group property sheet: Object Permissions tab

Initially, the tab shows tables. You can also select procedures, extended stored procedures, or views from the Permissions for drop-down list.

3. To grant or revoke a permission on an object, click in the grid box for the object and type of permission. Continue clicking in the box to cycle through the options:

- Grant
- Revoke
- No permissions

You can grant execute permission for stored procedures. For tables and views, you can grant the following types of permissions:

- Sel—permission to select a table or view
- Ins—permission to insert a row
- Upd—permission to update a row
- Del—permission to delete a row
- Ref—permission to create referential constraints

When you make a permission change for the group, the Changed column displays a check mark (✓) or deletes a check mark, indicating that you have altered the original settings.

► **Note**

Sybase Central reports on permissions explicitly granted or revoked and those that users obtain by being members in a group. It does not show permissions associated with having an role, such as sa_role, associated with their logins.

Shortcuts



Drag and drop the group icon onto the object icon.



Right click a grid box. Select a permission from the popup menu.

Granting and Revoking Command Permissions for a Group

Command permissions govern the ability to execute create commands for a database or database objects.

Privileges	Only a System Administrator can grant permission to create a database. Database Owners can assign command permissions to groups in databases they own.
------------	--

► **Note**

Permission to create a database applies only to users of the *master* database.

To manage command permissions:

1. Display the Group property sheet, as described in “Displaying Group Properties” on page 14-13.
2. Select the Commands tab.



Figure 14-11: Group property sheet: Command Permissions tab

3. Assign a permission by clicking in the status box for the command you want to change. Each click changes the status to the next option in the cycle described in the Legend.

4. Click Apply to change the permissions in the database.

► **Note**

Sybase Central reports on permissions explicitly granted or revoked and those that users obtain by being members in a group. It does not show permissions associated with having an role, such as sa_role, associated with their logins.

Shortcuts



Right click a grid box. Select a permission from the popup menu.

Creating and Executing DDL Scripts for Groups

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

15

Managing Roles

What's in This Chapter

Chapter 12, “Controlling Access to Adaptive Server” describes how to manage access to Adaptive Server using roles and permissions. You can enhance your ability to manage access to Adaptive Server through the use of user-defined roles. Most properties of user-defined roles also apply to system roles. This chapter describes the following:

- Creating a user-defined role
- Setting passwords for roles
- Assigning roles to a login
- Creating hierarchies of roles
- Specifying mutual exclusivity for roles
- Setting permissions for roles

For more information about user-defined roles, see *Adaptive Server System Administration Guide*.

Privileges	Only the System Security Officer can create and manage roles.
------------	---

► **Note**

For the System Security Officer to create and manage roles or for a login to activate roles using Sybase Central, the System Security Officer or login must have select permission on the *sysroles*, *sysserverroles*, and *sysloginroles* tables. One way to accomplish this is to give the *public* group in the *master* table these permissions.

Creating a User-defined Role

To create a user-defined role:

1. Select the Roles folder.
2. From the File menu, choose New; then choose User-defined Role from the cascading menu. The Create a New Role wizard opens.

The Create a New Role wizard asks for the following information:

Table 15-1: Inputs for Create a New Role wizard

Input	Description
Name	Name of role
Password	Optionally, provide a password for this role.

Shortcuts



Select the Roles folder icon. In the right pane, double-click the Add Role icon.

Deleting a User-defined Role

Before you delete a role, you need to remove all assignments to logins. See “Removing A Role Assignment from A Login” on page 15-6.

To delete a user-defined role, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

You cannot delete a system role.

Displaying Role Properties




To display the properties of a role:

1. Select the icon of the role you want to display.
2. From the File menu, choose Properties.



Figure 15-1: Role property sheet

Shortcuts

	Double-click the role icon.
	Select the role icon. Then select the Properties toolbar button.
	Click the right mouse button over the role icon. Choose Properties from the shortcut menu.

Setting A Password For A Role

A user cannot use a role until it is activated. You can require that a user enter a password before activating a role. This increases the level of security for the role.

To set a password for a role:

1. Display the Role property sheet as described in “Displaying Role Properties” on page 15-3.
2. If the role has a password, click the Change Password button on the General tab, in the Password group box. If the role does not have a password assigned, check the Required box. In either case, the Change Password for a Role dialog box opens.



Figure 15-2: Change Password for a Role dialog box

3. Enter a password for the role in the New Role Password box.
4. Confirm the password by entering it in the Confirm New Password box.
5. Click OK.

Assigning A Role to A Login

You must assign a role to a login for it to be in effect.

To assign a role to a login:

1. Display the Role property sheet as described in “Displaying Role Properties” on page 15-3.
2. Select the Logins tab.

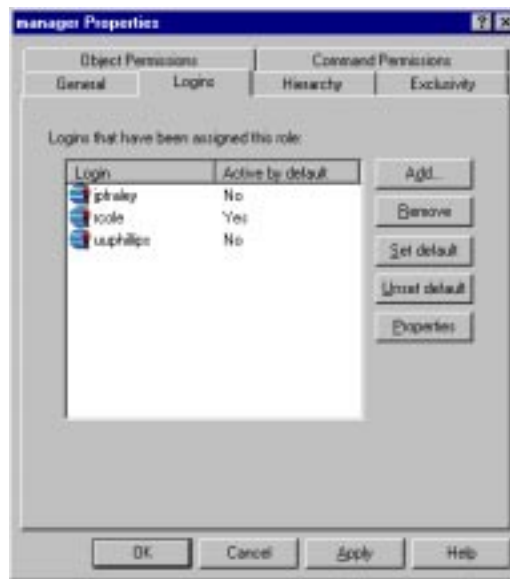


Figure 15-3: Role property sheet: Logins tab

3. Click the Add button. The Add Role to Login dialog box opens.



Figure 15-4: Add Role to Login dialog box

4. From the list, select the login to which you want to assign this role.
5. Click OK.

You can also assign roles to logins on the Login property sheet.

Removing A Role Assignment from A Login

To remove a role assignment from a login:

1. Display the Role property sheet as described in “Displaying Role Properties” on page 15-3.
2. Select the Logins tab.
3. In the list, select the login that you do not want to be assigned to this role.
4. Click the Remove button.

Activating and Deactivating A Role

Before you can use a role, you must activate it. By default, system roles are active at login. If a role requires a password, you cannot set it to be activated at login.

Privileges	To activate a role, a login must have select permission on the <i>sysroles</i> , <i>sysserverroles</i> , and <i>sysloginroles</i> tables.
------------	--

To activate a role:

1. Select the icon of the Adaptive Server installation you are connected to.
2. From the File menu, choose Role Activation. The Role Activation dialog box opens. It lists all roles that are assigned to your login.



Figure 15-5: Role Activation dialog box

3. In the list, select the role you want to activate. If the role requires a password, the Password box becomes active.
4. If the role requires a password, enter the password.
5. Click the Activate Role button.

Shortcut



Click the right mouse button over the Adaptive Server icon. Choose Activate Roles from the shortcut menu.

Activating A Role By Default

If a user-defined role does not require a password, you can activate it automatically when a user logs in. To activate a role automatically:

1. Display the Role property sheet as described in “Displaying Role Properties” on page 15-3.
2. Select the Logins tab.
3. Select the login for which you want to activate the role automatically.
4. Click Set Default.
5. Click Apply.

Removing Automatic Role Activation

To remove automatic activation of a role (that is, to require a login to manually activate the role):

1. Display the Role property sheet as described in “Displaying Role Properties” on page 15-3.
2. Select the Logins tab.
3. Select the login for which you want to remove automatic activation.
4. Click Unset Default.
5. Click Apply.

Deactivating A Role

To deactivate a role:

1. Select the icon of the Adaptive Server installation you are connected to.
2. From the File menu, choose Role Activation.
3. In the list, select the role you want to deactivate.
4. Click the Deactivate Role button.

Specifying Mutual Exclusivity

The mutual exclusivity feature lets you ensure that a login with a given role either cannot be assigned to another role or cannot use another role at the same time. Defining roles as mutually exclusive lets you enforce a policy of separation of duties. Mutual exclusivity is often used to prevent collaboration between job-related capabilities. For example, your business model might specify that separate transactions should be required to initiate a payment and to authorize a payment, and that no single individual should be capable of executing both transactions.

The System Security Officer can define mutual exclusivity at the membership level or at the activation level. When two roles are defined as being mutually exclusive at the membership level, the System Security Officer cannot grant the same login to both roles. When two roles are defined as being mutually exclusive at the activation level, the System Security Officer can grant a user both roles, but the user cannot activate both roles at the same time.

To specify mutual exclusivity for a role:

1. Display the property sheet of the role as described in “Displaying Role Properties” on page 15-3.
2. Select the Exclusivity tab.

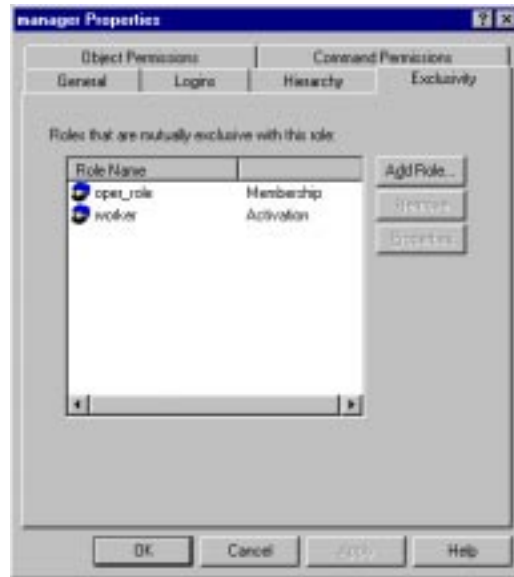


Figure 15-6: Role property sheet: Exclusivity tab

3. Click Add Role. The Add Mutually Exclusive Roles dialog box opens.

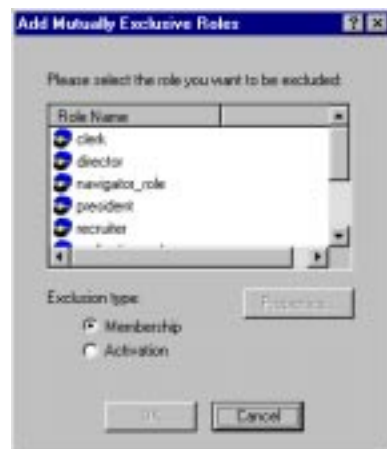


Figure 15-7: Add Mutually Exclusive Roles dialog box

4. In the list, select the role you want to make mutually exclusive with the role you are editing.
5. For Exclusion Type, select Membership or Activation. Membership means that no login can be assigned both of the roles. Activation means that no user can activate both roles at the same time.
6. Click OK. The new exclusivity mapping is added to the list on the Exclusivity tab.
7. Click Apply.

Specifying A Role's Hierarchy

A role hierarchy defines roles that contain or are contained by other roles, with the permissions and restrictions associated with the contained or containing roles.

Roles can have overlapping responsibilities and privileges. Users with different roles might need to perform common operations, since many organizations have functions performed by all employees. In such cases, it is more efficient and simpler to organize roles into hierarchies, thereby containing roles with the needed privileges rather than defining roles with the same set of privileges for different users.

To specify a hierarchy for a role:

1. Display the property sheet for the role as described in "Displaying Role Properties" on page 15-3.
2. Select the Hierarchy tab.



Figure 15-8: Role property sheet: Hierarchy tab

3. Click Assign.

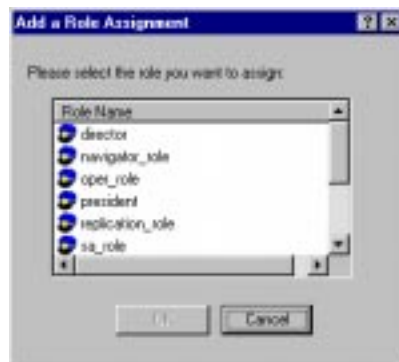


Figure 15-9: Add a Role Assignment dialog box

4. In the Add A Role Assignment dialog box, select the role that you want to contain in the role you are editing.

5. Click OK. The contained role is listed in the Hierarchy tab. If any roles are contained by the role you just added, they are displayed in a hierarchical list.

► **Note**

You can assign a role only to the root role in the Hierarchy tab. You cannot assign roles to roles contained by the root role. To assign roles to contained roles, open the contained role's property sheet. See "Example of Creating a Role Hierarchy" on page 15-13 for an illustration of opening a contained role's property sheets.

Example of Creating a Role Hierarchy

The following example demonstrates creating a four-level role hierarchy starting from the top and using the Properties button on the Hierarchy tab to open property sheets for successive contained roles. It also demonstrates how to expand the hierarchy view.

The example uses four roles: *president*, *director*, *manager*, *worker*. In this simple hierarchy, we assume that each role has specific permissions plus all the permissions of subordinates. Therefore, we will create a hierarchy in which each role contains the subordinate roles.

Figure 15-10 illustrates the results of this example. The procedure is as follows:

1. Open the property sheet for the role *president*.
2. On the Hierarchy tab, assign the role *director*.
3. In the list, select *director* and click Properties. The property sheet for *director* opens.
4. On the Hierarchy tab for *director*'s property sheet, assign the role *manager*.
5. In the list, select *manager* and click Properties.
6. On the Hierarchy tab for *manager*'s property sheet, assign the role *worker*.

You now have the series of opened property sheets as illustrated. As you begin closing the property sheets, notice that roles in the list have a plus or minus sign next to them if they contain other roles. You can expand or contract the hierarchy by clicking on the plus and minus signs just as you can expand the object tree in the left pane of the Sybase Central main window.

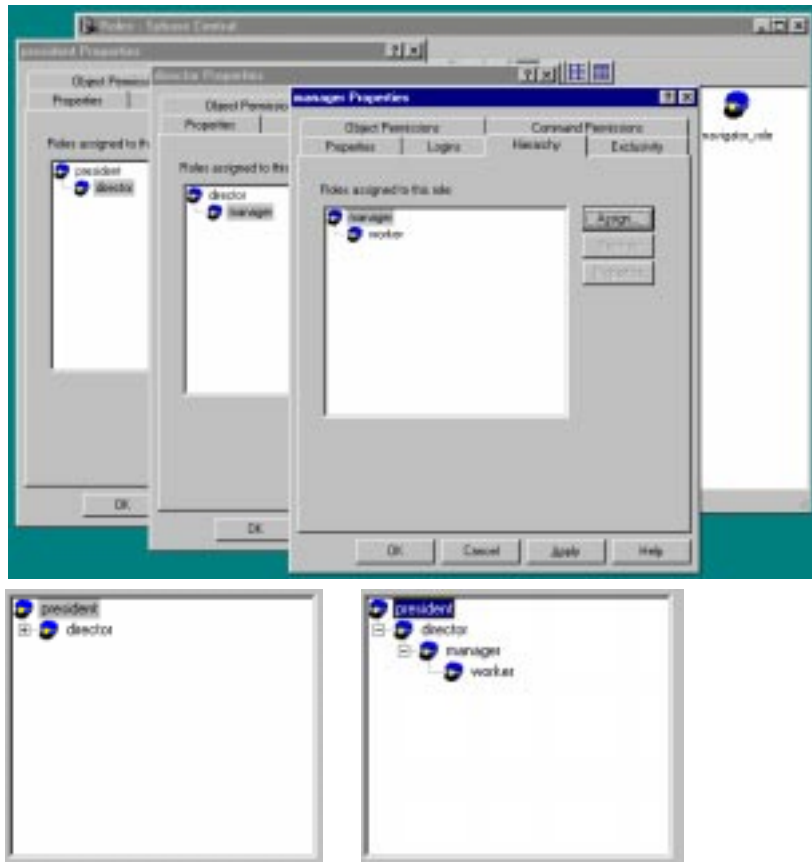


Figure 15-10: Illustration of role hierarchy

Managing Object Permissions for A Role

To grant or revoke object permissions for a role:

1. Display the property sheet for the role as described in “Displaying Role Properties” on page 15-3.
2. Select the Object Permissions tab.



Figure 15-11: Role property sheet: Object Permissions tab

3. Select the database in which you want to set permissions.
4. Select the type of object for which you want to set permissions. Initially, the Permissions For drop-down list shows tables. You can also select stored procedures, extended stored procedures, or views from the list.
5. To grant a permission on an object, click in the appropriate box for the type of permission. You can grant execute permission on stored procedures. For tables and views, you can grant the following types of permissions:
 - Sel—permission to select a table or view
 - Ins—permission to insert a row
 - Upd—permission to update a row

- Del—permission to delete a row
- Ref—permission to create referential constraints

When you make a permission change, the box displays a symbol (or deletes the current symbol), indicating that you have altered the original settings. The symbols are described in the legend at the bottom of the tab. Continue clicking in the box to cycle through the available options until you select the one that you want.

Shortcuts



Right click a grid box. Select a permission from the popup menu.

► **Note**

You can also set permissions for roles on objects from the Permissions tab of the table, stored procedure, extended stored procedure, and view property sheets.

Managing Command Permissions for A Role

To grant or revoke command permissions for a role:

1. Display the Role property sheet as described in “Displaying Role Properties” on page 15-3.
2. Select the Command Permissions tab.



Figure 15-12: Role property sheet: Command Permissions tab

3. Select the database in which you want to set command permissions.
4. Assign a permission by clicking in the Status box for the command you want to change. Each click changes the status to the next option in the cycle described in the Legend.
5. Click Apply to change the permissions in the database.

Shortcuts



Right click a grid box. Select a permission from the popup menu.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

16

Managing Defaults

What's in This Chapter

A **default** is a value that Adaptive Server inserts in a table column when a row of data is inserted and no value is provided for the column. Once you create a default, you can bind it to table columns and to user datatypes.

This chapter covers the following:

- Creating a default
- Displaying default properties
- Binding and unbinding a default
- Displaying dependencies

Creating a Default

Privileges	Only a database owner or a user with create default permission can create a default.
------------	---

To create a default:

1. In a database, select the Defaults folder.
2. From the File menu, choose Create; then, choose Default from the cascading menu. The New Default Properties dialog box opens.



Figure 16-1: New Default Properties dialog box

3. Enter a name for the default.
4. Accept the default owner or click Change and choose a different owner.
5. In the Default Expression box, enter a SQL expression for the default. The expression is the text that would follow the as keyword in a create default command.
6. Click OK.

Shortcuts



Select the Defaults folder. In the right pane double-click the Add Default icon.

Deleting a Default

Before you delete a default, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference a default, check its dependencies. See “Displaying Default Dependencies” on page 16-4.

To delete a default, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying Default Properties

To display the properties of a default:

1. Select the icon of the default.
2. From the File menu, choose Properties.

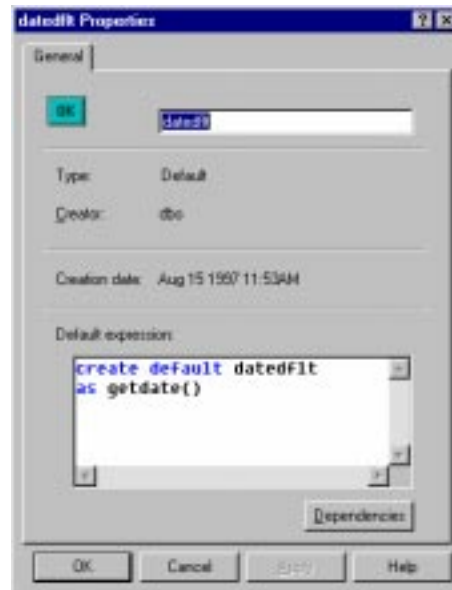





Figure 16-2: Default property sheet

Shortcuts

	Double-click the default icon.
	Select the default icon. Then select the Properties toolbar button.
	Click the right mouse button over the default icon. Then choose Properties from the shortcut menu.

Displaying Default Dependencies

To display dependencies for a default:

1. Select the default whose dependencies you want display.
2. From the File menu, choose Dependencies.





Figure 16-3: Dependencies dialog box for defaults

3. Select Tables or User Datatypes in the Object Type drop-down list.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the default icon. Choose Dependencies from the shortcut menu.

Binding a Default

You can bind a default to one or more table columns, or to all columns having a given user-defined datatype. The column has the default value if no other value is provided.

Columns that do not allow null values should have a default bound to them. Otherwise, Adaptive Server produces an error message whenever anyone inserts a row and fails to make an entry for that column.

Before binding a default, verify that:

- The column is large enough for the default.
- The default value is allowed by any rules that are bound to the column or user datatype.

► **Note**

A default bound to a column overrides a default bound to a user-defined datatype.

Binding a Default to a Table Column

To bind a default to a table column:

1. Select the Columns folder for the table you are working with.
2. Select a column.
3. Open the property sheet, as described in “Viewing the Properties of a Column” on page 22-30.



Figure 16-4: Column Properties: Rules and Defaults tab

4. Select the Rules and Defaults tab.
5. In the Default group box, select Named to bind an existing default. Select Clause to enter the text of a default.
6. If you selected Named, select a default from the drop-down list. If you selected Clause, enter the text of the default expression in the text box.

► **Note**

The default clause must be a valid constant expression that is compatible with the datatype of the column. The constant expression cannot include the name of any columns or other database objects, but can include built-in functions that do not reference database objects. Enclose strings in double quotes.

Binding a Default to a User-defined Datatype

To bind a default to a user-defined datatype:

1. Select the User-defined Datatypes folder.
2. Select a user-defined datatype.
3. Open the property sheet as described in “Displaying User-defined Datatype Properties” on page 20-4.
4. Select the Advanced Options tab.



Figure 16-5: User Datatype Properties: Advanced Options tab

5. Select a default from the Default drop-down list.
6. Optionally, select In Future Only. If you select In Future Only, the default is applied only to columns of this type added to a table after you bind the default. It does not apply to columns of this type currently in a table. If you do not select In Future Only, the default is applied to columns of this type currently in tables and columns of this type added later.

Unbinding A Default from A Table Column

To unbind a default from a table column:

1. Select the Columns folder for the table you are working with.
2. Select a column.
3. Open the property sheet for the column as described in “Viewing the Properties of a Column” on page 22-30.
4. Select the Rules and Defaults tab.
5. In the Default group box, select None.
6. Click Apply.

Unbinding A Default from A User-defined Datatype

To unbind a default from a user-defined datatype:

1. Select the User-defined Datatype folder.
2. Select a user-defined datatype.
3. Open the property sheet as described in “Displaying User-defined Datatype Properties” on page 20-4.
4. Select the Advanced Options tab.
5. In the Default drop-down list, select None.
6. Click Apply.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

17

Managing Stored Procedures

What's in This Chapter

A **stored procedure** is a named collection of SQL statements and flow control statements. Once you create a stored procedure, you can use it repeatedly without the need to enter the SQL statements individually each time you want to repeat a procedure. This chapter describes the following:

- Creating a procedure
- Displaying procedure properties
- Updating user and group permissions on a stored procedure

Creating a Stored Procedure

A stored procedure that performs a select, execute, or data modification command must be owned by the same user as the object acted upon.

Privileges	Only a database owner or a user or group with create procedure permission can create a stored procedure.
------------	---

To create a stored procedure:

1. In the database hierarchy, select the Stored Procedures folder.
2. From the File menu, choose New; then, choose Procedure from the cascading menu. The Create a New Procedure wizard opens.

The Create a New Procedure wizard asks for the following information:

Table 17-1: Inputs to Create a New Procedure wizard

Input	Description
Name	Name for the new stored procedure.
Owner	The user who owns the stored procedure.
Recompile	Whether or not to recompile each time the stored procedure is run.
Group	Whether or not the stored procedure is to be part of a group. If so, the group number.
Code	Enter the code for the stored procedure in the editor.

Recompile

Select the recompile option if you expect that the execution of the stored procedure may be different each time. For example, recompile if the data passed in its parameters changes so much that a query plan produced at execution would differ greatly from a plan that is stored.

Shortcuts



Select the Stored Procedures folder in the appropriate database. In the right pane, double-click the Add Procedure icon.

Deleting a Stored Procedure

Before you delete a stored procedure, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference a stored procedure, check its dependencies. See “Displaying Stored Procedure Dependencies” on page 17-5.

To delete a stored procedure, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Viewing Stored Procedure Code

To edit a stored procedure:

1. Select the stored procedure you want to edit.
2. From the File menu, choose Open. The code editor opens with the code for the stored procedure displayed.

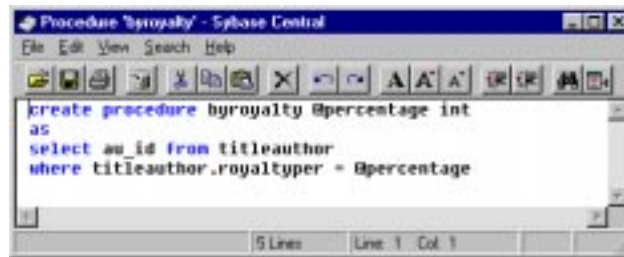


Figure 17-1: Code Editor for a stored procedure

You can edit the code, but you cannot execute the edited code against Adaptive Server. The value of editing the code would be to save it to a file or copy it to the clipboard to reuse in creating a new stored procedure. For example, you could delete this stored procedure and create it again using the edited code.

Shortcuts



Double-click the stored procedure icon.



Select the stored procedure. From the popup menu, choose Open.

Displaying Stored Procedure Properties

To display the properties of a stored procedure:

1. Select the icon for the stored procedure you want to display.
2. From the File menu, choose Properties.



Figure 17-2: Stored procedure property sheet

Shortcuts

	<p>Press the Alt key and double-click the icon for the stored procedure.</p>
	<p>Select the stored procedure icon. Then select the Properties toolbar button.</p>
	<p>Click the right mouse button over the stored procedure icon. Choose Properties from the shortcut menu.</p>

Displaying Stored Procedure Dependencies

To display dependencies for a stored procedure:

1. Select the stored procedure whose dependencies you want to display.
2. From the File menu, choose Dependencies.



Figure 17-3: Dependencies dialog box for a stored procedure

3. From the Object Type drop-down list, select the type of object you want to view. The options for the Referenced By tab are stored procedures and triggers. The options for the References tab are stored procedures, extended stored procedures, tables, views, and user datatypes.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the stored procedure icon. Then choose Dependencies from the shortcut menu.

Setting Permissions for A Stored Procedure

Sybase Central lets you update permissions for stored procedures as follows:

- Grant and revoke `execute` permission on a stored procedure.
- Grant permission to specific users, groups, or roles, or grant it using the `with grant` option so the recipient can also grant the permission to other users.
- Revoke the permission from specific users, groups, or roles, or revoke it using the `with cascade` option to revoke it from the named user and all users who acquired permission from the named user, directly or indirectly.

To update stored procedure permissions:

Privileges	A stored procedure owner can grant and revoke <code>execute</code> permission on the stored procedure.
------------	--

1. Display the stored procedure property sheet, as described in “Displaying Stored Procedure Properties” on page 17-4.
2. Click the Permissions tab.



Figure 17-4: Stored procedure property sheet: Permissions tab

3. In the Permissions For drop-down list, select Users, Groups, or Roles.
4. To change permissions for an object, click in the Exec column until it indicates the correct permission. The Legend at the bottom of the dialog box explains the symbols.
5. To see the properties of an object, select it in the list, then click the Properties button.
6. Click Apply.

Shortcuts



Right click a grid box. Select a permission from the popup menu.

Managing Extended Stored Procedures

Extended stored procedures let you call external functions from Adaptive Server. The external functions must be capable of calling a C language function and manipulating C language datatypes. Once you create an extended stored procedure, you use it as you would use any stored procedure. Extended stored procedures are contained in the Extended Stored Procedures folder.

Extended stored procedures can:

- Take input parameters
- Return a status value indicating success or failure and the reason for failure
- Return values of output parameters
- Return result sets

On platforms that support dynamic link libraries (DLLs), the external functions are compiled into DLLs. On platforms that do not support DLLs, the external functions are compiled into a single shared library named *libxp*.

Creating An Extended Stored Procedure

Privileges	Only a System Administrator can create extended stored procedures.
------------	--

To create an extended stored procedure:

1. Select the Extended Stored Procedures folder.
2. From the File menu, choose New. From the cascading menu, choose Extended Stored Procedure.
3. The Create A New Extended Stored Procedure wizard opens.

The Create A New Extended Stored Procedure wizard asks for the following information:

Table 17-2: Inputs to Create A New Extended Stored Procedure wizard

Input	Description
Name	A name for the extended stored procedure.
DLL name	<p>Name of the DLL associated with this extended stored procedure. Enter the name only, not the pathname.</p> <p>The DLL must be in a directory that is part of the PATH variable on an NT system or the LD_LIBRARY_PATH environment variable on a UNIX system.</p> <p>The DLL does not need to exist when you create the extended stored procedure, but it must exist when you try to run the extended stored procedure.</p>

Shortcuts



Select the Extended Stored Procedures folder in the appropriate database. In the right pane, double-click the Add Extended Stored Procedure icon.

Deleting an Extended Stored Procedure

Before you delete an extended stored procedure, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference an extended stored procedure, check its dependencies. See “Displaying Stored Procedure Dependencies” on page 17-5.

To delete an extended stored procedure, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Viewing Properties of an Extended Stored Procedure

To view or edit the properties of an extended stored procedure:

1. Select the extended stored procedure that you want to edit.
2. From the File menu, choose Properties. The property sheet for the extended stored procedure opens.



Figure 17-5: Extended Stored Procedure property sheet

Shortcuts



Double-click the extended stored procedure icon.



Select the extended stored procedure icon. Then select the Properties toolbar button.



Click the right mouse button over the extended stored procedure icon. Choose Properties from the shortcut menu.

Adding Parameters to an Extended Stored Procedure

The Create an Extended Stored Procedure wizard does not support inclusion of parameters for the extended stored procedure. If you want to add parameters to an extended stored procedure, generate DDL for the extended stored procedure, then edit the DDL code. To apply the changes, execute the DDL using `isql` or SQL Advantage. See “Creating and Executing DDL Scripts” on page 17-13.

Setting Permissions for an Extended Stored Procedure

You set permissions for Extended Stored Procedures the same way you do for stored procedures. See “Setting Permissions for A Stored Procedure” on page 17-6.

Extended Stored Procedure Dependencies

The dependencies dialog box lists procedures or triggers that reference the extended stored procedure.

► **Note**

The Dependencies dialog box does not display objects the extended stored procedure references because this information is not stored in Adaptive Server.

To view the dependencies of an extended stored procedure:

1. Select the extended stored procedure.
2. From the File menu, choose Dependencies.

Configuring Adaptive Server to Use Extended Stored Procedures

Extended stored procedures are run by an Open Server application called XP Server. When you run an extended stored procedure, the performance of Adaptive Server can be affected. You can set the following Adaptive Server configuration parameters to control the effect of running extended stored procedures:

- `xp_cmdshell_context`
- `esp_execution_stacksize`
- `esp_execution_priority`
- `esp_unload_dll`

`xp_cmdshell_context` sets the security context for the operating system command to be executed using the `xp_cmdshell` system extended stored procedure. If `xp_cmdshell_context` is zero, the permissions of the operating system account under which Adaptive Server is running are the permissions used to execute an operating system command from `xp_cmdshell`. This can allow users to execute operating commands that they would not ordinarily be able to execute under the security context of their own operating system accounts.

`esp_execution_stacksize` sets the size of the stack, in bytes, to allocate for ESP execution. Reset this parameter if your extended stored procedure functions require a larger stack size than the default 34816 bytes.

`esp_execution_priority` sets a priority from 0 through 15. The default priority is 8. If you set a high priority, the Open Server scheduler runs the extended stored procedure before other threads in its queue.

`esp_unload_dll` unloads a DLL from memory after the extended stored procedure that requested it terminates. If you set `esp_unload_dll`, all DLLs are unloaded. You can also unload DLLs individually using the `sp_freeldll` system procedure. Using `esp_unload_dll` minimizes the amount of memory used by XP Server.

For information about setting configuration parameters, see “Configuring Adaptive Server” in Chapter 4, “Managing the Adaptive Server Operating Environment”.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

18

Managing Rules

What's in This Chapter

A **rule** defines the acceptable values for a table column or a user-defined datatype. For example, the *pubs2* sample database rule, *title_idrule*, allows a column to contain the value “BU” followed by any four digits. Once you create a rule, you can bind it to table columns and to user-defined datatypes.

This chapter describes the following:

- Creating a rule
- Displaying rule properties
- Binding a rule to a column
- Binding a rule to a user-defined datatype
- Deleting a binding
- Displaying dependencies

Creating a Rule

Privileges	Only a database owner or a user or group with <code>create rule</code> permission can create a rule.
------------	--

To create a rule:

1. In a database hierarchy, select the Rules Folder.
2. From the File menu, choose New; then, choose Rule from the cascading menu. The New Rule Properties dialog box opens.



Figure 18-1: New Rule Properties dialog box

3. Enter a name for the rule.
4. Accept the default creator or click the Change button and choose a creator from the Select a Creator dialog box.



Figure 18-2: Select a Creator dialog box

5. Enter the rule in the Rule Expression box. Enter the part of the rule that follows the `as` keyword in the create rule statement.

The expression can be any expression valid in a `where` clause and can include arithmetic operators and relational operators. The expression takes one argument, prefixed by the `@` symbol, and refers to the value that is entered via the `update` or `insert` command. The expression cannot reference any column or other database object.

6. Click OK.

Shortcuts



Select the Rules folder. In the right pane, double-click the Add Rule icon.

Deleting a Rule

Before you delete a rule, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference a rule, check its dependencies. See “Displaying Rule Dependencies” on page 18-5.

To delete a rule, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying Rule Properties

To display the properties of a rule:

1. Select the icon of the rule for which you want to display properties.
2. From the File menu, choose Properties.



Figure 18-3: Rule property sheet

Shortcuts

- | | |
|--|--|
| | Double-click the rule icon. |
| | Select the rule icon. Then select the Properties toolbar button. |
| | Click the right mouse button over the rule icon. Choose Properties from the shortcut menu. |

Displaying Rule Dependencies

To display dependencies for a rule:

1. Select the icon of the rule whose dependencies you want to display.
2. From the File menu, choose Dependencies.



Figure 18-4: Dependencies dialog box for a rule

3. In the Object Type box, select the object type you want to view. The options are tables and user-defined datatypes.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the rule icon. Choose Dependencies from the shortcut menu.

Binding a Rule

Before binding a rule to a table column or user-defined datatype, consider the following:

- When you bind a rule to a table column or user-defined datatype that already has a rule bound to it, the new rule takes effect and the old one is unbound.
- Rules bound to columns take precedence over rules bound to user-defined datatypes.
- Rules must be compatible with the datatype of the column.
- Rules do not apply to data already in the database when they are created.
- Rules do not override column definitions.
- Make sure that any default value bound to a column or user-defined datatype is compatible with the rule. A default that conflicts with the rule is not inserted.
- You cannot bind a rule to a *text*, *image*, or *timestamp* column.
- Do not assign a value to a user-defined datatype variable that would be rejected by a rule bound to a column of the same datatype.

Binding a Rule to a Table Column

To bind a rule to a table column:

1. Select the Columns folder for the table you are working with.
2. Select the column to which you want to bind a rule.
3. Open the property sheet for the column as described in “Viewing the Properties of a Column” on page 22-30.
4. Select the Rules and Defaults tab.



Figure 18-5: Column property sheet: Rules and Defaults tab

5. In the Rule drop-down list, select the rule you want to bind to the table column.

Binding a Rule to a User-defined Datatype

To bind a rule to a user-defined datatype:

1. Select the User-defined Datatypes folder for the database in which you are working.
2. Select the user-defined datatype to which you want to bind a rule.
3. Open the property sheet as described in “Displaying User-defined Datatype Properties” on page 20-4.
4. Select the Advanced Options tab.



Figure 18-6: User-defined Datatype property sheet: Advanced Options tab

5. Select a rule from the Rule drop-down list.
6. Optionally, select In Future Only. If you select In Future Only, the rule is applied only to columns of this type added to a table after you bind the rule. It does not apply to columns of this type currently in a table. If you do not select In Future Only, the rule is applied to columns of this type currently in tables and columns of this type added later.

Unbinding A Rule

Before deleting a rule, you must unbind it from any table columns or user-defined datatypes to which it is bound.

Unbinding A Rule From A Table Column

To unbind a rule from a table column:

1. Select the Columns folder for the table you are working with.
2. Select the column from which you want to delete a binding.
3. Open the property sheet for the column as described in “Viewing the Properties of a Column” on page 22-30.
4. Select the Rules and Defaults tab.
5. In the Rule drop-down list, select None.
6. Click Apply.

Unbinding A Rule From A User-defined Datatype

To unbind a rule from a user-defined datatype:

1. Select the User-defined Datatypes folder.
2. Select the user-defined datatype from which you want to delete a binding.
3. Open the property sheet as described in “Displaying User-defined Datatype Properties” on page 20-4.
4. Select the Advanced Options tab.
5. In the Rule drop-down list, select None.
6. Click Apply.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

19

Managing Triggers

What's in This Chapter

A **trigger** is a special type of procedure attached to a table column. The trigger goes into effect when a user changes the table with an insert, update, or delete command. The trigger executes immediately after the data modification statements are completed. Because triggers are more versatile than rules and referential integrity constraints, they are often used to do the following:

- Enforce referential integrity.
- Cascade changes through related tables in a database, such as deleting a record.
- Enforce restrictions that are more complex than those defined with rules. Unlike rules, triggers can reference columns or database objects.
- Perform “what if” analyses, such as comparing a table before and after a data modification, and performing an action based on the results of the comparison.

Only the owner of an object can create a trigger on it. However, a trigger can modify an object owned by another user. If a trigger activates a trigger on another object, the owner of the first trigger must have privileges on the target object.

This chapter describes the following:

- Creating a trigger
- Displaying trigger properties
- Displaying dependencies

Creating a Trigger

When you create a trigger, you specify the table affected and the command (**insert**, **update**, or **delete**) that activates the trigger. You also specify the action the trigger is to take.

Privileges	Only a table owner or database owner can create a trigger.
------------	--

To create a trigger:

1. In the table you are working with, select the Triggers folder.
2. From the File menu, choose New; then, choose Trigger from the cascading menu. The Create a New Trigger wizard opens.

The Create a New Trigger wizard asks for the following information:

Table 19-1: Inputs to Create a New Trigger wizard

Input	Description
Name	Name for the new trigger.
Owner	Specify the user who owns the trigger.
Event	Specify the events that can fire the trigger: Insert, Delete, Update, Update of a column.
Code	Enter the code for the trigger in the editor. This text can consist of any valid SQL statements that would follow the as keyword in a create trigger statement.

Shortcuts



Select the Triggers folder. In the right pane, double-click the Add Trigger icon.

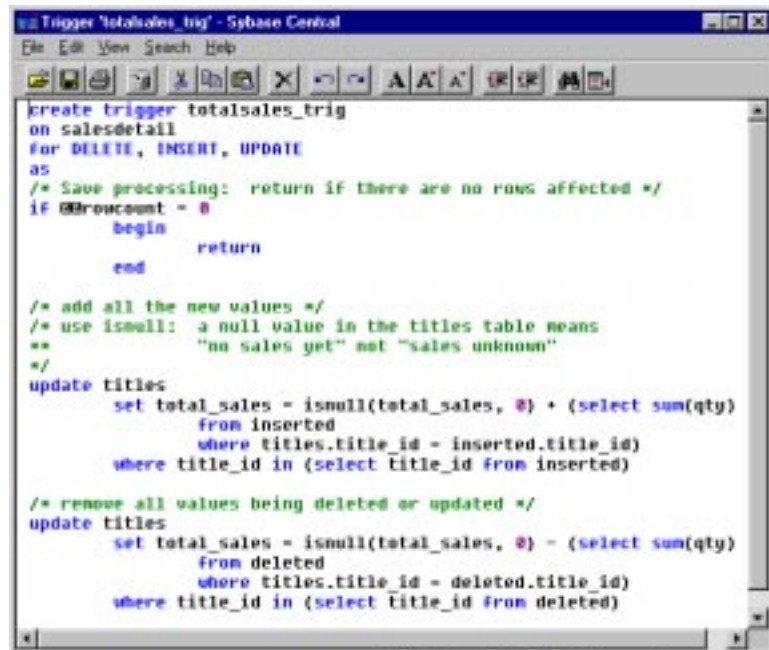
Deleting a Trigger

To delete a trigger, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Editing A Trigger 's Code

To edit the code for a trigger:

1. Select the trigger.
2. From the File menu, choose Open. The code editor opens with the trigger's code displayed.



```
Trigger 'totalsales_trig' - Sybase Central
File Edit View Search Help
[Icons]
create trigger totalsales_trig
on salesdetail
for DELETE, INSERT, UPDATE
as
/* Save processing: return if there are no rows affected */
if @@rowcount = 0
begin
return
end


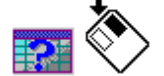
/* add all the new values */
/* use isnull: a null value in the titles table means
** "no sales yet" not "sales unknown"
*/
update titles
set total_sales = isnull(total_sales, 0) + (select sum(qty)
from inserted
where titles.title_id = inserted.title_id
where title_id in (select title_id from inserted)

/* remove all values being deleted or updated */
update titles
set total_sales = isnull(total_sales, 0) - (select sum(qty)
from deleted
where titles.title_id = deleted.title_id
where title_id in (select title_id from deleted)
```

Figure 19-1: Code editor for a trigger

You can edit the code, but you cannot execute the edited code against Adaptive Server. The benefit of editing the code could be to save it to a file or copy it to the clipboard to reuse in creating a new trigger. For example, you could delete this trigger and create it again using the edited code.

Shortcuts

	<p>Double-click the trigger icon.</p>
	<p>Click the right mouse button over the trigger icon. Choose Open from the shortcut menu.</p>

Displaying Trigger Properties

To display the properties of a trigger:

1. Select the icon of the trigger.
2. From the File menu, choose Properties.

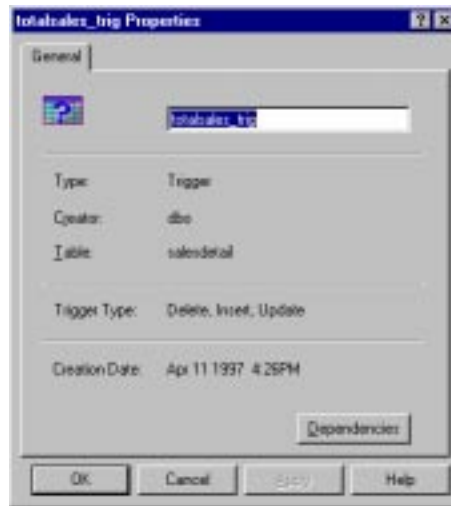


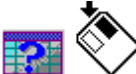


Figure 19-2: Trigger property sheet

Shortcuts

	Hold down the Alt key and double-click the trigger icon.
	Select the trigger icon. Then select the Properties toolbar button.
	Click the right mouse button over the trigger icon. Choose Properties from the shortcut menu.

Displaying Trigger Dependencies

To display objects that the trigger references:

1. Select the trigger icon.
2. From the File menu, choose Dependencies.

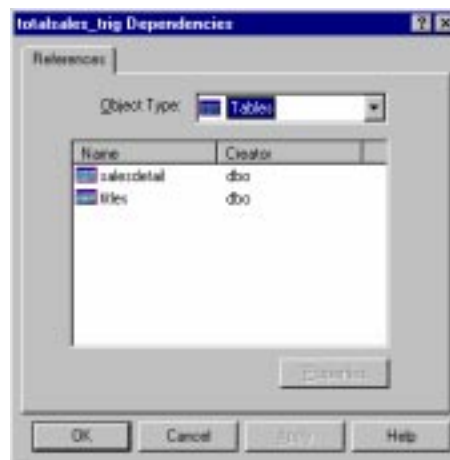



Figure 19-3: Dependencies dialog box for a trigger

3. Choose an object type from the Object Type drop-down list. The choices are procedures, extended stored procedures, tables, or views.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the trigger icon; then, choose Dependencies from the shortcut menu.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

20

Managing User-defined Datatypes

What's in This Chapter

The datatype of a table column specifies the kind of information (characters, numbers, or dates) the column holds and how the data is stored. For example, the integer (*int*) datatype stores whole numbers in the range of minus 2 billion to plus 2 billion. The tiny integer (*tinyint*) datatype stores whole numbers between 0 and 255.

A **user-defined datatype** is a customized datatype based on system datatypes. You can give a name to a frequently used datatype definition to save time. Unlike system datatypes, user-defined datatypes are case-sensitive. In addition, if a user-defined datatype has *precision*, *scale*, or *identity*, you cannot overwrite these attributes when you use the datatype to define a table column. When you define a user-defined datatype, you can bind it to rules and defaults.

This chapter describes the following:

- Creating a user-defined datatype
- Displaying user-defined datatype properties
- Editing a user-defined datatype
- Displaying dependencies

Creating a User-defined Datatype

To use a user-defined datatype in more than one database, create it in the *model* database template so it is known to all new databases you create.

Privileges	A table owner can define user-defined datatypes.
------------	--

To create a user-defined datatype:

1. In the database hierarchy, select the User-defined Datatypes folder.
2. From the File menu, choose New; then, choose User-defined Datatype from the cascading menu. The New User-defined Datatype property sheet opens.

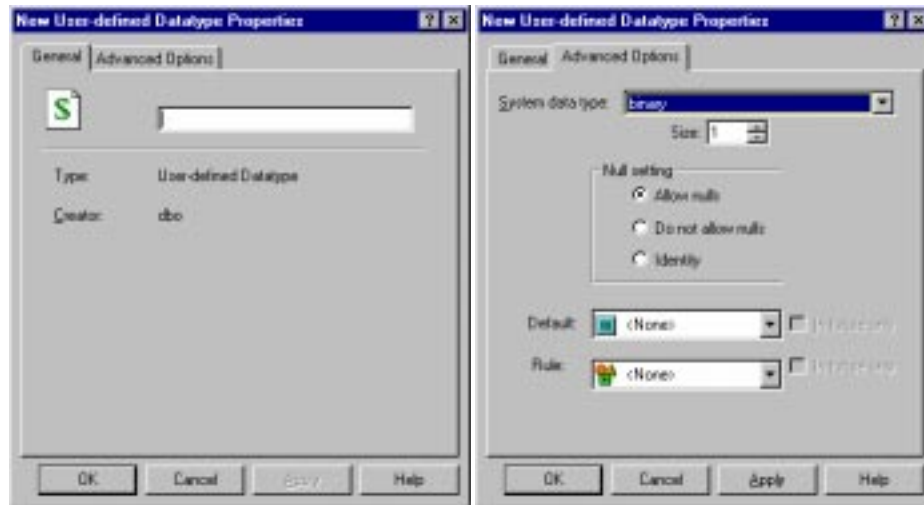


Figure 20-1: New User-defined Datatype dialog box

3. In the Name box, enter a name for the user-defined datatype.
4. Click the Advanced Options tab.
5. From the System Data Type list, select a system datatype on which to base the user-defined datatype.
6. If the datatype is *text*, *char*, *varchar*, *binary*, *varbinary*, *nchar*, or *nvarchar*, enter a maximum length for the datatype in the Size box.
7. If the datatype is *decimal*, *numeric*, or *float*, enter the precision in the Precision box. For *decimal* or *numeric*, precision is the maximum number of decimal digits that can be stored in the column. The range is 1 to 38 with a default of 18. For *float*, precision is the number of significant binary digits in the mantissa.
8. If the datatype is *decimal* or *numeric*, enter the scale in the Scale edit box. Scale is the maximum number of digits to the right of the decimal point. The scale must be less than or equal to the precision. The range is 10 to 38 with a default of 0.

9. In the Null Setting group box, specify how the user-defined datatype handles null values. To specify that the datatype can accept null values, select Allow Nulls. To specify that the datatype does not accept null values, select No Nulls. To specify that the datatype can use the system identity property to identify each row in the table, select Identity. Only one column can have the identity property.
10. To bind the user-defined datatype to a default, select a default from the Default list.
11. To bind the user-defined datatype to a rule, select a rule from the Rule list.
12. Click OK.

Shortcuts



Select the User-defined Datatypes folder. In the right pane, double-click the Add User-defined Datatype icon.

Deleting a User-defined Datatype

Before you delete a user-defined datatype, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference a user-defined datatype, check its dependencies. See “Displaying User-defined Datatype Dependencies” on page 20-5.

To delete a user-defined datatype, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying User-defined Datatype Properties

To display the properties of a user-defined datatype:

1. Select the icon of the user-defined datatype you want to display.
2. From the File menu, choose Properties.




Figure 20-2: User-defined Datatype property sheet

Shortcuts

- | | |
|--|--|
| | Double-click the icon of the user-defined datatype. |
| | Select the user-defined datatype icon. Then select the Properties toolbar button. |
| | Click the right mouse button over the icon of the user-defined datatype. Choose Properties from the shortcut menu. |

Binding Rules to User-defined Datatypes

To bind a rule to a user-defined datatype, follow the procedure in “Binding a Rule to a User-defined Datatype” on page 18-7. To unbind a rule from a user-defined datatype, follow the procedure in “Unbinding A Rule From A User-defined Datatype” on page 18-9.

Displaying User-defined Datatype Dependencies

To display objects that a user-defined datatype references and that reference a user-defined datatype:

1. Select the user-defined datatype whose dependencies you want to display.
2. From the File menu, choose Dependencies.

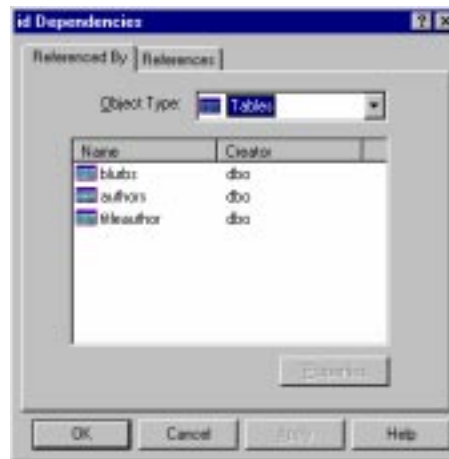


Figure 20-3: Dependencies dialog box for a user-defined datatype

3. To display objects that the user-defined datatype references, click the References tab. To display objects referenced by the user-defined datatype, select the Referenced By tab.
4. In either tab, from the Object Type drop-down list, select the type of object you want to view.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the user-defined datatype icon. Choose Dependencies from the shortcut menu.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

21 Managing Views

What's in This Chapter

A **view** is a subset of columns from one or more tables. For example, if you frequently want to see the expense column from Table A and the income column from Table B, you can create a named view to give you this information quickly. Once you create a view, you can set permissions for it.

This chapter describes the following:

- Creating a view
- Displaying view properties
- Displaying view dependencies
- Updating user and group view permissions

Creating a View

Privileges	Only a database owner or a user or group with create view permission can create a view.
------------	--

To create a view:

1. In the database hierarchy, select the Views folder.
2. From the File menu, choose New; then, choose View from the cascading menu. The Create View wizard opens.

The Create View wizard asks for the following information:

Table 21-1: Inputs to the Create View wizard

Input	Description
Name	Name for the new view.
Owner	The user that owns the view.
Code	Enter code for the view in the editor.

Shortcuts



Select the Views folder. In the right pane, double-click the Add View icon.

Deleting a View

Before you delete a view, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference a view, check its dependencies. See “Displaying View Dependencies” on page 21-5.

To delete a view, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying View Properties

To display the properties of a view:

1. Select the icon of the view you want to display.
2. From the File menu, choose Properties. The View property sheet opens.



Figure 21-1: View property sheet

Shortcuts



Hold down the Alt key and double-click the view icon.



Select the view icon. Then select the Properties toolbar button.



Click the right mouse button over the view icon. Choose Properties from the shortcut menu.

Editing a View

To edit the code of a view:

1. Select the view that you want to edit.
2. From the File menu, choose Open. The code editor opens with the text of the view displayed.

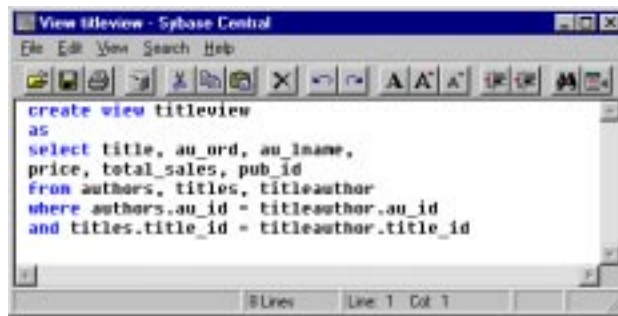


Figure 21-2: Code editor for a view

You can edit the code, but you cannot execute the edited code against Adaptive Server Enterprise. The benefit of editing the code could be to save it to a file or copy it to the clipboard to reuse in creating a new view. For example, you could delete this view and create it again using the edited code.

Shortcuts



Double-click the view icon.



Click the right mouse button over the view icon. Choose Open from the shortcut menu.

Displaying View Dependencies

To display objects that a view references and that reference a view:

1. Select the view whose dependencies you want to display.
2. From the File menu, choose Dependencies.



Figure 21-3: Dependencies dialog box for a view

3. To display objects that the view references, select the References tab. To display objects referenced by the view, select the Referenced By tab.
4. In either tab, from the Object Type drop-down list, select the type of object you want to view.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the user-defined datatype icon. Choose Dependencies from the shortcut menu.

Granting Permissions on a View

You can grant and revoke permissions on a view. You can grant the permission to specific users, groups or roles, or you can grant the permission using the **with grant** option so the recipient can also grant the permission to other users. You can revoke the permission from specific users, groups, or roles, or you can revoke the permission using the **with cascade** option to revoke it from the named user and all users who acquired the permission from the current user (directly or indirectly).

Privileges	A view owner can grant and revoke object permissions on a view.
------------	---

To manage permissions on a view, you can see which users, groups, and roles have permissions. You can also see all users, groups, and roles in the database and those that do not have any permission on the view.

To display view permissions:

1. Display the View property sheet, as described in “Displaying View Properties” on page 21-3.
2. Select the Permissions tab.

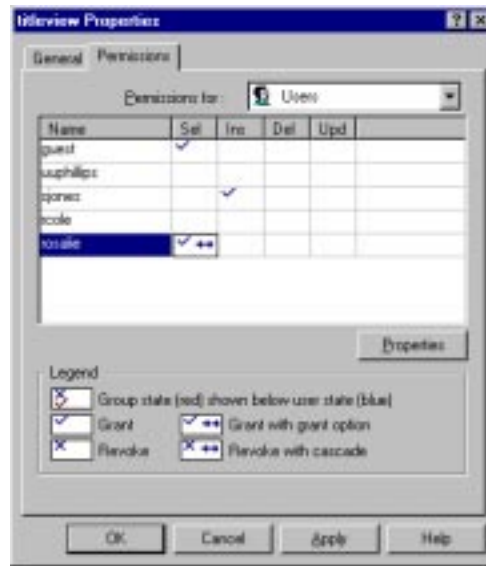


Figure 21-4: View property sheet: Permissions tab

3. In the Permissions For drop-down list, select Users, Groups, or Roles. The dialog box shows users, groups, or roles and the current state of their permissions.
4. Select the object whose permissions you want to change.
5. To grant or revoke permissions, click in one of the columns for the object you have chosen. The permissions cycle through the options listed on the Legend.

You can grant the following types of permissions:

- Sel—permission to select the view
- Ins—permission to insert a row in the view
- Up—permission to update a row in the view
- Del—permission to delete a row in the view
- Ref—permission to create a referential constraint that references the view

► **Note**

Sybase Central reports on permissions explicitly granted or revoked and those that users obtain by being members in a group. It does not show permissions associated with having an administrative role, such as System Administrator, associated with their logins.

Shortcuts



Right click a grid box. Select a permission from the popup menu.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

22

Managing User Tables

What's in This Chapter

Tables are the heart of a database. Consisting of columns and rows that contain data, a table is the mechanism for storing the data in a database:

- A **system table** stores information that allows the database to perform its services.
- A **user table** stores and provides access to user data. A special kind of user table called a **proxy table**, which accesses data on remote servers, is described in Chapter 23, “Managing Proxy Tables.”

System table definitions are not typically updated. This chapter describes how to use Sybase Central to create and update user table definitions.

The chapter describes the following tasks:

- Creating a table
- Displaying table properties
- Displaying table dependencies
- Updating permissions on tables
- Placing a table on a segment
- Creating constraints on a table
- Creating referential constraints on a table
- Recompiling procedures and triggers
- Checking table consistency
- Updating statistics
- Managing columns

Understanding Constraints

When you create a table, you define the characteristics of its columns. You can specify **constraints** on the data a column contains. Constraints allow you to control the type of data that can be added to a table column. Sybase Central lets you specify the following types of constraints:

- The value that appears in the column if no value is entered on an insertion or update. This type of constraint is called a **default**.
- Criteria against which Adaptive Server checks data entered on an insertion or update. This type of constraint is called a **rule** or **check constraint**.
- Key values in another table that must match key values in the row being inserted or updated. This type of constraint is called a **referential constraint**.

Implementing Rules and Check Constraints

You can create a rule or check constraint in the following ways:

- As part of the definition of a table column. This type of constraint, called a column check constraint, applies only to the column on which you define it.
- As part of the definition of a table. This type of constraint, called a table check constraint, can apply to all columns of a table.
- As a database object. This type of constraint is called a rule. Once you create a rule, you can bind it to table columns and to user-defined datatypes.

The advantage of the first two methods is that you specify the table or column and its constraint in the same step. The advantage of the third is that you define the rule once and apply it as needed to multiple table columns.

This section describes the method by which you create a column check constraint on a table column. For information about creating a table check constraint, see “Creating Check Constraints on a Table” on page 22-19. For information about creating a rule and binding it to a table column, see “Creating a Rule” on page 18-1 and “Binding a Rule” on page 18-6.

Implementing Referential Constraints

Referential constraints apply to a table as a whole. To create referential constraints, first create the table and then modify it by creating a foreign key. For details, see “Creating Referential Constraints on a Table” on page 22-21.

Constraints Summary

The following table summarizes the types of constraints you can place on table data and lists how to implement them.

Table 22-1: Types of table constraints

Constraint type	Column level?	Table level?	Where to implement
Default	Yes	No	<ul style="list-style-type: none"> Rules and Defaults tab of Column property sheet. See “Binding a Default to a Column” on page 22-32.
Check	Yes	Yes	<ul style="list-style-type: none"> Edit Check Constraints dialog box in table editor for column level constraints. See “Specifying Column Check Constraints” on page 22-8 New Check Constraints dialog box in Check Constraints folder for table level constraints. See “Creating Check Constraints on a Table” on page 22-19.
Rule	Yes	No	<ul style="list-style-type: none"> Rules and Defaults tab of Column property sheet. See “Binding a Rule to a Column” on page 22-31.
Referential	No	Yes	<ul style="list-style-type: none"> Create a New Foreign Key wizard in Foreign Keys folder. See “Creating Referential Constraints on a Table” on page 22-21.

Using the Table Editor

The table editor is the tool you use to create and edit tables and their columns.

Once you open the table editor, you can do the following:

- Enter a name for a new table and choose a creator in the toolbar.
- Change an existing table's name.
- Create, edit, and delete columns.
- Manage check constraints for columns.
- Edit the primary key of the table.

The table editor does not have a menu bar. To access its features you use toolbar icons or selections on a popup menu. Figure 22-1 illustrates the table editor toolbar.



Figure 22-1: Table editor toolbar

Table 22-2 is the key to Figure 22-1 and describes the purpose of each icon in the toolbar.

Table 22-2: Table editor toolbar icons and key to Figure 22-1

Key	Icon	Purpose
A	Add Column	Adds a new, empty row at the bottom of the table editor for defining a new table column.
B	Delete Column	Deletes the currently selected table column from the table editor. You can only delete columns that have not yet been saved (committed) to the table.
C	Column Properties	Opens a property sheet for the selected table column. The selected column must have been already saved (committed) to the table.
D	Edit Primary-key Columns	Allows you to edit the primary key for this table.
E	Edit Check Constraints	Allows you to manage the check constraints on columns in this table.
F	Save Table	Saves (commits) your changes to this table. The table editor stays open.

Table 22-2: Table editor toolbar icons and key to Figure 22-1

Key	Icon	Purpose
G	Save and Close	Saves (commits) your changes to this table, then closes the table editor.
H	Help	Opens online help for the table editor.

Opening the Table Editor

When you create a new table, the table editor opens in response to the Create New Table command. To open the table editor for an existing table:

1. Select a table.
2. From the File menu, choose Edit Columns.

Shortcut



Click the right mouse button over the table icon. Choose Edit Columns from the shortcut menu.

Creating a Table

Privileges	Only a database owner or a user with create table permission can create a table.
------------	---

To create a table:

1. In a database you are working in, select the User Tables folder.
2. From the File menu, choose New; then, choose Table from the cascading menu. The Table Editor opens.




Figure 22-2: Table editor

3. In the Name box, enter a name for the table.
4. From the Owner list, choose an owner for the table. The default is “dbo”.
5. Specify column definitions and check constraints as described on the following pages.
6. After you specify columns and check constraints, click the right mouse button and choose Save and Close from the shortcut menu.

Shortcuts



Select the User Tables folder. In the right pane, double-click the Add Table icon.

Specifying Column Definitions

For each column in the table:

1. In the Column Name box, enter a name for the column.
2. From the Data type drop-down list, select a datatype for the column.
3. Enter the specifications appropriate to the datatype:
 - If the datatype is *char*, *nchar*, *varchar*, *nvarchar*, *binary*, or *varbinary*, enter a length for the datatype in the Size/Prec. column.
 - If datatype is *decimal*, *numeric*, or *float*, enter the precision in the Size/Prec. column. For *decimal* or *numeric*, precision is the maximum number of decimal digits that can be stored in the column. Specify a precision of 1 to 38. For *float*, precision is the number of significant binary digits in the mantissa.
 - If datatype is *decimal* or *numeric*, enter the scale in the Scale edit box. Scale is the maximum number of digits that can be stored to the right of the decimal point. The scale must be less than or equal to the precision. Specify a scale of 0 to 38.
4. Click in the Allow Nulls column to set null options:
 - To allow NULL values in the column, indicated by a check mark
 - To disallow NULL values in the column, the column in the table editor is blank.
 - To allow the column to contain a system-generated, sequential value that identifies each row in the table, click the right mouse button and choose Identity from the shortcut menu. Only one column in a table can have this characteristic. The datatype for the column must be numeric for the Identity option to be available.

Specifying Column Check Constraints

To specify a column check constraint:

1. With the column highlighted, click the Check Constraint icon in the toolbar. The Edit Check Constraints dialog box opens.



Figure 22-3: Edit Check Constraints dialog box

2. Click New.
3. In the Constraint Name box, enter the constraint name.
4. In the Search Condition box, enter the text of the constraint. Check constraint text specifies a search condition that a value must pass before it is inserted into the table. For format information, see “Check Constraint Format” on page 22-9.
5. Click Add to List.
6. Click OK.

Check Constraint Format

The Search Condition consists of the `check` keyword followed by a search condition, enclosed in parentheses. The search condition specified by a check constraint can include any of the following:

- A list of constant expressions introduced with `in`
- A range of constant expressions introduced with `between`
- A set of conditions, introduced with `like`, that can contain wildcard characters

An expression can include arithmetic operators and Transact-SQL built-in functions. The search condition cannot contain subqueries, aggregate functions, or a host variable or parameter.

For example, the following constraint on the `pub_id` column specifies that the value of `pub_id` must be 1389, 0736, 0877, 1622, 1756, or any 4-digit number beginning with 99:

```
check (pub_id in ("1389", "0736", "0877", "1622",  
"1756") or pub_id like "99[0-9][0-9]")
```

Specifying Primary Keys

A primary key is a column or combination of columns that uniquely identifies a row. It cannot be NULL and it must have a unique index. A table with a primary key is eligible for joins with foreign keys in other tables. The primary key table can be thought of as the master table in a master-detail relationship. There can be many such master-detail groups in a database.

In the `pubs2` database, for example, the `title_id` column is the primary key of `titles`. It uniquely identifies the books in `titles`, and joins with `title_id` in `titleauthor`, `salesdetail`, and `roysched`. The `titles` table is the master table in relation to `titleauthor`, `salesdetail`, and `roysched`. The diagram in Chapter 1, "The `pubs2` Database," of *Adaptive Server Enterprise Reference Supplement* shows these relationships.

To specify a primary key on a column:

1. Click the right mouse button. From the shortcut menu, choose Edit Primary-key Columns. The Edit Primary-key Columns dialog box opens.

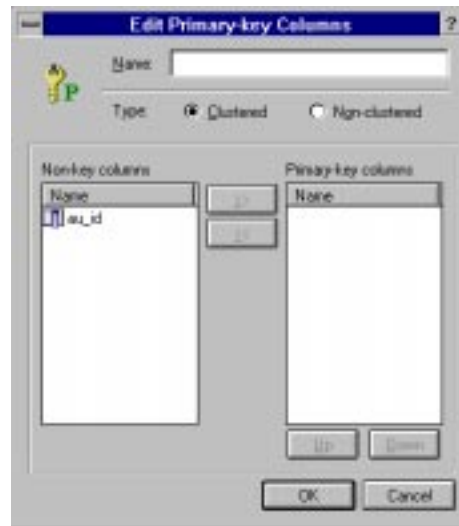


Figure 22-4: Edit Primary-key Columns dialog box

2. Enter a name for the primary key.
3. Specify whether the index to be created for the primary key will be clustered or nonclustered. For definitions of types of indexes, see Chapter 24, “Managing Indexes.”
4. In the Non-key Columns list, select the columns to be part of the key.
5. Click the right-pointing arrow to move the selected columns to the Primary Key Columns list.
6. Click OK.

Creating User Tables When Component Integration Services is Enabled

If you are using Component Integration Services, you can specify a default location for proxy tables using `sp_defaultloc`. If a default location is specified, you cannot create user tables. Any attempt to create a user table will result in creation of a new proxy table/remote table pair.

There is no visible indication in Sybase Central that a default location is specified, so you must be aware of your environment when you create user tables.

Deleting a Table

Before you delete a table, be sure that no other objects reference it. If any objects reference it, you must edit those objects to avoid errors. To find out if other objects reference a table, check its dependencies. See “Displaying Table Dependencies” on page 22-14.

When you delete a table, Adaptive Server deletes any indexes and triggers associated with the table and unbinds any rules or defaults that are bound to its columns.

Privileges	A table owner can delete a table.
------------	-----------------------------------

To delete a table, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying Table Properties


To display the property sheet for a table:

1. Select the icon of the table you want to display.
2. From the File menu, choose Properties.




Figure 22-5: Table property sheet

Shortcuts

-
- 

Select the table icon, and select the Properties toolbar button.

 - 

Click the right mouse button over the table icon, and choose Properties from the shortcut menu.
-

Adding Columns to a Table

To add columns to a table:

1. Select the table.
2. From the File menu, choose Edit Columns. The Table Editor opens.

Key	Column Name	Data Type	Size/Prec.	Scale	Allow Nulls	Check Constraints
	au_id	id	11			
	au_fname	varchar	40			
	au_lname	varchar	20			
	phone	char	12			
	address	varchar	40		✓	
	city	varchar	20		✓	
	state	char	2		✓	
	county	varchar	12		✓	
	postalcode	char	10		✓	

Figure 22-6: Table Editor

3. For each column you add, follow the instructions in “Specifying Column Definitions” on page 22-7.

► **Note**

When you add a column to a table, you cannot specify No Nulls. This is because rows that are already in the table will have null values for the new column when it is created.

4. To create constraints on new columns, follow the instructions in “Specifying Column Check Constraints” on page 22-8.

Setting Cache Strategy

You can select a cache strategy for a table. This selection overrides the default strategy, determined by the Adaptive Server optimizer, for reading data pages from a table into the buffers in data cache. You set cache strategy in the index property sheet.

1. In the Indexes folder for a table, select the index on which you want to set the cache strategy.
2. Open the Index property sheet as described in “Displaying Index Properties” in Chapter 24, “Managing Indexes”.
3. Select the Miscellaneous tab.
4. Choose one or both of the following options:
 - Most Recently Used Replacement—This selection specifies that Adaptive Server uses the most recently used (MRU) strategy for determining where in cache to place data pages when reading in new data.

If you clear the check box, Adaptive Server reads new pages into the MRU end of the chain of buffers in cache. Subsequent reads move the pages along the chain towards the least recently used (LRU) end until they are flushed out by new reads at the MRU end. If you select Most Recently Used Replacement, Adaptive Server reads new pages into the LRU end. They are used and immediately flushed when a new page enters the MRU end.

This strategy is advantageous when a page is needed only once for a query. It tends to keep such pages from flushing out other pages that can potentially be reused while still in cache.

Displaying Table Dependencies

When you create referential constraints (see “Creating Referential Constraints on a Table” on page 22-21), you set up dependencies between a table and other objects. The Dependencies dialog box lets you view these relationships.

You can display defaults, rules, user-defined datatypes, and tables that a table references, and you can display procedures, views, triggers, and tables that reference the table.

1. Select the icon of the table whose dependencies you want to display.
2. From the File menu, choose Dependencies.

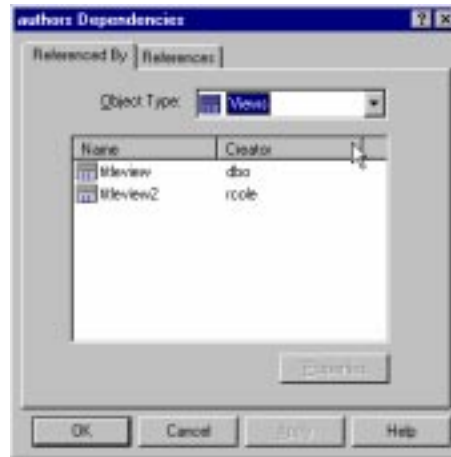



Figure 22-7: Dependencies dialog box for a table

3. To display objects that the table references, select the References tab. To display objects referenced by the table, select the Referenced By tab.
4. In either tab, from the Object Type drop-down list, select the type of object you want to view.

Shortcuts



Click the Dependencies button on the General tab of the property sheet.



Click the right mouse button over the trigger icon. Choose Dependencies from the shortcut menu.

Setting Permissions for A Table

You can grant and revoke permissions on a table. You can grant the permission to specific users, groups or roles, or you can grant the permission using the **with grant** option so the recipient can also grant the permission to other users. You can revoke the permission from specific users, groups or roles, or you can revoke the permission using the **with cascade** option to revoke it from the named user and all users who acquired the permission from the current user (directly or indirectly).

Privileges	A table owner can grant and revoke object permissions on a table.
------------	---

To set table permissions:

1. Display the Table property sheet, as described in “Displaying Table Properties” on page 22-12.
2. Click the Permissions tab.



Figure 22-8: Table property sheet: Permissions tab

3. From the Permissions For drop-down list, select Users, Groups, or Roles.

4. To grant or revoke permissions, click in one of the columns for the user, group, or role you have chosen. The permissions cycle through the options listed on the Legend.

You can grant the following types of permissions:

- Sel—permission to select the view
- Ins—permission to insert a row in the view
- Up—permission to update a row in the view
- Del—permission to delete a row in the view
- Ref—permission to create a referential constraint that references the view

► **Note**

Sybase Central reports on permissions explicitly granted or revoked and those that users obtain by being members in a group. It does not show permissions associated with having a role, such as sa_role, associated with their logins.

Shortcuts



Drag and drop the table icon on a user icon to grant all permissions to the user.



Drag and drop the table icon on a group icon to grant all permissions to the group members.



Right click a grid box. Select a permission from the popup menu.

Placing a Table on a Segment

Using a segment to put a table on a specific database device can increase Adaptive Server performance and can give increased control over placement, size, and space usage of database objects. For example:

- If you put a table on one device, and its nonclustered indexes on a device on another disk controller, the time required to read or write to the disk can be reduced.
- If you put tables and indexes only on specific segments, those objects cannot grow beyond the space available on the devices represented by the segments.
- If you put tables and indexes only on specific segments, other objects cannot contend for space with them.
- If a large, heavily-used table is split across devices on two separate disk controllers, read and write time may be improved.


To place a table on a segment:

1. Display Table Properties, as described in “Displaying Table Properties” on page 22-12.
2. Select the Miscellaneous tab.



Figure 22-9: Table property sheet: Miscellaneous tab

3. Select a segment from the Segment drop-down list.
4. Click Apply.

Shortcut

Drag and drop the table icon onto the segment icon.

Creating Check Constraints on a Table

This section describes creating check constraints on a table. For general information about constraints see “Understanding Constraints” on page 22-2. For instructions on creating check constraints on a table column, see “Specifying Column Check Constraints” on page 22-8.

To create a table check constraint:

1. In the table you are working with, select the Check Constraints folder.
2. From the File menu, choose New; then choose Check Constraint from the cascading menu. The New Check Constraint Properties dialog box opens.



Figure 22-10: New Check Constraint Properties dialog box

3. Enter a name for the constraint.
4. In the Check Constraint box, enter the text of the constraint. Check constraint text specifies a search condition that any value must pass before it is inserted into the table. For format information, see ““Check Constraint Format” on page 22-9.”
5. Click Apply.

Shortcuts



Select the Check Constraints folder. In the right pane, double-click the Add Check Constraint icon.

Creating Referential Constraints on a Table

Referential constraints allow you to ensure that data inserted in one table has matching values in another table. For example, in the *pubs2* database, a row inserted into the *salesdetail* table, recording the sale of books, must have a valid *title_id* in the *titles* table. Currently, *pubs2* enforces this referential integrity with a trigger. However, a referential constraint on the *salesdetail* table, referencing the *title_id* column in the *titles* table, could accomplish the same enforcement of referential integrity.

A referential constraint involves two tables:

- A **referencing** table on which data insertions are performed. This table contains one or more columns that serve as a foreign key.
- A **referenced** table against which insertions are checked. This table contains one or more columns that serve as a primary key.

When an insertion occurs on the referencing table, the value in the foreign key columns must match a value in the primary key columns on the referenced table. You can define a referential constraint on either the referenced or the referencing table.

Privileges	Only a database or table owner who has <code>references</code> permission on the referenced table can create a referential constraint.
------------	--

To create a referential constraint:

1. In the table hierarchy, select the Foreign Keys folder.
2. From the File menu, choose New; then choose Foreign Key from the cascading menu. The Create a New Foreign Key wizard opens. Complete the wizard.

The Create a New Foreign Key wizard asks for the following information:

Table 22-3: Inputs to Create a New Foreign Key wizard

Input	Description
Database	Database in which the foreign key table is located.
Table	Table in which the foreign key is located.
Reference type	The foreign key references either a primary key or a uniqueness constraint.
Columns to match	The local columns to match against the foreign columns.
Name	A name for the foreign key.

Shortcuts



Select the Foreign Keys folder. In the right pane, double-click the Add Foreign Key icon.

Example of a Foreign Key

Figure 22-11 illustrates the relationship of a primary key and foreign key. *titleauthor* is the foreign table; *authors* is the primary table. *titleauthor.au_id* is the foreign key; *authors.au_id* is the primary key. To create this relationship:

1. Start the Create a New Foreign Key wizard under *titleauthor*.
2. On the first page of the wizard, specify *pubs2* as the database and *authors* as the table.
3. On the second page of the wizard, select *au_id* as the foreign column and *au_id* as the primary column.
4. Complete the rest of the wizard.

titleauthor (referencing table)			authors (referenced table)		
au_id	title_id	au_ord	au_id	au_lname	au_fname
172-32-1176			172-32-1176		
213-46-8915			213-46-8915		
213-46-8915			228-95-7766		
486-29-1786			274-80-9391		
486-29-1786			486-29-1786		

Foreign key column
Datatype must match referenced column.

Referenced columns must have a primary key or uniqueness constraint.

Figure 22-11: Primary key - foreign key relationship

Displaying the Properties of A Foreign Key

You can change the name of a foreign key. No other properties can be changed. To display the properties of a foreign key:

1. In the Foreign Keys folder for the table you are working in.
2. Select the foreign key.
3. From the File menu, choose Properties.

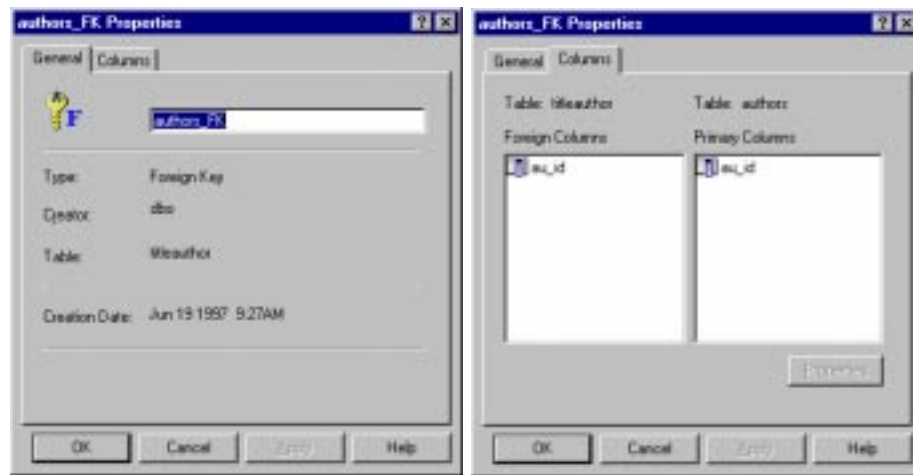


Figure 22-12: Foreign Key property sheet

Creating Unique Constraints

A unique constraint helps maintain data integrity. It ensures that for specified columns, no two rows in a table will have the same non-null value. Unique constraints are created as indexes on a table.

To create a unique constraint:

1. In the table you are working in, select the Unique Constraints folder.
2. From the File menu, choose New; then from the cascading menu, choose Unique Constraint. The Create a New Unique Table Constraint wizard opens.

The Create a New Unique Table Constraint wizard asks for the following information:

Table 22-4: Inputs to Create New Segment wizard

Input	Description
Name	Name for the new constraint.
Columns	Select the columns that are part of the constraint.
Segment	Specify the segment that the constraint will be on.
Fill factor	Optionally, specify adjustment of the fill factor and a value. Mutually exclusive with Maximum Rows Per Page.
Rows per page	Optionally, specify a value for maximum rows per page. Mutually exclusive with Fill Factor.
Clustered index	If there are no clustered indexes on this table, you can designate this constraint to be clustered.

Shortcut



Select the Unique Constraints folder. In the right pane, double-click the Add Unique Constraint icon.

Viewing the Properties of a Unique Constraint




You can change the name of a unique constraint and you can change the segment it is placed on. To change the segment of a unique constraint, follow the procedure described in “Placing an Index on a Segment” on page 24-7. To view the properties of a unique constraint:

1. Select the unique constraint you want to view.
2. From the File menu, choose Properties.



Figure 22-13: Unique Constraint property sheet

Shortcut

	Double-click the unique constraint icon.
	Select the unique constraint icon, and select the Properties toolbar button.
	Click the right mouse button over the unique constraint icon. Choose Properties from the shortcut menu.

Recompiling Procedures and Triggers

After you define a table, you can recompile each procedure and trigger that uses the table. The recompilation is done the next time the procedures or triggers run. As you add indexes or change the database in ways that affect its statistics, compiled procedures and triggers may lose efficiency. Recompiling procedures and triggers can optimize efficiency.

To recompile procedures and triggers:

1. Select the table for which you want to recompile procedures and triggers.
2. From the File menu, choose **Recompile**. The next time the procedures and triggers associated with this table are run, they will be recompiled.

You can also force a procedure to recompile every time it is run by choosing the recompile option when you create the procedure.

Checking Table Consistency

The Database Consistency Checker (**dbcc**) is a set of utility commands for checking the logical and physical consistency of a database. The **dbcc** command has several options to check table consistency.

This section gives general instructions for executing the **dbcc** command for table consistency and then describes each command option.

To check table consistency:

1. Select the icon of the table you want to check.
2. From the File menu, choose **Check Consistency**. The **Create DBCC Command(s)** wizard opens.
3. Complete the wizard.

The command option descriptions that follow include sample report output.

Shortcuts



Click the right mouse button over the icon of the table.
Choose **Check Consistency** from the shortcut menu.

Consistency Check Options

The following consistency check options are available:

- Overall Table Consistency
- Table Allocations
- Reindex
- Fix Text

Overall Table Consistency

Overall Table Consistency corresponds to the `dbcc checktable` command. This option checks that index and data pages are correctly linked, indexes are properly sorted, all pointers are consistent, and the data rows on each page have entries in the first page of an allocation map. The entries match their respective locations on the page.

To skip checking nonclustered indexes, select the Ignore Non-clustered Indexes box. If you leave the box unselected, the `dbcc` command checks all indexes.

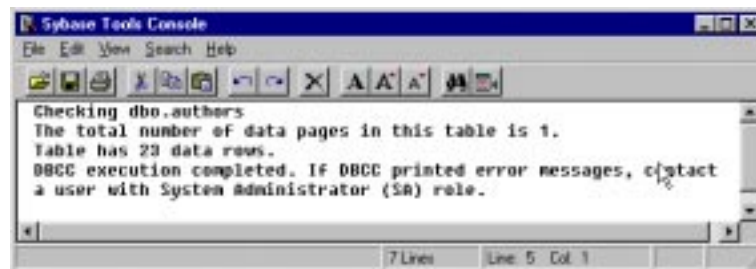


Figure 22-14: Overall Table Consistency - sample output

Table Allocations

Table Allocations corresponds to the `dbcc tablealloc` command. This option checks to see if index and data pages are correctly allocated, that no page is allocated that is not used, and that no page is used that is not allocated. The following report types are available:

- To report allocation pages listed in the Object Allocation Map pages for the table, select `Optimized`.
- To report all types of allocation errors, select `Full`.
- To report pages that are referenced but not allocated in the extent, select `Fast`.

To fix allocation errors, select `Fix Allocation Errors`. Sybase Central automatically places the database in single-user mode while executing `dbcc tablealloc` and then returns the database to multiple-user mode when processing is complete.

```

Sybase Tools Console
File Edit View Search Help
[Icons]
Checking dbo.authors
The total number of data pages in this table is 1.
Table has 23 data rows.
DBCC execution completed. If DBCC printed error messages, contact
a user with System Administrator (SA) role.

*****
TABLE: dbo.authors          OBJID = 16883888
INDID=1  FIRST=569         ROOT=377      SORT=1
      Data level: 1.  1 Data pages allocated and 1 Extents allocated.
      Indid   : 1.  1 Index pages allocated and 1 Extents allocated.
INDID=2  FIRST=265         ROOT=265      SORT=1
      Indid   : 2.  1 Index pages allocated and 1 Extents allocated.
TOTAL # of extents = 3

Statistical information for this run follows:
Total # of pages read = 6
Total # of pages found cache = 6
Total # of physical reads = 0
Total # of saved I/O = 0
DBCC execution completed. If DBCC printed error messages, contact a user
with System Administrator (SA) role.
[Icons]
24 Lines  Line 23 Col 1

```

Figure 22-15: Table Allocations - sample output

Reindex

Reindex corresponds to the `dbcc reindex` command. This option checks the integrity of indexes on user tables. It drops and rebuilds indexes it suspects are corrupt.

Fix Text

Fix Text corresponds to the `dbcc fix_text` command. This option upgrades *text* values after an Adaptive Server character set has been changed to a multibyte character set.

Managing Columns in the Columns Folder

In addition to managing columns in the table editor, you can manage columns in the Columns folder. From the Column property sheet, you can manage bindings to rules, defaults, and check constraints.

Viewing the Properties of a Column

To view the properties of a column:

1. Select the column that you want to view.
2. From the File menu, choose Properties. The Column property sheet opens.



Figure 22-16: Column Properties: General tab

The General tab lists descriptive information about the column.

Binding a Rule to a Column

To bind a rule to a column:

1. Select the column that you want to bind a rule to.
2. Open the property sheet as described in “Viewing the Properties of a Column” on page 22-30.
3. Select the Rules and Defaults tab.



Figure 22-17: Column Properties: Rules and Defaults tab

4. In the Rule drop-down list select the rule you want to bind to this column.
5. Click OK.

Unbinding a Rule from a Column

To unbind a rule from a column:

1. Select the column that you want to unbind a rule from.
2. Open the property sheet as described in “Viewing the Properties of a Column” on page 22-30.
3. Select the Rules and Defaults tab.
4. In the Rule drop-down list select None.

5. Click OK.

Binding a Default to a Column

To bind a default to a column:

1. Select the column that you want to bind a default to.
2. Open the property sheet as described in “Viewing the Properties of a Column” on page 22-30.
3. Select the Rules and Defaults tab. (See Figure 22-17.)
4. In the Default group box, select an existing default in the Named box, or enter the code for a default in the Clause text box. The code is the SQL that would come after the as keyword in a create default command.
5. Click OK.

Unbinding a Default from a Column

To unbind a default from a column:

1. Select the column that you want to bind a default from.
2. Open the property sheet as described in “Viewing the Properties of a Column” on page 22-30.
3. Select the Rules and Defaults tab. (See Figure 22-17.)
4. In the Default group box, select None.
5. Click OK.

Deleting Check Constraints

In the Check Constraints tab of the Column property sheet, you can view the SQL code of the check constraints on a column and you can delete a check constraint.

To delete a check constraint:

1. Select the column that has the constraint you want to delete.
2. Open the property sheet as described in “Viewing the Properties of a Column” on page 22-30.
3. Select the Check Constraints tab.

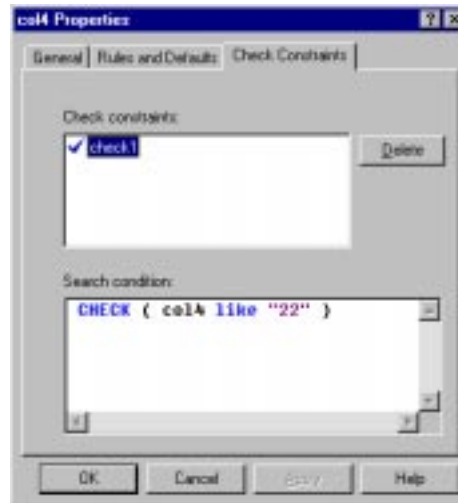


Figure 22-18: Column Properties: Check Constraints tab

4. In the Check Constraints list, select the constraint you want to delete. The code of the constraint is displayed in the Search Condition box.
5. Click Delete.

Creating and Executing DDL Scripts

When you generate DDL for a table, it includes the definitions of the columns and any constraints or keys. Generate DDL for indexes or triggers directly from those objects.

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

23

Managing Proxy Tables

What's in This Chapter

Adaptive Server supports a special kind of table called a proxy table using Component Integration Services. A proxy table exists in an Adaptive Server database, just like any other user table. However, unlike other user tables, the data for a proxy table exists on a remote server. This chapter explains how to manage proxy tables.

Understanding Proxy Tables

A proxy table is a user table that allows you to access data in a remote table, view, RPC, or file. Sybase Central does not support creation of proxy tables based on files. A proxy table has all the attributes of a user table, such as columns, indexes, triggers, and so forth, but it does not contain any data locally.

You can create proxy tables as follows:

- Map a proxy table to an existing remote table or view (relationship A in Figure 23-1.) The columns in these tables must match the properties of the columns in the remote table or must be the Adaptive Server equivalent of the remote column. Component Integration Services treats views as if they were tables without any indexes.
- Create a proxy table that maps to an RPC (relationship B in Figure 23-1.) The result set of the RPC is treated as a read-only table.
- Create a proxy table and a remote table at the same time (relationship C in Figure 23-1.). The columns on the remote table are created with the properties you define for the proxy table.




Figure 23-1: Proxy table mapping options

After you create a proxy table, you work with it as you would with any other table. For complete information about proxy tables, see *Component Integration Services User's Guide for Adaptive Server and OmniConnect*.

Component Integration Services Options

Adaptive Server is integrated with Component Integration Services, however the extent of services available depends on the options purchased for your system. A basic installation allows Adaptive Server to create proxy tables on a remote Adaptive Server. The extended option allows you to create proxy tables on heterogeneous remote servers, that is, remote servers that are non-Sybase databases.

In addition to Adaptive Server, Sybase Central can connect to and manage an OmniConnect server. An OmniConnect server is similar to Adaptive Server, except it cannot store data in user tables. It can only use proxy tables.

Creating Proxy Tables

The basic Adaptive Server installation does not allow you to create proxy tables. You must set configuration parameters and create supporting objects first.

Prerequisites for Creating Proxy Tables

Before you create a proxy table, you must do the following:

- Set the `enable cis` configuration parameter to 1 and restart Adaptive Server. See “Setting Configuration Parameters” in Chapter 4, “Managing the Adaptive Server Operating Environment”.
- Determine the type of RPC handling you want to use and if applicable, set the `cis rpc handling` configuration parameter.
- Add a remote server to which you will connect using Component Integration Services. See “Adding a Remote Server” in Chapter 6, “Managing Remote Servers”.
- Map a local login to a remote login for the remote server. The type of mapping depends on the type of RPC handling you use. See “Managing Remote Server Logins” in Chapter 6, “Managing Remote Servers”.

Creating a Proxy Table

To create a proxy table:

1. In a database, open the Proxy Tables folder. (If the Proxy Tables folder is not visible, you may need to configure Adaptive Server to use Component Integration Services. See “Prerequisites for Creating Proxy Tables” on page 23-3.
2. From the File menu, choose New; then from the cascading menu, choose Proxy Table.
3. Follow the instructions of the wizard.

The Create a New Proxy Table wizard asks for the information in Table 23-1. Additional details about creating proxy tables are provided in the following sections.

Table 23-1: Inputs to the Create a New Proxy Table wizard

Input	Description
Table basis	Options: - Existing table or view - RPC - New proxy table and new remote table.
Remote server	The remote server on which the remote object exists or will be created.
Remote database	If applicable, the remote database in which the remote object exists.
Remote table or view	The existing tables or views to which the proxy tables will be mapped.
Remote procedure	A remote procedure, the result set of which the proxy table will be mapped. You can only specify one for each time you run the wizard. The wizard lists remote procedures if it can obtain them from the remote server. If it cannot list procedures, you can specify one manually.
Naming convention	For multiple proxy tables, a convention for naming the tables. Names must conform to naming conventions for Adaptive Server objects. For heterogeneous servers, be aware of how a prefix or suffix will interact with the names on the remote server.
Name case	For multiple proxy tables, choose upper or lower case
Table editor	For RPC mapping and new remote tables.

4. If applicable, complete the definition of the table in the table editor.

Shortcuts



Select the Proxy Tables folder. In the right pane, double-click the Add Proxy Table icon.

Considerations for Creating Proxy Tables

If you create multiple proxy tables based on existing tables or views, the proxy tables are created with column definitions identical to those of the remote objects. In all other cases, you must complete the definition of the proxy table in the table editor.

The table editor for proxy tables is the same as that for user tables. For information about using the table editor, see “Using the Table Editor” in Chapter 22, “Managing User Tables”.

The following notes apply to creating proxy tables:

- If the proxy table is based on an existing table or view, you cannot change any column definitions. You can specify the name for the proxy table and create new columns.
- If you are creating a proxy table based on an RPC, base the table definition on the result set that you expect from the RPC.
- If you are creating a proxy table and remote table at the same time, you have full options for creating the proxy table. The remote table is created using the definitions for the proxy table or their equivalents on heterogeneous servers.
- If a column name in a remote table or view is an Adaptive Server keyword or contains an illegal character, creation of the proxy table fails. To workaroud this problem, create the proxy table using isql or SQL Advantage and use the `set quoted identifier` parameter.
- Adaptive Server tables can have a maximum of 255 columns. If a remote table has more than 255 columns, you cannot create a proxy table using Sybase Central. To create a proxy table against a remote table with more than 255 columns, use SQL Advantage and specify up to 255 columns from the remote table to include in the proxy table.
- If the server class you specified for a remote server is incorrect creation of proxy tables mapped to that server may fail. For example, if a remote server is a DB2 server and you specify server class “generic”, you may not be able to create proxy tables mapped to that remote server. If an error of this sort occurs, the error messages sent by Adaptive Server may be confusing and may not identify the actual error.

Creating Existing Tables Compared to Creating New Tables

It is important to understand the ramifications of creating a new proxy table and new remote table compared to creating a proxy table based on an existing remote table. If you create a proxy table based on an existing table and you decide to delete the proxy table, there is no effect on the remote table. However, if you create a new proxy table and remote table, when you delete the proxy table the remote table is also deleted.

If you create new remote tables based on proxy tables, consider naming the proxy table with a special convention that indicates that deleting it will delete the remote table.

Deleting a Proxy Table

◆ **WARNING!**

If you delete a proxy table that was created with the new local and remote table option, the remote table is also deleted. You will lose all data stored in the remote table.

If you try to delete a proxy table that was created with the new local remote option and the remote server is inaccessible, you get an error message indicating that the remote server is inaccessible. As a result, the proxy table is deleted, but the remote table is not deleted.

To delete a proxy table, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Verifying the Status of a Proxy Table

It is possible that after you create a proxy table, the remote object could be deleted or its definition could change. You can verify that the remote object still exists with the correct definition.

To verify a proxy table:

1. Select the proxy tables you want to verify.
2. From the File menu, choose Verify Table.

Sybase Central compares the proxy table to the remote object. Messages are written to the tools console.

The Verify Table command checks for the following problems:

- The remote object cannot be found.
- The column definitions in a remote table or view do not match the column definitions in the proxy table. (This does not include differences in case for column names.)
- The datatype of a remote column is not compatible with the datatype of the local column.
- The datatype of the remote column is compatible with the local column, but it is not the default mapping for this datatype.
- The datatype of the remote column is compatible with the local column, but there may be data loss when data is converted to the local datatype.

Notes:

- Sybase Central does not verify the indexes on a proxy table. To verify indexes, verify the table from the command line.
- Sybase Central verifies the existence of a remote procedure. It cannot verify that the columns in a proxy table match the result set of a call to the procedure.

Managing Proxy Tables

The procedures for editing properties of a proxy table are identical to those for a user table. You can also work with columns, triggers, indexes and constraints the same way you would a user table. However, the effect of your activities may be different than they would be in a user table. For example, permissions, indexes and new columns are propagated to (created on) remote tables, but triggers and constraints are not. For information about managing user tables, see Chapter 22, "Managing User Tables." For information about the effect of changes made to proxy tables, see *Component Integration Services User's Guide for Adaptive Server and OmniConnect*.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in "Generating DDL Scripts" on page 3-6.

To execute DDL files, use one of the procedures described in "Executing DDL Scripts" on page 3-6.

24 Managing Indexes

What's in This Chapter

This chapter describes the following:

- Creating an index
- Displaying index properties
- Placing an index on a segment
- Checking index consistency
- Updating statistics

An **index** is created on one or more table columns and points to the place where the column data is stored on disk. Indexes speed up data retrieval and are useful for enforcing referential integrity. A table can have more than one index.

The index attributes you can define when you create an index with Sybase Central are:

- Clustered or nonclustered
- Unique or nonunique

A **clustered index** is sorted on an ongoing basis so that the physical order of rows is the same as the logical order. The bottom or leaf level of the index contains the actual data pages of the table. Finding data using a clustered index is faster than using a nonclustered index.

Create a clustered index before creating nonclustered indexes because nonclustered indexes are rebuilt when a clustered index is created. A table can have only one clustered index.

A **nonclustered index** stores key values and pointers to data pages rather than the data itself. The physical order of the rows is not the same as the indexed order. A table can have up to 249 nonclustered indexes.

A **unique index** is one in which no two rows can have the same index value. A unique index is useful as an integrity check on unique data.

A **nonunique index** is one in which two or more rows can have the same index value.

You can combine these characteristics; for example, you can create a unique, nonclustered index.

Once you create an index, you can change the segment on which it is allocated and you can check index consistency. When you modify the index, the old index is deleted from the database and replaced with the new version.

Creating an Index

Privileges	Only a table owner or a user with create index permission on a table can create an index.
------------	--

To create a new index:

1. In the table in which you want to create the index, select the Indexes Folder.
2. From the File menu, choose New; then, choose Index from the cascading menu. The Create a New Index wizard opens.

The Create a New Index wizard asks for the following information:

Table 24-1: Inputs to the Create a New Index wizard

Input	Description
Name	Name for the new index.
Columns	The columns on which to create the index.
Segment	Database segment on which to store the index. A clustered index must be created on the same segment as its table, otherwise, the table moves to the new segment.
Cache	Cache to which you want to bind the index.
Unique	Specifies that the index is unique and no two rows can have the same key (index) value. When this feature is checked, the system checks for duplicate key values when the index is created and each time data is added with an insert or update command. You cannot create a unique index on a column that includes duplicate values or more than one null value. This kind of index makes sense when the data itself is unique.
Clustered	Specifies that the index is clustered and has data, rather than pointers, in its leaf pages. A table can have only one clustered index.
Duplicate key	Accept or reject duplicate key values in a unique index.

Table 24-1: Inputs to the Create a New Index wizard

Input	Description
Duplicate row	Accept or reject rows with duplicate values in a non-unique clustered index.
Fill factor	Set the fill factor as a percentage. Mutually exclusive with setting rows per page. See “Fill Factor” on page 24-4.
Rows per page	Set the rows per page. Mutually exclusive with setting the fill factor. See “Maximum Rows per Page” on page 24-5.
Cache strategy	MRU or large buffer prefetch. See “Cache Strategy” on page 24-5.
Is data sorted	Indicate if data in the table has already been sorted. See “Presorted Data” on page 24-4.

Shortcuts

Select the Indexes folder. Double-click the Add Index icon.

► Note

You can also create a unique index as a unique constraint. See “Creating Unique Constraints” in Chapter 22, “Managing User Tables”.

◆ WARNING!

Use the Duplicate Keys and Duplicate Rows features with caution. An update that creates a duplicate key row or a duplicate row can result in data being overwritten.

Fill Factor

The default fill factor is determined by the value of the **fill factor** configuration parameter. There is seldom a reason to change the **fill factor** parameter.

If the default value is 0, you can enter a percentage from 0 to 100. If the default value is other than 0, only percentages from 1 to 100 are valid. A fill factor of 0 does not mean the pages are empty, but means there is some room for additional data. Use a fill factor of 100 only for read-only tables to which no additional data will **ever** be added. Small fill factors cause each index to take more storage space and result in index pages that are not completely full.

► **Note**

The fill factor percentage is relevant only at the time an index is created. As the data changes, the pages are not maintained at any particular level of fullness.

Presorted Data

To indicate that data has been presorted for faster index building, check Sorted Data. This option works only if the table data has been loaded in presorted format with the bulk copy utility and the index is one of the following:

- Clustered
- Unique, nonclustered
- Nonunique, nonclustered, and there are no duplicate keys

Cache Strategy

You can select a cache strategy for the index. This selection overrides the default strategy, determined by the Adaptive Server optimizer, for reading data pages from an index into the buffers in data cache. The following selections are available:

- **Most Recently Used Replacement (MRU)**—This selection specifies that Adaptive Server uses the most recently used strategy for determining where in cache to place data pages when reading in new data.

If you clear the check box, Adaptive Server reads new pages into the MRU end of the chain of buffers in cache. Subsequent reads move the pages along the chain towards the least recently used (LRU) end until they are flushed out by new reads at the MRU end. If you select Most Recently Used Replacement, Adaptive Server reads new pages into the LRU end. They are used and immediately flushed when a new page enters the MRU end.

This strategy is advantageous when a page is needed only once for a query. It tends to keep such pages from flushing out other pages that can potentially be reused while still in cache.

- **Large Buffer Prefetch**—This selection applies if one or more large buffer pools is defined in the default cache or, if the index is bound to a named cache, in the named cache. A large buffer pool is one that has buffers larger than the 2K default, as specified in the Cache property sheet. If you select Large Buffer Prefetch, the Adaptive Server optimizer can fetch data in I/Os of as many as eight 2K data pages at a time instead of the default of one page at a time.

This strategy is advantageous for data that is stored and accessed sequentially; for example, it can improve performance for queries that scan the leaf level of a nonclustered index.

Maximum Rows per Page

Optionally, to override the default assigned by Adaptive Server for the maximum number of index rows to store on a leaf level page, enter a number in the Maximum Rows Per Page box. This option is mutually exclusive with the Fill Factor option. Also, unlike the Fill Factor option, you can change the Maximum Rows Per Page value at any time.

The default value, 0, causes Adaptive Server to create clustered indexes with full pages and nonclustered indexes with full leaf pages.

For a clustered index, values between 0 and 256 are valid. For a nonclustered index, the largest allowable Maximum Rows Per Page value depends on the size of the index key. To approximate the largest allowable value, subtract 32 from the page size and divide the resulting number by the index key size.

Deleting an Index

To delete an index, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

Displaying Index Properties

To display index properties:

1. Select the icon of the index you want to display.
2. From the File menu, choose Properties.

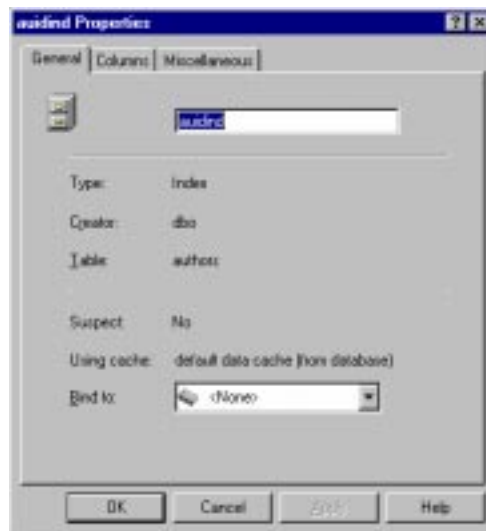





Figure 24-1: Index property sheet

Shortcuts

	Double-click the index icon.
	Select the index icon. Then select the Properties toolbar button.
	Click the right mouse button over the index icon. Choose Properties from the shortcut menu.

Placing an Index on a Segment

Using a segment to put an index on a specific database device can improve Adaptive Server performance and can give increased control over placement, size, and space usage of database objects. For example:

- If you put a table on one device, and its nonclustered indexes on a device on another disk controller, the time required to read or write to the disk can be reduced.
- If you put tables and indexes only on specific segments, those objects cannot grow beyond the space available on the devices represented by the segments.
- If you put tables and indexes in their own segments that map to reserved devices, other objects cannot contend for space with them.

To place an index on a segment:

1. Display the Index property sheet, as described in “Displaying Index Properties” on page 24-6.
2. Click the Miscellaneous tab.



Figure 24-2: Index property sheet: Miscellaneous tab

3. From the Segment drop-down list, select the name of a segment. This is the segment on which the index puts new data.
4. Click Apply. Future growth of the index now occurs on this segment.

Setting Maximum Rows Per Page

To override the default assigned by Adaptive Server for the maximum number of index rows to store on a data page:

1. Select the index on which you want to set the rows per page.
2. Open the property sheet for the index as described in “Displaying Index Properties” on page 24-6.
3. Select the Miscellaneous tab.
4. Enter a value in the Rows per Page box.
5. Click OK.

For a table with clustered indexes, values between 0 and 256 are valid. For a table with nonclustered indexes, the largest allowable Rows Per Page value depends on the size of the index key. To approximate the largest allowable value, subtract 32 from the page size and divide the resulting number by the index key size.

Checking Index Consistency

The Database Consistency Checker (dbcc) is a set of utility commands for checking the logical and physical consistency of a database. The dbcc command has an `indexalloc` option to check allocation structures for an index. You can check the current index to see that all data pages are correctly allocated, that no page is allocated that is not used, and that no page is used that is not allocated. This section gives general instructions for executing the dbcc command for index consistency and then describes each command option.

Privileges	Only a database owner can check index consistency.
------------	--

To check index consistency:

1. Select the icon of the index you want to check.
2. From the File menu, choose Check Consistency. The Create DBCC Command(s) wizard opens.
3. To fix allocation errors, check Fix Allocation Errors.
4. The report types correspond to the dbcc `indexalloc` command. The following report types are available:
 - Select Optimized to report allocation pages listed in the Object Allocation Map pages for the index. (For information about the Object Allocation Map, see *System Administration Guide*.)
 - Select Full to report all types of allocation errors.
 - Select Fast to report pages that are referenced but not allocated in the extent.
5. Click OK to check index consistency.

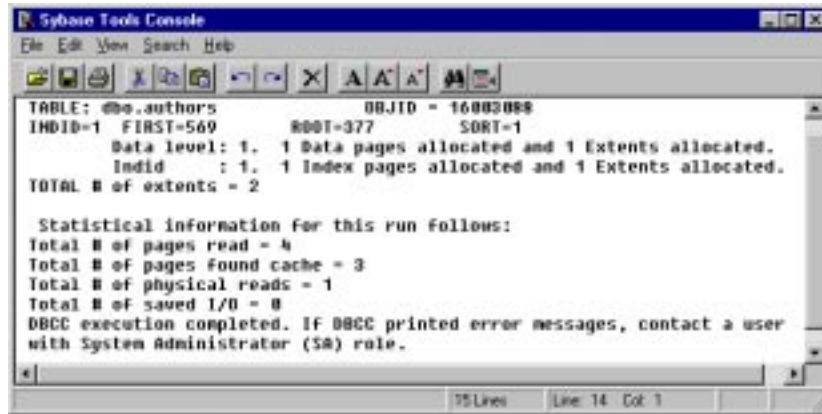


Figure 24-3: Output from Full option of index consistency check

You can copy and paste report output into a word processor or text editor.

Shortcuts



Click the right mouse button over the index icon. Choose Check Consistency from the shortcut menu.

Updating Statistics

When you update statistics, Adaptive Server updates information about the distribution of key values in indexes associated with a table. Update statistics if a great deal of data in an indexed column has been added, changed, or removed. Update statistics after adding new rows to a table whose rows have been deleted with `truncate table`.

Privileges	Only a table owner or a database owner can update index statistics.
------------	---

To update statistics:

1. Select the icon of the index you want to update.
2. From the File menu, choose Update Statistics.

Shortcuts



Click the right mouse button over the index icon. Choose Update Statistics from the shortcut menu.

► **Note**

After running `update statistics`, recompile the table for which the index is defined. This ensures that the procedures and triggers that use the index will use the new key distribution.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

25 Application Tuning

What's in This Chapter

Application tuning lets you tune Adaptive Server performance at the application level. This chapter discusses the following:

- Application tuning concepts
- Engine groups
- Execution classes

Application Tuning Concepts

Most performance-tuning techniques enhance performance at the system or query level. Adaptive Server also lets you tune at the application level. You can tell Adaptive Server the performance trade-offs you want to make among individual connections and sessions in mixed work load environments. For example, if one or more execution objects (client applications, Adaptive Server logins, or stored procedures) have a higher priority than others, you can improve response time and throughput for them by assigning execution attributes. Adaptive Server considers the execution attributes when it places the execution object in one of three priority run queues. You can also use engine affinity to constrain execution objects to specific Adaptive Server engines.

Each client application or login can initiate many Adaptive Server tasks, so in a single-application environment, you can tune for performance at the login and application level, enhancing performance for chosen connections or sessions. In a multiple-application environment, such tuning can improve performance for selected applications and improve performance for some running Adaptive Server processes.

You establish which client applications, logins, and stored procedures receive execution preference when you create a hierarchy among execution objects.

Assigning execution preferences can take on varying levels of complexity depending on the mix of applications in the Adaptive Server environment.

What is Execution Preference?

Suppose you decide that slower response time is okay when login A uses `isql`, but that login B must be able to get results from `isql` as fast as possible. In this case, you would be willing to sacrifice response time for login A to enhance response time for login B.

You can do this by assigning an execution class, which is a collection of attributes representing task priority and engine affinity. In the current example, you can tell Adaptive Server that you want it to give preference to login B when it executes `isql` by binding Login B to the preferred execution class EC1 with a scope of `isql`.

Execution Class Basics

You implement execution preference by managing two Adaptive Server objects: execution classes and engine groups. There are two types of execution classes: predefined and user defined.

Predefined Execution Classes

There are three predefined execution classes, EC1, EC2, and EC3. They are composed of a fixed mix of execution attributes such that:

- EC1 has attributes for above-average performance
- EC2 has attributes for average performance
- EC3 has attributes for below-average performance

You can associate an execution class with an execution object by binding the execution object to the execution class. If you do not assign an execution class to an execution object, it has the execution attributes of EC2 by default.

It is a good idea to apply execution class changes one at a time, analyze the effect on the system, and verify that the results are as you expected.

User-defined Execution Classes

You can define your own execution classes. You might want to do so:

- If EC1, EC2, and EC3 do not accommodate your critical execution object's needs.
- To associate execution objects with a particular group of engines as described in “Engine Affinity Basics” on page 25-3.

You can bind an execution object to a user-defined execution class just as you would with a predefined execution class.

Engine Affinity Basics

Engine affinity is an attribute of an execution object. It means the execution object can execute only on the engines in a particular **engine group** (one or more Adaptive Server engines.) Conceptually, it is similar to binding. This capability only has value on multiprocessor systems running multiple Adaptive Server engines.

You create engine affinity by associating an execution class with an engine group.

In a multiprocessor environment running multiple Adaptive Server engines, any available engine can process the next execution object in the runnable task queue. If you associate execution objects in an execution class with an engine group (a subset of the total number of engines), you limit their access to processing resources. This makes engines that are not part of the engine group more available for use by the remaining execution objects, which presumably can execute on any Adaptive Server engine. The performance of the unaffiliated objects improves because they have the benefit of the resources that the objects in the execution class are deprived of.

If you want to set up engine affinity, you need to perform the following steps:

1. Create an engine group as described in “Creating an Engine Group” on page 25-9 (or use the predefined engine groups LASTONLINE and ANYENGINE.)
2. Bind execution objects to an execution class as described in “Binding an Execution Object to an Execution Class” on page 25-17.
3. Associate the execution class with an engine group as described in “Associating an Execution Class with an Engine Group” on page 25-17.

Figure 25-1 illustrates the use of engine affinity. Engine E2 and engine E3 are part of engine group G1. A user-defined execution class called X1 is associated with engine group G1. The logins *mgr1* and *mgr2* are bound to X1. The logins *support1* and *support2* are not bound to an engine group.




Figure 25-1: Using an engine group

Assume that *mgr1* and *mgr2* do not need immediate access to resources, but *support1* and *support2* need ready access to resources. You could bind *support1* and *support2* to EC1 for the highest priority, or not bind them to an execution class, leaving them with the attributes of EC2 by default. In either case, they have access to any engine when it becomes available.

To restrict resource access for *mgr1* and *mgr2*, create engine group G1 with a subset of the available engines. Then create execution class X1 and give it an affinity to G1. Bind *mgr1* and *mgr2* to X1. They get CPU resources only when engine E2 or engine E3 are available. Since *mgr1* and *mgr2* cannot use engine E0 or engine E1, the probability that these engines will be available for use by *support1* and *support2* increases.

Scheduling with Execution Classes and Engine Groups

Each execution class is associated with a priority run queue. If the execution classes do not have an engine affinity, Adaptive Server processes tasks in the highest priority run queue before it processes those in the medium priority run queue. It processes the tasks in the low priority run queue last.

Execution classes are associated with the run queues as follows:

- Tasks assigned to EC1 are placed in the high priority run queue.
- Tasks assigned to EC2 are placed in the medium priority run queue.
- Tasks assigned to EC3 are placed in the low priority run queue.

When you create a user-defined execution class, you assign it a priority that places it in one of the run queues associated with the three pre-defined execution classes.

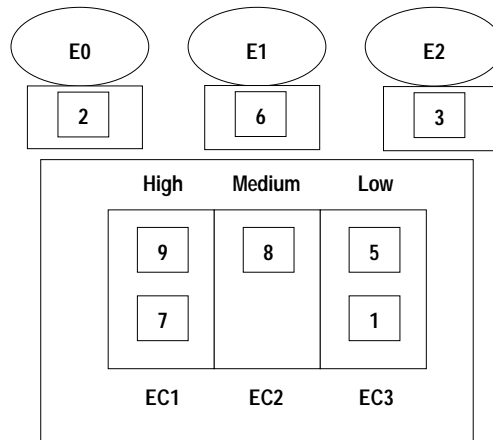


Figure 25-2: How execution class and engine affinity affect execution

Figure 25-2 illustrates tasks in the three priority run queues (high, medium, and low) waiting for access to three engines (E0, E1, and E2) in Adaptive Server. The engines are running tasks 2, 6, and 3. The figure is the basis for example 1 and example 2.

Example 1:

The tasks are bound to execution classes *EC1*, *EC2*, and *EC3*. Assume that there are no user-defined execution classes. Therefore, tasks in the high and medium priority run queues (7, 8, and 9) have no engine affinity and the tasks in the low priority run queue (5 and 1) have affinity to engine E2. (By definition, *EC1* and *EC2* have no engine affinity and *EC3* has affinity to the engine with the highest number.)

Assuming no new tasks, as engines become available tasks get processed in the order 9, 7, 8, 5, 1. However, tasks 5 and 1 get processed only if engine E2 is available. If one of the other tasks is assigned to engine E2, neither 5 nor 1 get processed even if engine E0 or engine E1 is available. If a new task enters the high or medium queues and is waiting for processing time when engine E2 becomes available, the new task will get processed and tasks 5 and 1 will continue to wait.

► **Note**

When you use engine affinity, it is possible that some tasks will not get processed.

Example 2:

Assume the same set of engines and tasks as in example 1. For this example, assume a user-defined execution class *class1* with affinity to an engine group that contains engine E1.

Assume that task 9, which is first in the high priority run queue (A), is in *class1*. This means it can get processed only on engine E1. If engine E1 is busy, but engine E2 is not busy, then engine E2 will process task 7, the next task in the high priority run queue, which has no affinity. If the next engine available is engine E0, it will process task 8, which is in the medium priority queue. If the next engine available is engine E2, it will process task 5, which is in the low priority queue.

Notice that although an application tuner took care to assign the preferred execution class *EC1* to task 9, assigning engine affinity temporarily lowered its execution preference to below that of a task with *EC3*. This effect might be highly undesirable, or it might be just what was intended if the execution class and engine affinity assignments were well thought out.

Scope and Precedence Rules

When you bind a login or an application to an execution class, you can specify a scope for that binding. A scope limits the effect of the execution class to specific applications for a login or specific logins for an application. For example, you can bind a login to an execution class, but only when it executes `isql`. Or you could bind `isql` to an execution class, but only when it is run by the login `sa`.

There are no restrictions on the number of execution classes to which you can bind an execution object. There are also no restrictions on the number of execution classes you can associate with engine groups. Therefore, it is possible that you can create conflicts among different engine affinity and scope assignments. When conflicts occur, Adaptive Server uses scope and precedence rules. Use the rules in this order:

1. Use the precedence rule when the process involves multiple execution object types (procedures, logins, applications.)
2. Use the scope rules when there are multiple class definitions for the same execution object.

Precedence Rule

The execution class of a stored procedure overrides that of a login. The execution class of a login overrides that of a client application.




Figure 25-3: Precedence rule

The precedence of an execution object can be temporarily raised. It cannot be lowered. For example, if a stored procedure has a higher precedence than that of the client application process invoking it, the client process's precedence is temporarily raised to that of the stored procedure for the period of time during which the stored procedure runs. If an execution object invokes a stored procedure with a lower priority than its own, the execution object's priority is not lowered.

Figure 25-4 illustrates an application, *app_1*, that executes with medium precedence. When it calls stored procedure *sp_high*, which has a high precedence, it executes at that high precedence. When it invokes stored procedure *sp_low*, which has a low precedence, it continues to execute at medium precedence.




Figure 25-4: Example of precedence

Scope Rule

When an execution object is bound to multiple execution classes, the one with the most narrow scope has precedence. For example, if you bind login *sa* to *EC2* with a scope of *All* and to execution class *class1* with a scope of *isql*, the attributes associated with the assignment to *class1* take precedence when *sa* is executing *isql*.

Resolving Precedence Conflicts

Adaptive Server uses the following rules to resolve conflicting precedence when multiple execution objects and execution classes have the same scope.

- Execution objects not bound to a specific execution class are given the execution attributes of *EC2* by default, with no engine affinity.
- An execution object bound to an execution class has higher precedence than one with default attributes. (An object bound to *EC3* has precedence over an object with *EC2* attributes if they are competing for the last online engine.)
- If a client application and a login are bound to different execution classes, the login's execution class takes precedence over the client application's execution class (from the precedence rule, see "Scope and Precedence Rules" on page 25-6).
- If a stored procedure and a client application or login have different execution classes, Adaptive Server uses the execution class with the higher priority when it executes the stored procedure (from the precedence rule).
- If there are multiple definitions for the same execution object, the one with a narrower scope has the highest priority (from the scope rule).

Managing Engine Groups

An engine group is a set of one or more Adaptive Server engines. The significance of engine groups is the ability to create an affinity between an execution class and an engine group. By creating this affinity, you can tune Adaptive Server performance for applications, logins, and stored procedures.

There are two default engine groups: *LASTONLINE* and *ANYENGINE*. *ANYENGINE* contains all engines on Adaptive Server. *LASTONLINE* contains the engine with the highest number. For example, if Adaptive Server has four engines (0, 1, 2, 3), *LASTONLINE* contains engine 3.

Creating an Engine Group

To create an engine group:

1. Open the Engine Groups folder.
2. From the File menu, choose New; then choose Engine Group.
3. Follow the instructions of the wizard.

The Create a New Engine Group wizard asks for the following information:

Table 25-1: Inputs to Create a New Engine Group wizard

Input	Description
Name	Name for the new engine group.
Engines	The Adaptive Server engines to put in the engine group.

► **Note**

The number of available engines is controlled by the Adaptive Server configuration parameters `max available engines` and `min available engines` and is limited by the number of CPUs on the machine Adaptive Server is running on.

Shortcut



Select the Engine Groups folder. In the right pane double-click the Add Engine Group icon.

Deleting an Engine Group

You cannot delete the default engine groups -- *ANYENGINE* and *LASTONLINE*. To delete other engine groups, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

If you remove the engines from an engine group, the engine group is automatically deleted.

Displaying the Properties of an Engine Group

To display the properties of an engine group:

1. Open the Engine Groups folder.
2. Select the engine group you want to display.
3. From the File menu, choose Properties.

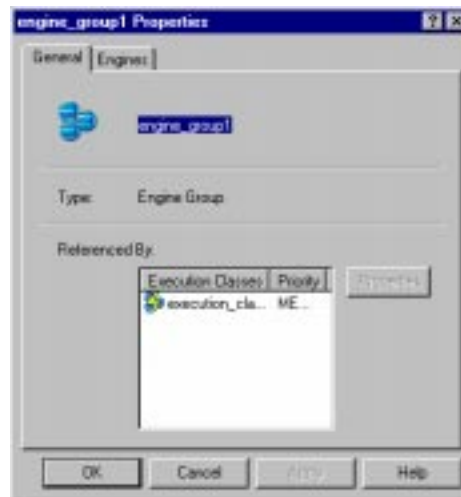


Figure 25-5: Engine Group property sheet

► **Note**

You cannot edit the properties of the default engine groups -- *ANYENGINE* and *LASTONLINE*.

Shortcuts



Double-click the engine group icon.



Select the engine group icon. Click the Properties button in the toolbar.



Click the right mouse button over the engine group icon. Choose Properties from the shortcut menu.

Changing the Engines in An Engine Group

To change the engines in an engine group:

1. Open the property sheet for the engine group you want to change as described in “Displaying the Properties of an Engine Group” on page 25-10.
2. Select the Engines tab.

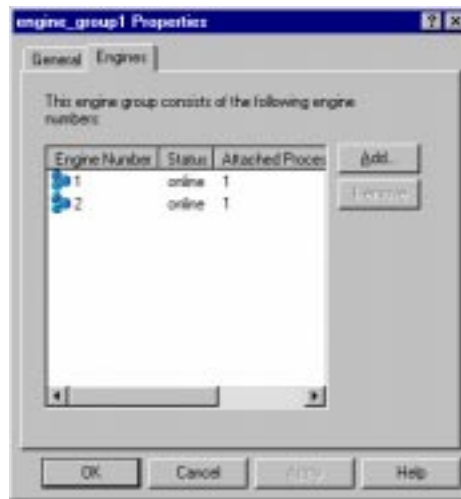


Figure 25-6: Engine Group property sheet: Engines tab

Adding An Engine to An Engine Group

To add an engine to an engine group:

1. Click the Add button.

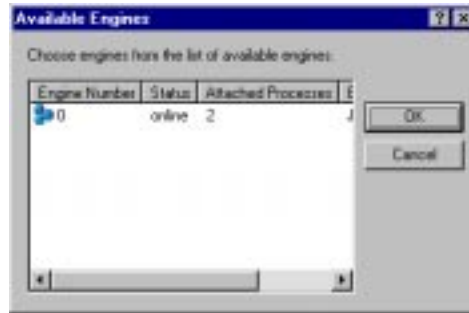


Figure 25-7: Available Engines dialog box

2. In the Available Engines dialog box select the engines you want to add to the engine group.
3. Click OK.

Deleting An Engine from An Engine Group

To delete an engine from an engine group:

1. In the list of engines on the Engines tab, select the engines you want to delete.
2. Click Remove.
3. Click OK.

► **Note**

If you remove all the engines from an engine group, the engine group is deleted.

Creating and Executing DDL Scripts for Engine Groups

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6. No DDL is created for the engine groups LASTONLINE and ANYENGINE.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

Managing Execution Classes

Execution classes let you set execution attributes for applications, Adaptive Server logins, and stored procedures. These attributes, including priority and engine affinity, affect the performance and throughput of the execution objects.

Creating an Execution Class

To create an execution class:

1. Open the Execution Classes folder.
2. From the File menu, choose New; then choose Execution Class.
3. Follow the instructions of the wizard.

The Create a New Execution Class wizard asks for the following information:

Table 25-2: Inputs to Create a New Execution Class wizard

Input	Description
Name	Name for the new execution class.
Priority	Priority for this execution class: high, medium, or low.
Engine groups	The engine groups to associate with the execution class.

Shortcut



Select the Execution Classes folder. In the right pane double-click the Add Execution Class icon.

Deleting an Execution Class

To delete an execution class, follow the procedure for deleting objects as described in “Deleting an Object” on page 3-4.

You cannot delete the engine groups *EC1*, *EC2*, and *EC3*.

Displaying the Properties of an Execution Class




To display the properties of an execution class:

1. Open the Execution Classes folder.
2. Select the execution class you want to display.
3. From the File menu, choose Properties.



Figure 25-8: Execution Class property sheet

Shortcuts

-
- | | |
|---|--|
|  | Double-click the execution class icon. |
|---|--|
-
- | | |
|---|--|
|  | Select the execution class icon. Then, click the Properties button in the toolbar. |
|---|--|
-
- | | |
|---|---|
|  | Click the right mouse button over the execution class icon. Choose Properties from the shortcut menu. |
|---|---|
-

Associating an Execution Class with an Engine Group

To change the engine group associated with an execution class:

1. Open the property sheet for the execution class as described in “Displaying the Properties of an Execution Class” on page 25-16.
2. On the General tab, select an engine group from the Engine Group drop-down list (Figure 25-8.)
3. Click OK.

Binding an Execution Object to an Execution Class

To bind an object to an execution class:

1. Open the property sheet for the execution class as described in “Displaying the Properties of an Execution Class” on page 25-16.
2. Select the Bindings tab.



Figure 25-9: Execution class property sheet: Bindings tab

3. Select the type of object you want to bind in the Show Object Bindings For drop-down list.
4. Click the Bind button.

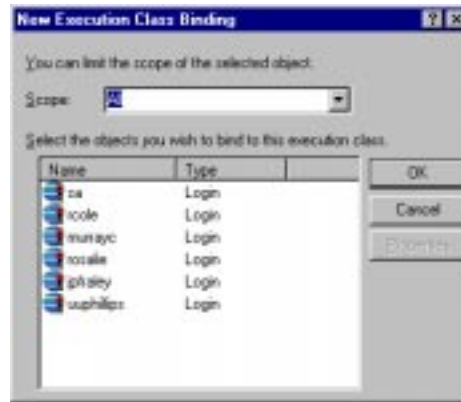


Figure 25-10: New Execution Class Bindings dialog box for logins

5. In the New Execution Class Bindings dialog box, select the objects you want to bind to the execution class.

If you are binding logins or applications you can specify a scope. For more information about specifying scope, see “Specifying the Scope of A Binding” on page 25-18.

If you are binding stored procedures, select the database in which they are located.

6. Click OK.

Specifying the Scope of A Binding

When you bind an application or a login to an execution class, you can specify a scope for the binding. You select the scope from the Scope drop-down list on the New Execution Class Binding dialog box.

The scope for a login specifies the range of applications to which this execution class binding applies. The options are:

- All – The execution class binding applies to all applications run by this login. A scope of All is equivalent to a scope of NULL in the `sp_bindexclass` stored procedure.
- Specific application – The execution class binding applies only when the login runs the specified application.

The scope for an application specifies the logins for which this execution class binding applies. The options are:

- All – The execution class binding applies to any login that runs this application. A scope of All is equivalent to a scope of NULL in the `sp_bindexclass` stored procedure.
- Specific login – The execution class binding applies when this application is run by the specific login.

For more information about scope, including an example and precedence rules, see “Scope and Precedence Rules” on page 25-6.

Unbinding an Execution Object from an Execution Class

To unbind an object from an execution class:

1. Open the property sheet for the execution class as described in “Displaying the Properties of an Execution Class” on page 25-16.
2. Select the Bindings tab.
3. Select the type of object you want to unbind in the Show Object Bindings For drop-down list.
4. In the list window, select the objects you want to unbind.
5. Click the Unbind button.
6. Click OK or Apply.

Creating and Executing DDL Scripts

To create DDL scripts, follow the procedure described in “Generating DDL Scripts” on page 3-6.

The DDL for engine groups *EC1*, *EC2*, and *EC3* does not include a create statement. It just has bindings.

To execute DDL files, use one of the procedures described in “Executing DDL Scripts” on page 3-6.

26

Interpreting Monitor Data

Whats in This Chapter

This chapter describes the information displayed in the Adaptive Server monitors. It also provides tips on using the information to improve the performance of Adaptive Server and its applications and where to look for further research. The chapter includes the following information:

- Summary of Adaptive Server Monitors
- Definition of Metrics
- Resources for More Information
- Descriptions of each monitor

The descriptions define each data item displayed in the monitor. They also include tuning charts showing major performance indicators and suggestions for improving performance.

Summary of Adaptive Server Monitors

The following table summarizes the Adaptive Server monitors.

Table 26-1: Adaptive Server monitors

Type and Level of Detail	Monitor Name	Description
Server-wide performance overviews	Performance Summary Monitor	Displays hit rates, device I/O, page I/O, and lock statistics.
	Performance Trends Monitor	Plots values for up to 60 sample intervals for statistics you select.
Server-wide activity summaries	Application Activity Monitor	Displays summary level CPU usage, page I/O, and lock requests for applications.
	Engine Activity Monitor	Displays CPU usage and OS yield rates by engine.
	Network Activity Monitor	Displays statistics about packets and bytes sent and received by Adaptive Server.
	Process Activity Monitor	Displays summary level CPU usage, page I/O, and lock requests for processes.
	Stored Procedure Activity Monitor	Displays summary level CPU usage, page I/O, and lock requests for stored procedures and triggers.
	Transaction Activity Monitor	Displays metrics on updates, deletes, inserts, and selects. Also shows page split metrics.
Server-wide detail: cache management	Cache Monitor	Displays page I/O and efficiency for the data cache and execution statistics for the procedure cache. Shows current run size for data and procedure caches.
	Data Cache Monitor	Displays activity and efficiency of named data caches.
	Memory Utilization Monitor	Displays memory partitions allocated to Adaptive Server in pie chart format.
Server-wide detail: I/O management	Device I/O Monitor	Displays physical I/O activity and hit rates by device.
	Object Page I/O Monitor	Displays data and index page I/O for database tables.
Server-wide detail: lock management	Object Lock Status Monitor	Displays locking SPIDs and login names, locked table names, storage page number, length of time held, and number of other locks being blocked.

Table 26-1: Adaptive Server monitors (continued)

Type and Level of Detail	Monitor Name	Description
Process detail	Process Current SQL Statement Monitor	Displays the currently executing SQL statement and query plan for the selected process.

Definition of Metrics

This section defines some common metrics displayed in monitors:

- #/second for Sample
- #/second for Cumulative
- Cumulative Counts
- Cumulative Data for Processes

The monitors present many metrics as a rate. Regardless of the length of a sample interval, the rate in all cases is the number of occurrences per second (#/second). This standardized rate makes data collection easier and comparisons more meaningful.

#/second for Sample

The per-second rates for a sample interval report the number of occurrences per second averaged over the most recently completed sample interval, calculated as:

$$\text{\#/sec for sample} = \frac{\text{Count for most recently completed sample interval}}{\text{Number of seconds in one sample}}$$

#/second for Cumulative

The per-second rates for a session (cumulative rates) report the number of occurrences per second averaged over the current session. A session starts when the monitor opens; a session ends when the monitor closes. The cumulative #/second rate is calculated as:

$$\text{\#/sec for session} = \frac{\text{Cumulative count for the session}}{\text{Number of seconds since monitor was opened}}$$

Rate Examples

For example, logical read rates for sample and cumulative are calculated as follows:

$$\text{Rate for sample} = \frac{\text{Count of logical page reads in most recent sample}}{\text{Number of seconds in a sample interval}}$$

$$\text{Rate for session} = \frac{\text{Count of logical page reads since monitor opened}}{\text{Number of elapsed seconds since monitor opened}}$$

Cumulative Counts

The totals metrics in the monitors are cumulative counts of activity. Each time a sample interval elapses, the monitor obtains the amount of activity that occurred during the sample interval and adds it to the previously reported total.

The total values show Adaptive Server activity accumulated from the beginning of the server connection that starts when you open that monitor.

Cumulative Data for Processes

Cumulative data that applies to processes, such as the page I/O for a process, is an accumulation of data from the time that the monitor was opened unless the process was created after the connection was established. In that case, the totals are accumulated from the time the process was created.

Resources for More Information

This section lists chapters in Sybase documentation that can help you interpret Adaptive Server monitoring statistics and tune your system appropriately.

Information about CPU Usage and Engines

For information about tuning for efficient CPU usage, configuring Adaptive Server engines on multiprocessor machines, and Adaptive Server task management, see:

- “Using CPU Resources Effectively” in *Adaptive Server Performance and Tuning Guide*
- “Monitoring Adaptive Server Performance with sp_sysmon” in *Adaptive Server Performance and Tuning Guide*
- “Setting Configuration Parameters” in *Adaptive Server System Administration Guide*
- “Managing Multiprocessor Servers” in *Adaptive Server System Administration Guide*

Information about Tuning Devices

For information about defining devices and segments, creating objects on specific segments, moving objects among segments and devices, splitting tables among devices, partitioning tables, and the performance implications of using these Adaptive Server features, see:

- “Controlling Physical Data Placement” in *Adaptive Server Performance and Tuning Guide*
- “Overview of Disk Resources” in *System Administration Guide*
- “Creating and Using Segments” in *System Administration Guide*
- “Monitoring Adaptive Server Performance with sp_sysmon” in *Sybase Adaptive Server Performance and Tuning Guide*

► **Note**

Segments are labels for one or more devices. Although you must create and maintain objects on segments (not devices), the chapter “Creating and Using Segments” in the *System Administration Guide* includes detailed procedures for moving objects to specific devices by using segments.

Information about Memory Management

For information about how Adaptive Server uses memory and how to configure Adaptive Server memory, see:

- “Configuring Memory” in *Adaptive Server System Administration Guide*
- “Setting Configuration Parameters” in *Adaptive Server System Administration Guide*
- “Memory Use and Performance” in *Adaptive Server Performance and Tuning Guide*

Information about Tuning Data Caches

For information about sizing and configuring named data caches, see:

- “Memory Use and Performance” in *Adaptive Server Performance and Tuning Guide*
- “Configuring Data Caches” in *Adaptive Server Administration Guide*

Information about Tuning Tables and Indexes

For information about designing and maintaining tables and indexes for maximum performance, see:

- “Creating Databases and Tables” in *Transact-SQL User’s Guide*
- “Creating Indexes on Tables” in *Transact-SQL User’s Guide*
- “Database Design and Denormalizing for Performance” in *Adaptive Server Performance and Tuning Guide*
- “Indexing for Performance” in *Adaptive Server Performance and Tuning Guide*
- “Data Storage” in *Adaptive Server Performance and Tuning Guide*

Information about Tuning SQL Statements

For information on analyzing query plans and tuning SQL statements, see:

- “Understanding Query Plans” in *Adaptive Server Performance and Tuning Guide*
- “The Adaptive Server Query Optimizer” in *Adaptive Server Performance and Tuning Guide*
- “Transact-SQL Performance Tips” in *Adaptive Server Performance and Tuning Guide*

Information about Processes and Task Management

For information on analyzing and troubleshooting Adaptive Server processes, see:

- “Diagnosing System Problems” in *System Administration Guide*
- Description of **time slice** parameter in “Setting Configuration Parameters” in *System Administration Guide*
- Description of task management in “Monitoring Adaptive Server Performance with sp_sysmon” in *Adaptive Server Performance and Tuning Guide*
- Description of execution class and engine affinity in “Execution Priorities and Engine Affinity” in *Adaptive Server Performance and Tuning Guide*.

Information about Tuning Transactions

For information on understanding the performance implications of various types of update, insert, and delete operations, see:

- “The Adaptive Server Query Optimizer” in *Adaptive Server Performance and Tuning Guide*
- “Understanding Query Plans” in *Adaptive Server Performance and Tuning Guide*
- “Locking on Adaptive Server” in *Adaptive Server Performance and Tuning Guide*

Information about Locks and Lock Contention

For information on lock types, lock contention, and lock promotion parameters, see:

- “Locking on Adaptive Server” in *Adaptive Server Performance and Tuning Guide*
- “Cursors and Performance” in *Adaptive Server Performance and Tuning Guide*
- “Setting Configuration Parameters” in *Adaptive Server System Administration Guide*

Information about Networks

For information on how Adaptive Server uses the network and how to tune the network for performance, see:

- “Networks and Performance” in *Adaptive Server Performance and Tuning Guide*
- “Setting Configuration Parameters” in *Adaptive Server System Administration Guide*

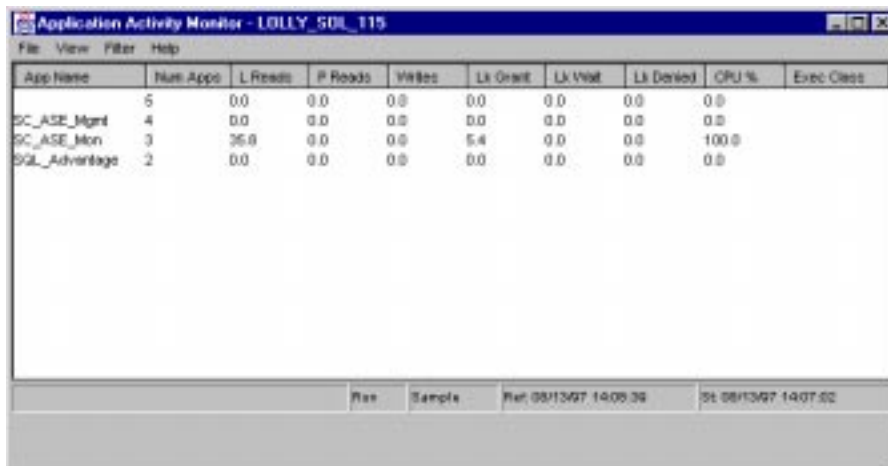
Application Activity Monitor

This monitor shows high level resource information for currently running Adaptive Server applications.

Use this monitor to:

- Identify the applications that are executed most often and consume the most resources. Improving the performance of these applications will have the greatest impact on overall system performance.
- Compare current performance against benchmarked values.
- Determine which applications tend to be I/O intensive, which are CPU intensive, and which are victims of or responsible for a large number of blocking locks.

The following figure shows the Application Activity Monitor. Data item explanations follow the figure. See *Table 26-2: Tuning using the Application Activity Monitor* for Adaptive Server tuning tips.



App Name	Num Apps	L Reads	P Reads	Writes	Lk Grant	Lk Wait	Lk Denied	CPU %	Exec Class
SC_ASE_Mgmt	5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
SC_ASE_Mon	4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
SC_ASE_Mon	3	36.0	0.0	0.0	5.4	0.0	0.0	100.0	
SQL_Advantage	2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	

Run Sample Run: 08/13/97 14:09:39 Stop: 08/13/97 14:07:52

Figure 26-1: Application Activity Monitor

The monitor shows:

- *Application Name* – the name of an Open Client Library application written using the `ct_lib` library. Adaptive Server knows application names only if the application provides it by setting the `CS_APPNAME` property in the `ct_con_props` function.

- *Number of Application Connections* (sample only) – the number of connections to Adaptive Server running this application during the sample interval. The number reflects the connection status at the time the sample interval ended.
- *Logical Reads* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page reads (whether satisfied from cache or from the database device) performed for processes executing this application.
- *Physical Reads* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page reads performed for processes executing this application that could not be satisfied from the data cache.
- *Writes* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page writes performed for processes executing this application.
- *Locks Granted* (sample or cumulative) – number of shared, exclusive, or update lock requests from processes executing this application that were either:
 - Granted immediately, or
 - Not needed because the requester already held a sufficient lock
- *Locks Granted after Wait* (sample or cumulative) – number of shared, exclusive, or update lock requests from processes executing this application that were granted after the requester waited.
- *Locks Denied* (sample or cumulative) – number of shared, exclusive, or update lock requests from processes executing this application that were not granted. Reasons for a denied lock are:
 - Didn't wait – the lock was not available immediately but the requester did not want to wait for the lock request to be queued.
 - Deadlock – the requester was selected as a deadlock victim.
 - Interrupted – the lock request was interrupted by an attention condition.
- *CPU Percent* (sample or cumulative) – the percentage of time (Adaptive Server ticks) that was consumed by processes executing this application, compared to the total time used by the Adaptive Server kernel to execute tasks during the sample interval.

- *Execution Class* (sample only) – the execution class configured for this application. Because of override features, the configured execution class does not necessarily reflect the priority and engine group at runtime. The following notations are used in the monitor:
 - **blank** means that no execution class is configured for this application. Adaptive Server uses the EC2 execution class as the default for scheduling these applications.
 - **an execution class name** is the execution class configured for the application in general, without consideration for specific logins. (That is, an execution class is configured for the application with null scope.)
 - **execution class name followed by an *** means that, in addition to the execution class configured for the application in general, additional execution classes are configured for specific logins. (That is, an execution class is configured for the application with null scope and at least one additional execution class is configured for the application with a specific scope.)
 - *** (with no execution class)** means that no execution class is configured for the application in general, but execution classes are configured for specific logins using this application. (That is, at least one execution class is configured for the application with a specific scope.)

To see all execution class bindings for applications, stored procedures, and logins, open the Execution Classes folder in Sybase Central and display the properties of individual execution classes.

More about Execution Classes

An execution class represents three execution attributes: task priority, time slice, and engine affinity. The following table shows the attribute values for the three system defined execution classes. There can also be user-defined execution classes.

Predefined Execution Class	Base Priority	Time Slice	Engine Affinity
EC1 - most preferred execution attributes	HIGH	$T = X n$	NONE
EC2 - average execution attributes	MEDIUM	$n < T < X n$	NONE
EC3 - least preferred execution attributes	LOW	$T = n$	Engine with highest ID

In the Time Slice column:

T = Time slice

X = internally set constant

n = value of the time slice Adaptive Server configuration parameter

System administrators can override the engine affinity implied by execution class.

► **Note**

In some cases, raising the execution class of a poorly performing application can make performance worse. It is important to evaluate how this application's execution attributes interact with those of other concurrently running procedures, logins, and applications.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-2: Tuning using the Application Activity Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Physical Reads</i> high	Investigate data cache issues.	<ul style="list-style-type: none"> • Object Page I/O Monitor • Data Cache Monitor
<i>Locks Denied</i> is high	Investigate lock contention issues.	<ul style="list-style-type: none"> • Object Lock Status Monitor
<i>CPU Utilization Percent</i> is low	Investigate why this application is not getting CPU time.	<ul style="list-style-type: none"> • <code>sp_sysmon</code> command (task management section) to see reasons for task context switches • Process Activity Monitor to see execution classes of concurrently running processes. • Execution Class property sheet - Bindings tab.

Cache Monitor

This monitor shows information about the procedure cache and data cache.

Use this monitor to:

- Determine efficiency of current configuration for data cache and procedure cache.
- Determine if enough memory is allocated to Adaptive Server and if the allocation of that space in the procedure cache and data cache is appropriate.

The following figure shows the Cache Monitor. Data item explanations follow the figure. See *Table 26-3: Tuning using the Cache Monitor* for Adaptive Server tuning tips. For background information, see “Information about Memory Management” on page 26-6.

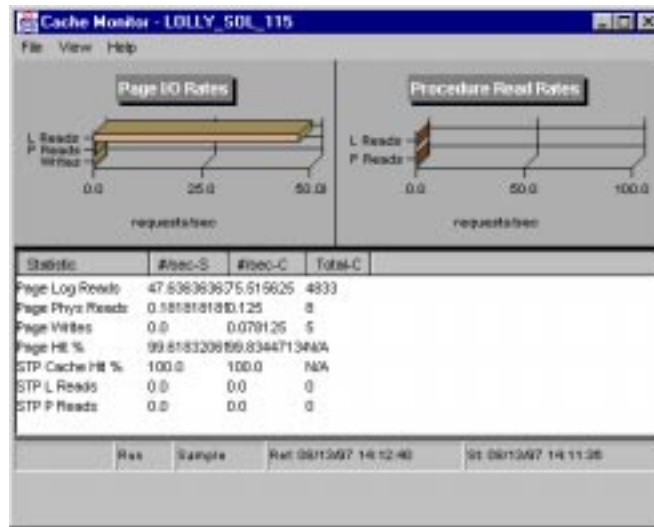


Figure 26-2: Cache Monitor

You can toggle the view to make the graphs reflect either sample or cumulative rates per second. The columns below the graphs always display both sample and cumulative data. The monitor shows:

- *Page I/O: Logical Reads* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page reads for objects associated with this data cache, whether satisfied from cache or from a database device.

The higher the value, the more activity on pages in this data cache.

- *Page I/O: Physical Reads* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page reads that could not be satisfied from cache. Physical reads require database device I/O.

Low rates are desirable because physical I/O is time-intensive. High values are normal at the beginning stages of use (for example, when Adaptive Server or an application starts) because data caches start empty and fill up during use.

- *Page I/O: Writes* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page writes.
- *Page Hit Percent* – percentage of times that a read could be satisfied from data cache rather than requiring a physical read from a database device. The displayed values are: hit rate for the sample interval and cumulative hit rate for the monitor session.

$$\text{Percent} = \frac{\text{Total logical reads} - \text{Total physical reads}}{\text{Total logical reads}}$$

High hit rates are desirable, since they indicate less time spent waiting on device I/O. Low hit rates indicate that the cache size is too small for the current load.

- *Procedure Cache Hit Percent* – percentage of times a query plan was found in the procedure cache rather than requiring a read from the *sysprocedures* table. The displayed values are: hit rate for the sample interval and cumulative hit rate for the monitor session.

High hit rates are desirable, since they indicate less time spent waiting on I/O. Low hit rates indicate that the procedure cache size is too small for the current load.

- *Procedure Logical Reads* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of requests to execute a stored procedure, whether the query plan was read from procedure cache or from the *sysprocedures* table.
- *Procedure Physical Reads* (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of requests to execute a stored procedure that required a read from the *sysprocedures* table.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-3: Tuning using the Cache Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Procedure Cache Hit Percent</i> and <i>Page Hit Percent</i> low	Allocate more memory to Adaptive Server. Test the results carefully. Allocating more memory to Adaptive Server reduces available memory elsewhere, possibly slowing down some processes.	
<i>Page Hit Percent</i> low	Increase size of the data cache or allocate more memory to Adaptive Server. Investigate configuration of named caches.	Data Cache Monitor
<i>Procedure Cache Hit Percent</i> low	Increase size of procedure cache or allocate more memory to Adaptive Server. Test the results carefully. Allocating more memory when a very large number of processes are running might have minimal gain.	

Data Cache Monitor

This monitor shows overall activity and efficiency levels of the 10 most active data caches (including the named data caches and the default cache) during the current session. Use this monitor to:

- Determine if memory is allocated efficiently among your named data caches and the default cache
- Determine if objects are bound to named data caches in an efficient scheme
- Decide where to make adjustments if inefficiencies are indicated
- Determine if buffer pools in a data cache are handling the large I/O requests of Adaptive Server

The following figure shows the Named Data Cache Monitor. Data item explanations follow the figure. See *Table 26-4: Tuning using the Data Cache Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Tuning Data Caches” on page 26-6.

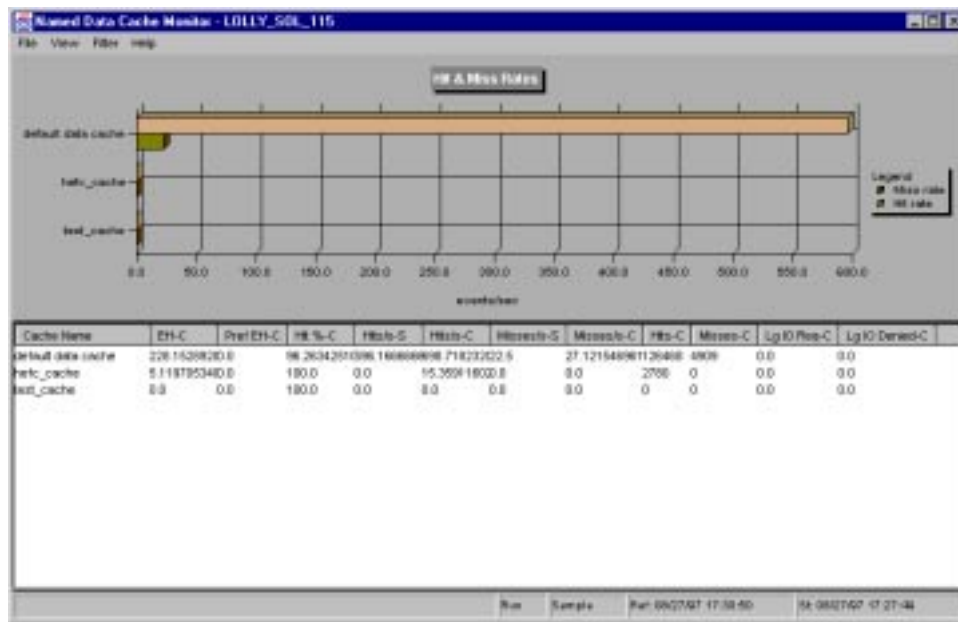


Figure 26-3: Data Cache Monitor

This monitor displays the following data:

- *Cache Name* – name assigned to a named data cache when it was configured. The default data cache holds pages for objects that are not explicitly bound to one of the named data caches.
- *Hit Percent* (cumulative) – percentage of times that a data page read could be satisfied from the cache without requiring a physical page read.

$$\text{Percent} = \frac{\text{Logical reads} - \text{Physical reads (since monitor opened)}}{\text{Logical reads (since monitor opened)}}$$

High hit rates are desirable, since they indicate less time spent waiting on device I/O. Low hit rates indicate that the named data cache size is too small for the current load.

► **Note**

Very large tables with random page access generally show a low cache hit ratio.

- *Prefetch Effectiveness* (sample or cumulative) -- a percentage comparing the count of pages being reused in large I/O buffers (the denominator) to the number of those pages that were ever referenced by Adaptive Server. When a buffer is reused, all of the pages in it are counted as reused. Buffers are reused when there are no free buffers in the pool to accept a new physical read from a database device.

$$\text{Percent} = \frac{\text{Number of reused pages referenced by Adaptive Server}}{\text{Pages per buffer} \times \text{Number of reused buffers}}$$

This metric indicates the effectiveness of large I/O grabs in the buffer pool.

High values are desirable, indicating that a high number of prefetched pages in large I/O grabs are actually being used. Low values indicate that large I/O in this data cache is not providing much benefit. Low values could also mean that a large table is fragmented enough to render large I/O ineffective.

► **Note**

Regardless of how many buffer pools are configured in a named data cache, Adaptive Server only uses two of them. It uses the 2K buffer pool and the pool configured with the largest-sized buffers. Prefetch effectiveness does not apply to the 2K buffer pool, since 2K grabs are not considered large I/O. Therefore, prefetch effectiveness applies to the largest buffer pool in the cache. For example, if a data cache has pools of buffers sized 2K, 8K, and 16K, the 8K pool is not used, and this metric reflects the effectiveness of large I/O in the 16K buffer pool.

- **Logical Reads** (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page reads for objects associated with this data cache, whether satisfied from cache or from a database device.

The higher the value, the more activity on pages in this data cache.

- **Physical Reads** (#/second for Sample, #/second for Cumulative, Cumulative Counts) – number of data page reads for objects associated with this data cache that could not be satisfied from the data cache. Physical reads require database device I/O.

Low rates are desirable because physical I/O is time-intensive. High values are normal at the beginning stages of use (for example, when Adaptive Server or an application starts) because data caches start empty and fill up during use.

- **Large I/O Requested Rate** (#/second for Sample, #/second for Cumulative) – number of times the Optimizer requested the Buffer Manager to load more than one contiguous page from a database device at a time.
- **Large I/O Denied Rate** (#/second for Sample, #/second for Cumulative) – number of times the Buffer Manager could not satisfy requests from the Optimizer to load more than one contiguous page from a database device into a data cache.

Low rates are desirable. High rates indicate that the Adaptive Server optimizer wants to perform large I/O and often cannot. See Table 26-4 for reasons why large I/O might be denied.

- **Efficiency** (session) – a measurement of space usage efficiency expressed as cache hit percent per second per megabyte.

$$\text{Relative space efficiency} = \frac{\text{Hit pct per second for this data cache}}{\text{MB in this data cache}}$$

This metric indicates relative efficiency of space allocations among data caches.

High values are desirable. The higher the value, the more efficient the space usage. Very high values mean that space is appropriately allocated to the named cache.

Very low values (low hit rates and large amounts of memory) could mean that space is being wasted. Redistributing that space (to the default cache, for example) might improve the *Hit Percent* of another named cache.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-4: Tuning using the Data Cache Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Hit Percent</i> low	Increase the size of the data cache. An optimally sized named cache retains a good percentage of often-used pages in cache, increasing hit rates. Adding space to a named data cache detracts from available procedure cache space.	<ul style="list-style-type: none"> Object Page I/O Monitor. <i>Logical Reads</i> identifies high-use objects. Caches property sheet – Configuration tab.
	Change which high-use objects are bound to the named cache. Optimal object binding reduces physical reads. Indexes and their table data can be bound to separate caches.	<ul style="list-style-type: none"> Object Page I/O Monitor. <i>Logical Reads</i> identifies high-use objects. Caches property sheet – Cache Bindings tab.
	Add buffer pools or resize existing buffer pools for large I/O.	<ul style="list-style-type: none"> Caches property sheet – Buffer Pool tab.

Table 26-4: Tuning using the Data Cache Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Efficiency</i> low relative to other named caches.	<p>Reduce the size of the cache or delete the cache and reallocate the space to other caches.</p> <p>Efficiency is low when hit rates are low and memory allocations are relatively high. This could mean that space is being wasted in this named cache. Redistributing that space (to the default cache, for example) might improve the hit rate of another cache.</p> <p>High efficiency means that space is appropriately allocated to this named cache.</p>	<ul style="list-style-type: none"> • This monitor. Compare hit rates and efficiency for other named caches.
<i>Prefetch Effectiveness</i> low	<p>Examine objects bound to the cache to determine if they are fragmented. If so, perform maintenance.</p> <p>Large I/O reads data in adjacent pages only.</p> <hr/> <p>If you need space elsewhere, delete this buffer pool and use a smaller sized pool (for example, use an 8K pool instead of a 16K pool.)</p> <hr/> <p>Check that not more than two buffer pools are allocated per cache.</p> <p>The Adaptive Server optimizer attempts to use the buffer pool with the largest sized buffers. If it cannot use that pool, it uses the 2K pool. All other pools are ignored.</p>	<ul style="list-style-type: none"> • Caches property sheet – Cache Bindings tab.
<i>Large I/O Denied Rate</i> above zero.	<p>Add buffer pools or change the size of existing buffer pools.</p> <p>Adaptive Server optimizer cannot perform large I/O if there are no buffers to handle the large reads.</p>	<ul style="list-style-type: none"> • Caches property sheet – Buffer Pool tab.

Table 26-4: Tuning using the Data Cache Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Physical Reads - Cumulative Count</i> higher than expectations	<p>Increase the size of the data cache.</p> <p>An optimally sized named cache retains a good percentage of often-used pages in cache, increasing hit rates.</p> <p>Adding space to a named data cache detracts from available procedure cache space.</p>	<ul style="list-style-type: none"> • Object Page I/O Monitor. <i>Logical Reads</i> identifies high-use objects. • Caches property sheet – Configuration tab.
	<p>Change which high-use objects are bound to the named cache.</p> <p>Optimal object binding reduces physical reads. Indexes and their table data can be bound to separate caches.</p>	<ul style="list-style-type: none"> • Object Page I/O Monitor <i>Logical Reads</i> identifies high-use objects.
	<p>Add buffer pools or resize existing buffer pools for large I/O.</p>	<ul style="list-style-type: none"> • Caches property sheet – Buffer Pool tab.
<i>Logical Reads</i> higher than expectations	<p>Reduce the number of page reads it takes to satisfy requests involving this table.</p> <ul style="list-style-type: none"> • Research long-running statements to reduce table scans • Research index design to reduce table scans 	<ul style="list-style-type: none"> • Process Current SQL Statement Monitor to tune queries. To identify processes using the table you are analyzing, set long sample intervals and toggle between the Process Activity Monitor and Object Page I/O Monitor.

Device I/O Monitor

This monitor shows buffer (not page) I/O activity on the database devices defined for Adaptive Server. A database device can be an operating system file, an entire disk, or a disk partition.

Individual database devices have different I/O rate capacities. Do not use activity level by itself as an indicator of potential or actual bottlenecks.

Device I/O counts are not the same as page I/O counts:

- **Device I/O** counts buffer reads and writes. The number of data pages included in one device read or write depends on the buffer size. For a 2K buffer pool (the default size if you are not using large I/O), one device I/O handles one 2K data page. For an 8K buffer pool, one device I/O handles four 2K data pages.
- **Data page I/O** counts data pages, regardless of the buffer size.

Use this monitor to:

- Identify bottlenecks caused by device contention.
- Identify bottlenecks caused when the maximum possible physical I/O rate for a device is approached.
- Evaluate the distribution of load among devices and decide whether to reallocate objects among the devices, or split tables.
- Determine when to add new devices.
- Monitor transaction log I/O, if the transaction log is isolated on its own device.

The following figure shows the Device I/O Monitor. Data item explanations follow the figure. See *Table 26-5: Tuning using the Device I/O Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Tuning Devices” on page 26-5.

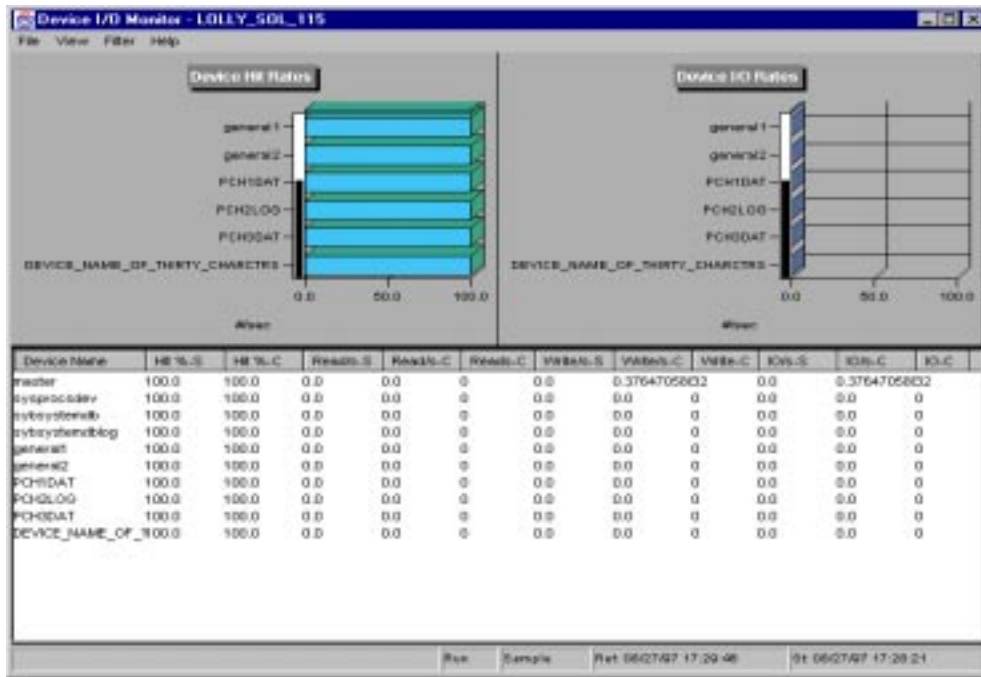


Figure 26-4: Device I/O Monitor

You can toggle the view to make the graph reflect either sample or cumulative rates per second. The columns below the graph are not affected by the toggle. This monitor displays the following data:

- *Device Name* – database device name assigned when the device was initialized for Adaptive Server use. The *master.sysdevices* table maps physical device names to Adaptive Server database device names. A database device can be an operating system file, an entire disk, or a disk partition.
- *Device Hit Percent* (sample and cumulative) – a percentage comparing the number of times access to the device was granted immediately to the number of total requests.

$$\text{Percent} = \frac{\text{I/O requests granted immediately}}{\text{Total I/O requests}}$$

Since requests can be made in one sample interval and completed in another, the rate might be skewed slightly downward. For the same reason, more requests could be granted than made during a sample interval, causing the rate to be greater than 100%.

High (but not much over 100%) rates are desirable. Low rates indicate that requests are not being granted. Values consistently over 100% indicate significant delays in granting requests.

- *Device Reads* (#/second for Sample, #/second for Sample, and Cumulative Counts) – number of buffers (not data pages) read from this database device.

The higher the value, the more read activity on the device. Compare rates for different devices to help in analyzing load balances across devices. High values that match or are approaching the physical limitation of the I/O device are undesirable, indicating that Adaptive Server might wait for I/O.

- *Device Writes* (#/second for Sample, #/second for Sample, and Cumulative Counts) – number of buffers (not data pages) written to this database device.

The higher the value, the more write activity on the device. Compare rates for different devices to help in analyzing load balances across devices. High values that match or are approaching the physical limitation of the I/O device are undesirable, indicating that Adaptive Server might wait for I/O.

- *Total Device I/O* (#/second for Sample, #/second for Sample, and Cumulative Counts) – the total of device reads and device writes.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-5: Tuning using the Device I/O Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Device Hit Percent</i> low and I/O rates are approaching the physical limits of the device	Add a new device or redistribute objects (see next row). Overloaded devices can have a high impact on performance. However, remember that different devices have different capacities. Use OS utilities to determine capacity limits for a specific device.	

Table 26-5: Tuning using the Device I/O Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Device Hit Percent</i> low or over 100%	Redistribute objects among devices. ¹ Some techniques are: <ul style="list-style-type: none"> • Place heavily used indexes on a different segment from their tables. • Separate heavily used tables by placing them on different devices. • Place <i>tempdb</i> on its own physical disk. • Place transaction logs on different devices from tables. • Split a large table over several devices. To split tables: <ul style="list-style-type: none"> - Separate text and data pages. - Partition a heap table. - Partially load a clustered index table. 	<ul style="list-style-type: none"> • <i>Totals</i> and <i>Average Rates</i> on the current monitor to determine which devices could handle additional objects. • Disk I/O Management section of <i>sp_sysmon</i> output.
Load among devices is unbalanced	Balance loads by redistributing objects. ¹ Follow tips for low <i>Hit Rate</i> .	
<i>Writes</i> very high on the transaction log's device	Configure a large I/O buffer pool for the transaction log's data cache (a 4K pool is recommended.) Use <i>sp_logiosize</i> to make sure the log I/O buffer size matches the buffer pool size.	

¹Although you must create and maintain objects on segments (not devices), the chapter "Creating and Using Segments" in the *System Administration Guide* includes detailed procedures for moving objects to specific devices by using segments. Segments are labels for one or more devices.

Engine Activity Monitor

This monitor shows details about the current CPU load caused by Adaptive Server task processing. The monitor isolates the time Adaptive Server spends processing tasks versus the total CPU time used by Adaptive Server.

When Adaptive Server has no tasks to process, it enters a loop that regularly checks for network I/O, completed disk I/Os, and tasks on the run queue. Operating system commands that check CPU activity usually include this looping activity as busy time, whereas Adaptive Server Monitor and `sp_sysmon` consider the time spent looping as idle time. Hence, this monitor reflects the actual task load on Adaptive Server engines.

Use this monitor to:

- Tune the following Adaptive Server configuration parameters:
 - `max online engines` (applies to SMP systems only)
 - `runnable process search count`
- Evaluate the task load on Adaptive Server

The following figure shows the Engine Activity Monitor. Data item explanations follow the figure. See *Table 26-6: Tuning using the Engine Activity Monitor* for Adaptive Server tuning tips. For more background information, see “Information about CPU Usage and Engines” on page 26-5.

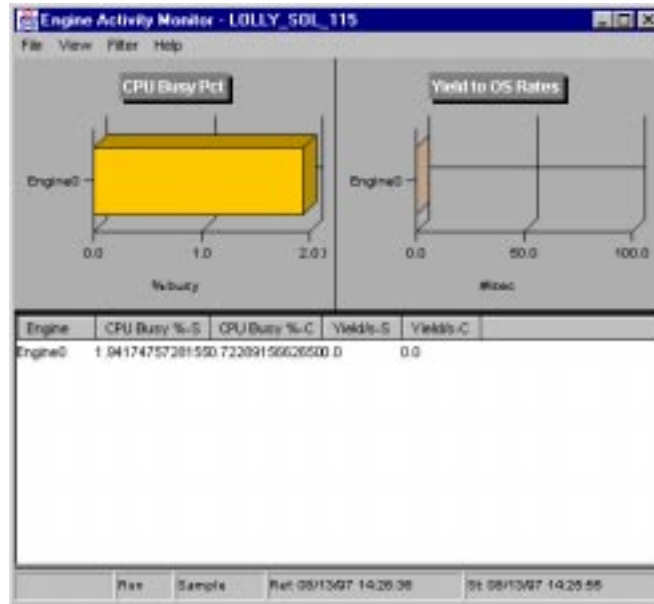


Figure 26-5: Engine Activity Monitor

You can toggle the view to make the graphs reflect either sample or cumulative rates or percentages. The columns below the graph are not affected by the toggle. This monitor displays the following data:

- *CPU Busy Percent* (sample and cumulative) – time used by the Adaptive Server kernel to execute tasks on an engine compared to the total time that the engine was available to Adaptive Server. The rest of the time that the engine is available to Adaptive Server is considered idle time.

$$\text{Percent} = \frac{\text{Time Adaptive Server spends executing tasks on engine}}{\text{Time engine is available to Adaptive Server}}$$

Values **lower than 60 percent** are desirable. Values approaching 60 percent indicate that the engine might be overloaded some of the time. If all engines are over 60%, adding another engine might improve performance.

Some engines show no activity if there are fewer processes than CPUs.

- *Yield to OS* (sample and cumulative) – number of times the engine yielded to the operating system.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-6: Tuning using the Engine Activity Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<p><i>CPU Busy Percent</i> low for all engines and <i>Yield to OS</i> high, indicating that engines are yielding voluntarily</p>	<p>Reduce the number of engines configured for Adaptive Server.</p>	<ul style="list-style-type: none"> • Network Activity Monitor to make sure your changes do not increase network latency.
<p><i>CPU Busy Percent</i> greater than 60% for all engines or <i>CPU Usage</i> medium to high and <i>Yield to OS</i> high, indicating that the engine is being starved out of CPU time</p>	<p>Add another engine by increasing the max online engines configuration parameter. First, however, check:</p> <ul style="list-style-type: none"> • Resource contention levels. Adding engines can increase contention for data caches, database devices, and locks. • Number of usable CPUs on the machine. Never configure more engines than CPUs. 	<ul style="list-style-type: none"> • Performance Summary Monitor

Memory Utilization Monitor

This monitor displays a pie chart showing Adaptive Server memory allocation. The information is based on Adaptive Server configuration parameter values, and is therefore static.

Use this monitor after you:

- Change Adaptive Server memory configuration and then reboot Adaptive Server
- Upgrade to a new Adaptive Server release

Use this monitor occasionally, to obtain an idea of how memory is allocated for Adaptive Server. During the period between your request to open the monitor and the appearance of the monitor, Adaptive Server performance may noticeably degrade.

The following figure shows the Memory Utilization Monitor. For more background information, see “Information about Memory Management” on page 26-6.




Figure 26-6: Memory Utilization Monitor

Network Activity Monitor

This monitor shows packet volume and packet sizes used for communication between Adaptive Server and its clients. The monitor also shows values of some Adaptive Server configuration parameters that affect network traffic. Use this monitor to:

- Monitor volume of Adaptive Server network activity and be proactive about preventing network overload.
- Determine effectiveness of current network configuration for Adaptive Server applications.
- Configure the default network packet size and max network packet size configuration parameters.

The following figure shows the Network Activity Monitor. Data item explanations follow the figure. See *Table 26-7: Tuning using the Network Activity Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Networks” on page 26-8.

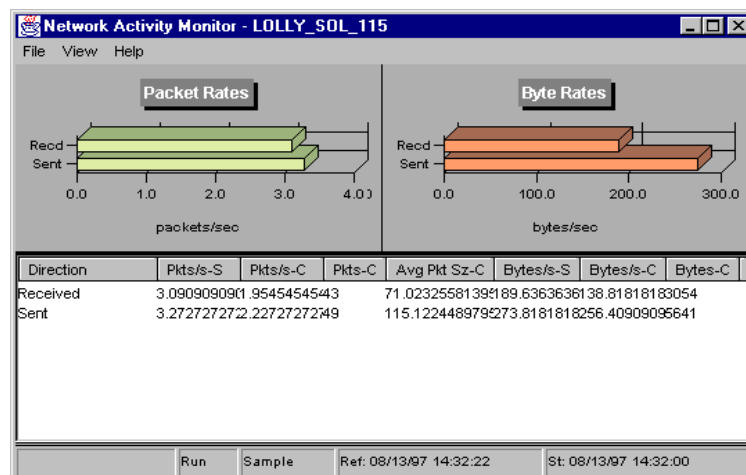


Figure 26-7: Network Activity Monitor

You can toggle the view to make the graphs reflect either sample or cumulative rates or percentages. The columns below the graph are not affected by the toggle. This monitor displays the following data:

- *Default Packet Size* – the runtime setting of the default network packet size configuration parameter. This parameter specifies the logical packet size, in bytes, used for all network traffic between Adaptive Server and its clients, unless a connection requests a different packet size when it logs in to Adaptive Server.

This parameter affects the amount of memory, in bytes, reserved at server boot time for network I/O buffers, which is:

$$3 * \text{default network packet size} * \text{number of user connections}$$

- *Maximum Packet Size* – the runtime setting of the maximum network packet size configuration parameter. This parameter specifies the maximum packet size, in bytes, that a connection can request when it logs in to Adaptive Server, overriding the default network packet size. For example, an isql command can use the -A option to request a specific packet size.

Connections that send or receive large amounts of data across the network can achieve significant performance improvement by using larger packet sizes. Generally, you want to keep the value of default network packet size small for users performing short queries, and allow users who send or receive large volumes of data to request larger packet sizes when they log on.

If you increase maximum network packet size, check that the additional network memory configuration parameter is set large enough to accommodate requests for larger packet sizes. If there is not enough additional network memory configured to handle a large packet request, the connection receives network I/O buffers equal to the default packet size even though it logged on requesting a larger size.

- *Packets Received (#/second for Sample, #/second for Cumulative, and Cumulative Counts)* – number of packets received by Adaptive Server from its clients.

Values are affected by the number of requests submitted, the length of the commands in the requests, and the number of bytes contained in each packet. On busy networks, the packet transfer rates can also depend on the amount of other traffic not related to Adaptive Server and the packet per second capacity of the network hardware.

- *Packets Sent* (#/second for Sample, #/second for Cumulative, and Cumulative Counts) – number of packets sent by Adaptive Server to its clients.

Values are affected by the length of the information that satisfies the requests and the number of bytes contained in each packet. On busy networks, traffic not related to Adaptive Server and the packet-per-second capacity of the network hardware has an impact on these rates.

- *Received Packet Size* (cumulative only) – the average number of bytes per packet received by Adaptive Server from its clients during the session. Calculated as:

$$\text{Average} = \frac{\text{Bytes received during this session}}{\text{Packets received during this session}}$$

If this value is close to the default network packet size, it means that on average, packets are full. It could mean that the packet size is not large enough to handle average requests, and many requests are being divided into multiple packets. This situation can affect Adaptive Server task-level performance. When a CPU task receives a packet in a multi-packet batch or command, it sleeps while waiting for the next packet in the batch or command.

- *Sent Packet Size* (cumulative only) – average number of bytes per packet sent by Adaptive Server during the session. Calculated as:

$$\text{Average} = \frac{\text{Bytes sent during this session}}{\text{Packets sent during this session}}$$

If this value is close to the default network packet size, it means that on average, packets are full. It could mean that the packet size is not large enough to handle average requests, and many requests are being sent in multiple packets. This situation can affect Adaptive Server task-level performance. Immediately after a CPU task sends a packet, it sleeps. A task can sleep many times if it is sending a large amount of data in small packets.

- *Bytes Received* (#/second for Sample, #/second for Cumulative, and Cumulative Counts) – number of bytes received by Adaptive Server from its clients.

These metrics indicate the length of client requests. Packet size does not affect them.

- *Bytes Sent* (#/second for Sample, #/second for Cumulative, and Cumulative Counts) – number of bytes sent by Adaptive Server to answer requests from clients.

These metrics indicate the amount of information required to satisfy requests. Packet size does not affect them.

The following table suggests performance tuning actions using the Network Activity Monitor.

Table 26-7: Tuning using the Network Activity Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<p><i>Average Received Packet Size</i> or <i>Average Sent Packet Size</i> equals (512 bytes) the <i>Default Packet Size</i> configuration parameter</p>	<p>Increase the max network packet size configuration parameter.</p>	<p>To view current configuration settings:</p> <ol style="list-style-type: none"> 1. In the Sybase Central main window, select Adaptive Server. 2. Choose File --> Configure.
	<p>Make sure that users or applications performing large data transfers log on to Adaptive Server requesting a packet size larger than the default size. Functions that can benefit from large packet sizes are:</p> <ul style="list-style-type: none"> • bulk copy • <i>readtext</i> and <i>writetext</i> commands • <i>select</i> statements resulting in large amounts of data 	
	<p>Increase the default packet size configuration parameter.</p> <p>Use multiples of 512 bytes. Consider the packet size of the underlying network.</p> <p>In general, default packet size should be large enough to handle average transfers.</p>	<ul style="list-style-type: none"> • <i>sp_sysmon</i> command. Examine task context switches due to Network Packet Received and Network Packet Sent.
	<p>Make sure the additional network memory configuration parameter is set large enough to accommodate requests for larger packet sizes. Otherwise, the connection receives network I/O buffers equal to the default packet size.</p>	

Table 26-7: Tuning using the Network Activity Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Average Received Packet Size</i> and <i>Average Sent Packet Size</i> are both at least 512 bytes less than <i>Default Packet Size</i>	Decrease default packet size configuration parameter if conserving Adaptive Server memory is important. Note: Never set default packet size less than 512 bytes.	
<i>Packets Received</i> plus <i>Packets Sent</i> is high relative to benchmarks or relative to total network traffic. (Use network monitoring tools to observe total network traffic.)	Reduce network traffic caused by Adaptive Server. Suggestions are: <ul style="list-style-type: none"> • Encourage use of stored procedures, views, and triggers. These features store code on Adaptive Server and require only a short command across the network to invoke them. • Use larger packet sizes by increasing default packet size or encouraging connections that request packet sizes larger than the default. • Set the tcp no delay configuration parameter to off (the default). The off setting allows TCP to batch small logical packets into large physical packets. • Tune the TCP packet size to better accommodate Adaptive Server average packet size. • Tune SQL statements to select (and send across the network) the specific information required. For example, an application should not select all columns in a table when it only needs information from one column. • Tune SQL statements that use cursors to request multiple rows in each fetch command if a single row does not fill up a network packet. • Run applications with large data transfers at non-peak hours or use a separate network for such operations. Run backups at non-peak hours. 	

Table 26-7: Tuning using the Network Activity Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
	<ul style="list-style-type: none"> • Use replication processing and two-phase commits only when necessary. These features contribute heavily to network traffic. • Place heavy network users on a separate network. • Configure multiple network listener ports for a single Adaptive Server. See platform specific configuration guides for more information. 	
<i>Total Network I/Os Delayed > zero</i>	<p>Add another engine (applies to SMP systems only).</p> <p>If <code>sp_sysmon</code> shows some engines having little or no activity, then adding another engine will not help.</p> <p>Consult with a network administrator to determine reason for a network bottleneck.</p>	<ul style="list-style-type: none"> • <code>sp_sysmon</code> shows network activity by engine.

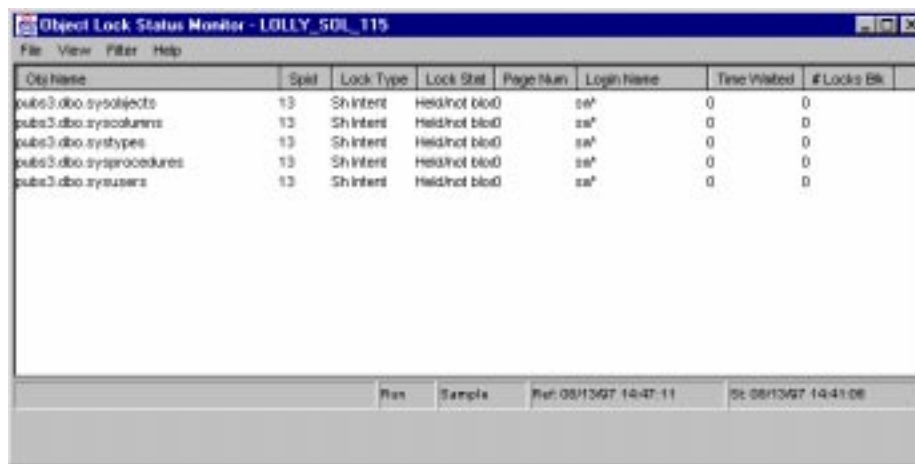
Object Lock Status Monitor

This monitor shows details about currently held or blocked locks on tables. The data is sorted using the following keys: database, object within database, page within object, and status within page. This ordering makes blocked and blocking processes show up next to each other in the list.

Use this monitor to:

- Determine which tables are currently experiencing lock contention.)
- Research the processes, login accounts, and tables associated with current locks.
- Determine which tables are experiencing contention because of a hot spot problem.
- Research time waited on locks.
- Research the types of locks that are blocking.

The following figure shows the Object Lock Status Monitor. Data item explanations follow the figure. See *Table 26-8: Tuning using the Object Lock Status Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Locks and Lock Contention” on page 26-8.



Obj Name	Spid	Lock Type	Lock Stat	Page Num	Login Name	Time Waited	#Locks Bk
pubs3.dbo.sysobjects	13	ShIntent	Held/not blo		sa*	0	0
pubs3.dbo.syscolumns	13	ShIntent	Held/not blo		sa*	0	0
pubs3.dbo.systypes	13	ShIntent	Held/not blo		sa*	0	0
pubs3.dbo.sysprocedures	13	ShIntent	Held/not blo		sa*	0	0
pubs3.dbo.sysusers	13	ShIntent	Held/not blo		sa*	0	0

Run Sample Run: 08/13/97 14:47:11 Stop: 08/13/97 14:41:08

Figure 26-8: Object Lock Status Monitor

The Object Lock Status Monitor shows sample data only. The information is a snapshot of the lock status at the end of the most recent sample interval. The monitor shows:

- *Object Name* – name of the table, in the format: *DatabaseName.Owner.TableName*. For temporary tables, the string `TempTable` is used.
- *SPID* – the identifier of a process within Adaptive Server, assigned by Adaptive Server. When a process terminates, Adaptive Server may immediately reassign the SPID to a new process. Therefore, use the SPID only to distinguish among processes that are running simultaneously.
- *Lock Type* – identifies the type of lock. **Page locks** generally are preferred over table locks, since table locks restrict the number of concurrent tasks that can access the data in the table. If a task requires an exclusive table lock, it must wait until all shared locks on the table are released. However, during a table scan, a table lock is more efficient (in terms of lock overhead) than a series of page locks.

The lock types are:

- *Shared Page Lock* – Adaptive Server applies this lock for read operations. If a shared lock has been applied to a data page, other transactions can also acquire a shared lock even when the first transaction is not finished. However, no transaction can acquire an exclusive lock on the page until all shared locks on it are released.

By default, Adaptive Server releases shared page locks after the scan is complete on the page. It does not hold them until the statement completes or until the end of the transaction. Transactions that need an exclusive page lock wait or “block” for the release of the shared page locks before continuing.

- *Shared Intent Table Lock* – Adaptive Server applies a shared intent table lock on any table that has an open cursor or a shared page lock on one of its pages. Setting a shared intent table lock prevents other transactions from subsequently acquiring a shared table lock on the table. The intent lock is held as long as the concurrent page lock is in effect.
- *Shared Table Lock* - similar to shared page lock except that it applies to the entire table.
- *Shared Address Lock* – similar to a shared page lock, except that it applies to non-leaf index pages.

- *Exclusive Page Lock* – Adaptive Server applies these locks for data modification operations. When a transaction gets an exclusive lock, other transactions cannot acquire a lock of any kind on the page until the exclusive lock is released at the end of the transaction. Those other transactions wait or “block” until the exclusive lock is released.
- *Exclusive Intent Table Lock* – Adaptive Server applies an exclusive intent table lock on any table that has an open cursor or an exclusive page lock on one of its pages. Setting an exclusive intent table lock prevents other transactions from subsequently acquiring an exclusive table lock on the table. An intent lock is held as long as the concurrent page lock is in effect.
- *Exclusive Table Lock* – similar to an exclusive page lock, except this lock applies to an entire table.
- *Exclusive Address Lock* – similar to an exclusive page lock, except that it applies to non-leaf index pages.
- *Update Page Lock* – Adaptive Server applies update locks during the initial portion of an update, delete, or fetch (for cursors declared for update) operation when the pages are being read. The update locks allow shared locks on the page, but do not allow other update or exclusive locks. This is an internal lock to help avoid deadlocks. Later, if the pages need to be changed and no other shared locks exist on the pages, Adaptive Server promotes the update locks to exclusive locks.
- *Lock Status* – one of the following:
 - Held/Blocking - the process owns the lock and at least one other process is waiting for that lock to be released.
 - Held/Not Blocking - the process owns the lock and no other processes are waiting for it.
 - Requested/Blocked - the process has requested a lock that another process currently holds.
 - Requested/Not Blocked - the process has requested a lock that is free.
- *Page Number* – the page number of the held or requested page.
- *Login Name* – the login account executing the SQL statement associated with the lock.
- *Time Waited* – amount of time a process waited for a lock. **Low** values are desirable.

- *Locks Blocked* (demand locks) – number of requested locks currently waiting for this lock to be released.

Low values are desirable. High values indicate that many processes are waiting in line for a lock on this page or table.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-8: Tuning using the Object Lock Status Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Locks Blocked</i> high or <i>Time Waited</i> high	Tune table to reduce lock contention: <ul style="list-style-type: none"> • Split a heavily used table into two tables to create the potential for more shared locking. (Separate frequently updated information from frequently accessed read-only columns.) • Decrease the number of rows stored per page by using <i>fillfactor</i> or <i>max_rows_per_page</i>. • In small tables or indexes experiencing heavy contention, use a low value for the <i>max_rows_per_page</i> option in the <i>create table</i>, <i>create index</i>, or <i>alter table</i> command. (A value of 1 in <i>max_rows_per_page</i> makes a lock on a page a lock on only one row. While this technique will reduce lock contention for that table, it causes a page split for every insert.) 	<ul style="list-style-type: none"> • Transaction Activity Monitor to monitor page splits.

Table 26-8: Tuning using the Object Lock Status Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
	<p>Examine blocking transactions to reduce the length of time they hold exclusive or update locks. Try to:</p> <ul style="list-style-type: none"> • Avoid user interactions within transactions. • Make transactions short. • Make sure update or delete transactions can use an index to locate rows. • Avoid cursors that acquire exclusive locks within transactions. (Such locks are held for the duration of the transaction.) • Make sure cursors within transactions include a close on endtran clause. (Without this clause, the cursor's locks are held beyond the transaction.) • Avoid overusing holdlock. • avoid overusing transaction isolation level 3. (Isolation levels 0 or 1 are less restrictive.) • Keep a transaction in a single batch. This reduces network latency time and thus releases locks quicker. 	<ul style="list-style-type: none"> • Note the SPID and login name associated with the locks you want to research. Use the Process Current SQL Statement Monitor to see the text and query plan of the currently running statement for a process.
The same <i>Page Number</i> appears many times	<p>Eliminate hot spots (many tasks requiring locks on the same page at the same time) in a table by:</p> <ul style="list-style-type: none"> • Partitioning or increasing the number of partitions on heap tables. (Partitioning provides more than one last page, reducing contention for exclusive locks on the last page.) • Creating a clustered index. 	

Object Page I/O Monitor

This monitor shows physical and logical page I/O statistics associated with tables and their indexes. It shows activity on **all** tables, including system tables, temporary work tables, and database tables. Use this monitor to:

- Determine which tables are the most heavily used and therefore most likely to contribute to I/O bottlenecks
- Determine which tables and indexes could benefit from being bound to a named data cache
- Determine which tables could benefit from index redesigning or index maintenance

This monitor collects a large amount of data. For most efficient use, consider the following:

- Sample data is most useful if you set long sample intervals; otherwise, the monitor refreshes before you have a chance to analyze the data
- Use the database filter to view data on tables from specific databases, rather than from all databases

The following figure shows the Object Page I/O Monitor. Data item explanations follow the figure. See *Table 26-9: Tuning using the Object Page I/O Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Tuning Tables and Indexes” on page 26-6.

Object Name	Pg L Reads/s	Pg P Reads/s	Pg W Writes/s	Idx L Reads/s	Idx P Reads/s
master.dbo.sysobjects	1.1818181818180.0	0.0	0.4545454545450.0	0.0	0.0
master.dbo.syscolumns	6.7272727272720.0	0.0	3.1818181818180.0	0.0	0.0
master.dbo.systypes	0.3636363636360.0	0.0	0.1818181818180.0	0.0	0.0
master.dbo.sysusers	4.1818181818180.0	0.0	2.0000000000000.0	0.0	0.0
master.dbo.sysdatabases	44.7272727272720.0	0.0	22.1818181818180.0	0.0	0.0
master.dbo.syslogins	2.0	0.0	0.0	1.0	0.0
master.dbo.sysmessages	0.0	0.0	0.0	0.0	0.0
tempdb.**TempOwner**.**TempObject**	1.0	0.0	0.0	0.0	0.0
tempdb.**TempOwner**.**TempObject**	1.0	0.0	0.0	0.0	0.0
tempdb.**TempOwner**.**TempObject**	1.1818181818180.0	0.0	0.0	0.0	0.0

Run Sample Ret: 08/13/97 14:49:54 01: 08/13/97 14:49:40

Figure 26-9: Object Page I/O Monitor

On the Object Page I/O Monitor, you can toggle between cumulative and sample data. The monitor shows:

- *Object Name* – name of the table, in the format: *DatabaseName.Owner.TableName*. For temporary tables, the string `TempTable` is used.
- *Data Page Logical Reads* (#/second for Sample or #/second for Cumulative) – number of data page reads for this table, whether satisfied from cache or from a database device.

The highest values indicate the tables with the most activity.

- *Data Page Physical Reads* (#/second for Sample or #/second for Cumulative) – number of data page reads for this table that could not be satisfied from the data cache.

The highest values indicate the tables causing the most device I/O activity for reads.

Low rates are desirable because physical I/O is time-intensive. High values are normal at the beginning stages of use (for example, when Adaptive Server or an application starts) because data caches start empty and fill up during use.

- *Data Page Writes* (#/second for Sample or #/second for Cumulative) – number of data pages written to the transaction log for this table.
- *Index Logical Reads* (#/second for Sample or #/second for Cumulative) – number of index page reads for this table, whether satisfied from cache or from a database device.

The highest values indicate the indexes that are being used most often for reads.

- *Index Physical Reads* (#/second for Sample or #/second for Cumulative) – number of index page reads for this table that could not be satisfied from the data cache.

The highest values indicate the indexes that are causing the most device I/O activity.

Low rates are desirable because physical I/O is time-intensive. High values are normal at the beginning stages of use (for example, when Adaptive Server or an application starts) because data caches start empty and fill up during use.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-9: Tuning using the Object Page I/O Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Data Page or Index Page Physical Reads or Writes</i> greater than expected	Bind the object to a named data cache or increase the size of the named data cache. Increase the size of the data cache that the object uses. Implement large I/O on the object by binding it to a large I/O buffer pool.	<ul style="list-style-type: none"> • Data Cache Monitor • Databases -->User Tables property sheet - Cache Binding tab.
<i>Index Page Physical Reads</i> greater than expectations	Research index usage on the table.	
<i>Data Page Logical Reads</i> much greater than expectations	Reduce the number of page reads it takes to satisfy requests involving this table. <ul style="list-style-type: none"> • Research long-running statements to reduce table scans • Research index design to reduce table scans 	<ul style="list-style-type: none"> • Process Current SQL Statement Monitor to tune queries. To identify the processes that are using the table you are analyzing, use a long sample interval and toggle between the Process Activity Monitor and Object Page I/O Monitor.

Performance Summary Monitor

This monitor summarizes the principle Adaptive Server performance indicators. Use this monitor to:

- Monitor the current overall performance of Adaptive Server
- Identify problem areas that warrant further investigation
- Begin to isolate the cause of a current performance problem

The following figure shows the Performance Summary Monitor. Data item explanations follow the figure. See *Table 26-10: Tuning using the Performance Summary Monitor* for Adaptive Server tuning tips.

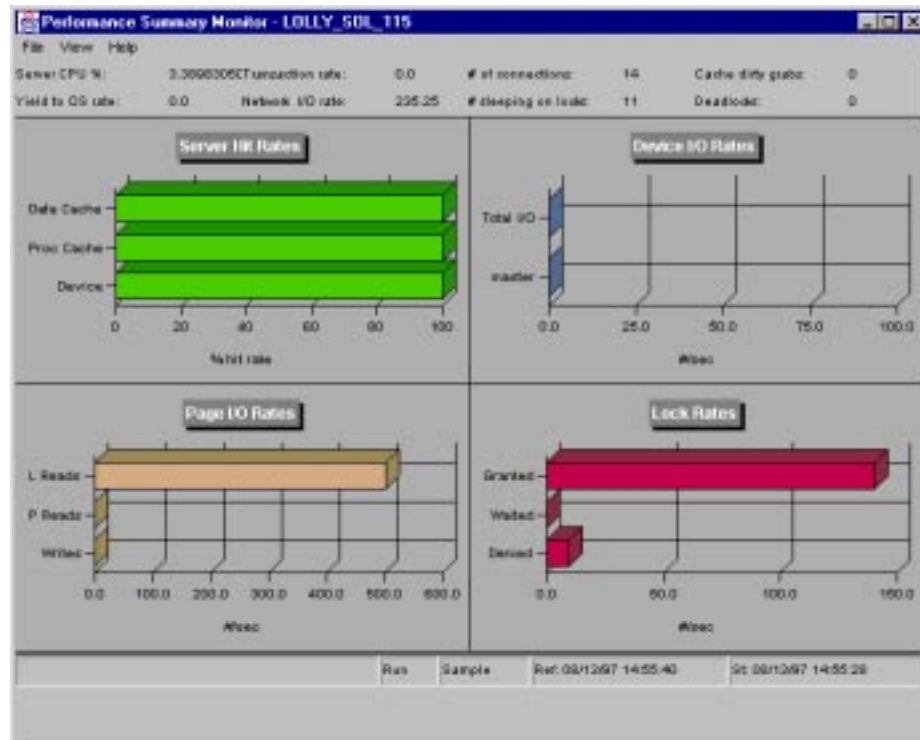


Figure 26-10: Performance Summary Monitor

You can toggle the view to show data for the last sample interval or cumulative data since the monitor was opened. (Two exceptions are the *Number of Connections* and *Number Sleeping on Locks* at the top of the monitor display. Those items always reflect connection status at the end of the most recent sample interval.) This monitor displays the following data:

- *Server CPU Percent* (sample or cumulative) – time used by the Adaptive Server kernel to execute tasks on an engine compared to the total time that the engine was available to Adaptive Server. The rest of the time that the engine is available to Adaptive Server is considered idle time. The displayed value includes all engines for the most recent sample interval or since the monitor opened.

$$\text{Percent} = \frac{\text{Time Adaptive Server spends executing tasks on all engines}}{\text{Time the engines were available}}$$

Values **lower than 60 percent** are desirable. Values approaching 60 percent indicate that the engine might be overloaded some of the time. If all engines are over 60%, adding another engine might improve performance.

Some engines show no activity if there are fewer processes than CPUs.

- *Yield to OS Rate* (#/second for Sample or #/second for Cumulative) – number of times the Adaptive Server engines yielded to the operating system.
Low values are generally desirable. High values mean that Adaptive Server is competing for CPU time with other, non-Adaptive Server processes running on the machine.
- *Transaction Rate* (#/second for Sample or #/second for Cumulative) – number of explicit, implicit, and ANSI standard transactions completed (committed). Multidatabase transactions are counted per database affected. Rollbacks are not counted. Transactions from all applications are counted, including applications written with Monitor Client Library. If a transaction starts in one sample interval and completes in the following sample interval, it is counted in the second sample interval only.
Figure 26-16 on page 26-69 illustrates implicit, explicit, and ANSI standard transactions.

The transaction count indicates total transaction throughput handled by Adaptive Server. In systems where activity level is predictable, you can compare benchmarked values to current activity to gauge the health of the system.

- *Network I/O Rate* (#/second for Sample or #/second for Cumulative) – combined number of bytes sent and received by Adaptive Server.
- *Number of Connections* (sample) – number of connections currently open on Adaptive Server at the end of the most recent sample interval. It includes all types of connections, such as client connections and remote procedure calls.

This metric provides a general understanding of the Adaptive Server environment and the work load during the current sample.

- *Number Sleeping on Locks* (sample) – number of connections waiting to obtain a lock when the sample was taken at the end of the most recent sample interval.

Low values are desirable. High values indicate high levels of object lock contention on data pages. High values could also indicate device contention.

- *Cache Dirty Grabs* (sample or cumulative) – the number of times Adaptive Server found a dirty buffer instead of a clean one when grabbing a buffer from the LRU end of a buffer pool within a data cache. In the case of a dirty buffer, Adaptive Server must wait for I/O on the dirty buffer to complete before it can complete a fetch of new pages from a database device.

A value of **zero** is desirable. Non-zero values mean that at least one buffer pool in the named cache has a wash area that is too small. Non-zero values indicate a serious performance hit.

- *Deadlock Count* (sample or cumulative) – number of locks denied because of a deadlock situation

Low values are desirable.

- *Data Cache Hit Rate* (sample or cumulative) – percentage of times that reads could be satisfied from any of the caches (named cache or the default cache).

$$\text{Percent} = \frac{\text{Total read requests} - \text{Requests requiring physical reads}}{\text{Total read requests}}$$

This percentage indicates how efficiently cache configuration is handling current activity. **High** values are desirable. Low values indicate that one or several named data caches or the default data cache has a low hit rate.

- *Procedure Cache Hit Rate* (sample or cumulative) – percentage of times that requests for procedure executions could find an idle query plan in procedure cache. If an idle query plan does not exist, Adaptive Server makes a physical read from a database device. The displayed value is a percentage for the most recent sample interval.

$$\text{Pct} = \frac{\text{Total proc execution reqs} - \text{Requests requiring physical reads}}{\text{Total procedure execution requests}}$$

This percentage indicates how efficiently the procedure cache is handling current activity. **High** values are desirable, since they indicate less time spent waiting on database device I/O. Low values mean that procedure cache might be too small.

- *Device Hit Rate* (sample or cumulative) – a percentage comparing the number of times access to devices were granted immediately to the number of total requests.

$$\text{Percent} = \frac{\text{I/O requests granted immediately}}{\text{Total I/O requests}}$$

Since requests can be made in one sample interval and completed in another, the rate might be skewed slightly downward. For the same reason, more requests could be granted than made during a sample interval, causing the rate to be greater than 100%.

High rates are desirable. Low rates indicate that requests are not being granted.

- *Total I/O* (sample or cumulative) – number of physical I/Os (that is, buffer reads and writes) involving all database devices. Depending on data cache configuration, a buffer can include a varied number of pages. The displayed values are: #/second for Sample for the sample interval or #/second for Sample since the monitor opened.
- *Most Active Device* (sample or cumulative) – number of physical I/Os (that is, buffer reads and writes) involving the most active database device. Depending on data cache configuration, a buffer can include a varied number of pages. The displayed values are: #/second for Sample for the sample interval or #/second for Sample since the monitor opened.
- *Logical Reads* (sample or cumulative) – number of data page reads, whether satisfied from cache or from a database device. The displayed values are: #/second for Sample for the sample interval or #/second for Sample since the monitor opened.
- *Physical Reads* (sample or cumulative) – number of data page reads that could not be satisfied from the data cache. A physical read requires database device I/O. The displayed values are: #/second for Sample for the sample interval or #/second for Sample since the monitor opened.

Low rates are desirable because physical I/O is time-intensive. High values are normal at the beginning stages of use (for example, when Adaptive Server or an application starts) because data caches start empty and fill up during use.

- *Writes* (sample or cumulative) – number of data page writes to the transaction log for a database. The displayed values are: #/second for Sample for the sample interval or #/second for Sample since the monitor opened.
- *Locks Granted Immediately* (sample or cumulative) – number of shared, exclusive, or update lock requests that were either:
 - Granted immediately, or
 - Not needed because the requester already held a sufficient lock
- *Locks Granted after Wait* (sample or cumulative) – number of shared, exclusive, or update lock requests that were granted after the requester waited.

Medium to low values (in proportion to Locks Granted) are desirable. High values indicates that lock contention is affecting response time, although requests are still being satisfied.

- *Locks Denied* (sample or cumulative) – number of shared, exclusive, or update lock requests that were not granted. Reasons for a denied lock are:
 - Didn't wait – the lock was not available immediately but the requester did not want to wait for the lock request to be queued.
 - Deadlock – the requester was selected as a deadlock victim.
 - Interrupted – the lock request was interrupted by an attention condition.

Low values (in proportion to Locks Granted and Locks Granted after Wait) are desirable. High values indicate that lock contention issues might be preventing tasks from completing.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-10: Tuning using the Performance Summary Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Yield to OS</i> high	Investigate whether Adaptive Server is getting enough CPU time.	• Engine Activity Monitor
<i>CPU Usage</i> greater than 60%	Add another engine by increasing the max online engines configuration parameter. First, however, check: <ul style="list-style-type: none"> • Resource contention levels. Adding engines can increase contention for data caches, database devices, and locks. • Number of usable CPUs on the machine. Never configure more engines than CPUs. 	• Check resource contention levels by examining other items displayed on this monitor.
<i>Cache Dirty Grabs</i> greater than zero	Investigate wash area size in the data caches. Perform this action as soon as possible for greatly enhanced performance.	• sp_sysmon command, Cache Management section, to see dirty grab statistics per buffer pool.
<i>Network I/O Rate</i> significantly lower than benchmarks	Investigate whether a network problem exists.	• Generic network management tools.

Table 26-10: Tuning using the Performance Summary Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Network I/O Rate</i> approaching the network's packet per second capability	Investigate reducing Adaptive Server network traffic.	• Network Activity Monitor
<i>Number Sleeping on Locks</i> high or <i>Deadlock Count</i> high	Investigate object lock contention. Investigate lock promotion configuration parameters: <ul style="list-style-type: none"> • lock promotion LWM • lock promotion HWM • lock promotion PCT You can tune these parameters server-wide or on a per-table basis.	• Object Lock Status Monitor
<i>Data Cache Hit Rate</i> low	Investigate data cache configuration and how objects are bound to data caches.	• Data Cache Monitor
<i>Procedure Cache Hit Percent</i> low	Investigate procedure cache size.	
<i>Device Hit Percent</i> low or <i>Device I/O Rate</i> approaching the speed limitations of your devices	Investigate device configuration and how objects are distributed on devices.	• Device I/O Monitor
<i>Locks Denied</i> or <i>Locks Granted after Wait</i> high	Investigate lock configuration issues. Investigate object lock contention.	• Object Lock Status Monitor

Performance Trends Monitor

This monitor plots values for selected statistics as they occur, for up to 60 sample intervals. You select the statistics you want to plot from a dialog box when you open the Performance Trends Monitor. Each statistic you choose is presented in a separate graph.

The x axis shows sample intervals, with 0 indicating the most recent sample interval, and negative values showing past intervals. The Y axis shows the values for the statistic you requested.

Use this monitor to:

- Capture benchmark data.
- Determine busy times or determine when bottlenecks tend to occur. For example, using a sample interval of 1/2 hour lets you easily plot trends for a 24-hour period.
- Scan for trouble indicators. For example, using a sample interval of one minute lets you review activity over the past hour at a glance.

The following figure shows a sample graph from the Performance Trends Monitor.

To set up this monitor:

1. Choose the Performance Trends Monitor from the Sybase Central window or from the File --> Monitors menu.
2. In the Select Data Items dialog box, choose one or more items that you want to see graphs for by clicking on them. Deselect an item by clicking on it again.

You can choose any number of items, but 4 to 6 is recommended. More than 6 graphs will not fit on most display terminals.

3. Click Okay. The Performance Trends Monitor opens.
4. To set the sample interval length, from the View menu on the Performance Trends Monitor, choose Sample Interval.

The monitor displays 60 sample intervals at a time. If you set sample interval to 10 seconds, you can see the activity trends for 10 minutes at a time. If you set sample interval to 30 minutes, you can see the activity trends for 30 hours at a time.

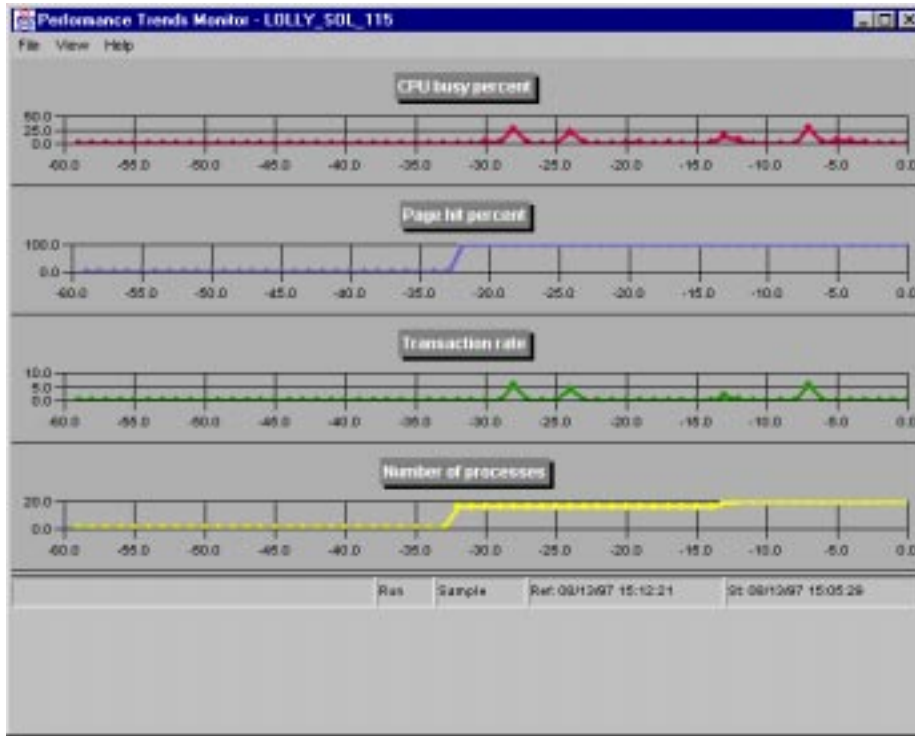


Figure 26-11: A Sample Performance Trends Monitor

Process Activity Monitor

This monitor shows high level resource information for currently running Adaptive Server processes. If a process consists of worker threads, the monitor adds all data for all worker threads and presents only the totals for the parent thread (the process).

Use this monitor to:

- Begin to troubleshoot a slow-running process.
- Monitor process activity and resource usage (CPU usage, I/O requests, and lock requests).
- Research locking problems.

The following figure shows the Process Activity Monitor. Data item explanations follow the figure. See *Table 26-11: Tuning using the Process Activity Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Processes and Task Management” on page 26-7.



The screenshot shows the 'Process Activity Monitor' window for user 'LOLLY_5DL_115'. The window contains a table with the following columns: Spid, Login Name, App Name, State, Con. Time, L. Requests, P. Requests, W. Requests, Lk. Grants, Lk. Waits, and Lk. Denied. The data is as follows:

Spid	Login Name	App Name	State	Con. Time	L. Requests	P. Requests	W. Requests	Lk. Grants	Lk. Waits	Lk. Denied
1	sa	SC_ASE_Mon	Rec Sleep	14	0.0	0.0	0.0	0.0	0.0	0.0
2	*		Sleeping	14	0.0	0.0	0.0	0.0	0.0	0.0
3	*		Sleeping	14	0.0	0.0	0.0	0.0	0.0	0.0
4	*		Sleeping	14	0.0	0.0	0.0	0.0	0.0	0.0
5	*		Sleeping	14	0.0	0.0	0.0	0.0	0.0	0.0
6	*		Sleeping	14	1.1429571420.0		0.3571429570.2857142860.0			0.0
7	sa	SC_ASE_Mon	Rec Sleep	14	18.642857140.0		0.0	3.0	0.0	0.0
8	sa	SC_ASE_Mgmt	Rec Sleep	14	0.0	0.0	0.0	0.0	0.0	0.0
9	sa	SC_ASE_Mgmt	Rec Sleep	14	0.0	0.0	0.0	0.0	0.0	0.0
10	sa	SybaseCentral	Rec Sleep	14	0.0	0.0	0.0	0.0	0.0	0.0
11	sa	SC_ASE_Mgmt	Remote IO	14	0.0	0.0	0.0	0.0	0.0	0.0
12	sa	SC_ASE_Mgmt	Remote IO	14	0.0	0.0	0.0	0.0	0.0	0.0
13	sa	SC_ASE_Mgmt	Rec Sleep	14	0.0	0.0	0.0	0.0	0.0	0.0

At the bottom of the window, there are fields for 'Res. Sample', 'Ret: 08/13/07 15: 50:07', and 'St: 08/13/07 15: 50:53'.

Figure 26-12: Process Activity Monitor

You can toggle the view to show data for the last sample interval or cumulative data (where applicable) since the monitor was opened. Note the following about how this monitor handles the volatile status of processes:

- For processes that were created before you opened this monitor, the statistics reflect process history since you opened the monitor. (Monitor Server does not collect process data until requested to do so by a client.) An asterisk (*) after the login name indicates processes that were already executing when you opened the monitor.
- For processes that were created after you opened this monitor, the statistics reflect process history since the process was created.
- If a process ends while this monitor is open, the monitor refresh at the next sample interval does not show the logged out process.

The monitor shows:

- *SPID* – the identifier for a Adaptive Server process, assigned by Adaptive Server. When a process terminates, Adaptive Server may immediately reassign the SPID to a new process. Therefore, use the SPID only to distinguish among processes that are running simultaneously.
- *Login Name* – the login name of the user who is executing the process. An asterisk (*) indicates that the process was executing before you initiated the gathering of statistics for it.
- *Application Name* – the name of an Open Client Library application written using the `ct_lib` library. Adaptive Server knows application names only if the application provides it by setting the `CS_APPNAME` property in the `ct_con_props` function.
- *Process State (sample)* – the process state at the end of the last sample interval. States are:
 - Alarm Sleep – waiting on an alarm.
 - Background – Adaptive Server process executing.
 - Bad Status – undetermined error condition.
 - Infected – tagged by Adaptive Server as unprocessable.
 - Lock Sleep – waiting on a lock acquisition.
 - Recv Sleep – waiting on a network read.
 - Runnable – waiting to run according to priority and availability of CPU.
 - Running – executing.
 - Send Sleep – waiting on a network send.

- Sleeping – paused for any other reason not listed here, such as: waiting on device I/O (physical reads) or waiting for client activity.
- Stopped – process terminated.
- Unknown – process state undeterminable.
- Remote I/O – process is waiting on a remote (OMNI) server to complete an operation.
- Synch Sleep – waiting to synchronize with some other server process(es) that are working in parallel to execute threads in a given query.
- *Connect Time* – time elapsed, in seconds, since the process was created. If the process was active before you opened the monitor, *Connect Time* is the time that has elapsed since the monitor opened.
- *Logical Reads* (#/second for Sample or #/second for Cumulative) – number of data page reads (whether satisfied from cache or from the database device) performed for this process.
- *Physical Reads* (#/second for Sample or #/second for Cumulative) – number of data page reads performed for this process that could not be satisfied from the data cache.
- *Writes* (#/second for Sample or #/second for Cumulative) – number of data page writes performed for this process.
- *Locks Granted* (sample or cumulative) – number of shared, exclusive, or update lock requests from this process that were either:
 - Granted immediately, or
 - Not needed because the requester already held a sufficient lockIn the cumulative view, the displayed value is a count for the process since the process was created or since the monitor was opened, whichever occurred later.
- *Locks Granted after Wait* (sample or cumulative) – number of shared, exclusive, or update lock requests from this process that were granted after the requester waited.
In the cumulative view, the displayed value is a count for the process since the process was created or since the monitor was opened, whichever occurred later.

- *Locks Denied* (sample or cumulative) – number of shared, exclusive, or update lock requests from this process that were not granted. Reasons for a denied lock are:
 - Didn't wait – the lock was not available immediately but the requester did not want to wait for the lock request to be queued.
 - Deadlock – the requester was selected as a deadlock victim.
 - Interrupted – the lock request was interrupted by an attention condition.
- *CPU Percent* (sample or cumulative) – the percentage of time a single process was in the Running state over the time all processes were in the Running state.
- *Blocking SPID* (sample) – the SPID holding the lock that is blocking this process.
- *Number Blocked* (sample) – indicates whether any locks requested by this process are waiting because they conflict with locks held by other processes. The Yes or No indicator reflects the locking status at the time the sample interval ended.
- *Execution Class* (sample) – the execution class this login is currently assigned to. An execution class represents three execution attributes: task priority, time slice, and engine affinity. For more information, see “More about Execution Classes” on page 26-12.

System administrators can override the engine affinity implied by execution class. The Process Activity Monitor shows which engine this process is currently using.

► **Note**

In some cases, raising the execution class of a poorly performing process can make performance worse. It is important to evaluate how this login's execution attributes interact with those of other concurrently running procedures, logins, and applications.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-11: Tuning using the Process Activity Monitor

If This is True	Consider These Actions
<i>Process State</i> is "Runnable" for an unacceptable amount of time or <i>CPU Utilization Percent</i> is low	Investigate engine activity. Investigate execution classes of this and other concurrently running processes. Consider raising execution attributes of I/O bound applications, logins, and stored procedures and/or lowering execution attributes of cpu-bound applications, logins, and stored procedures.
<i>Process State</i> is "Sleeping" for an unacceptable amount of time	Investigate the SQL statement to determine what the process is waiting for. If the process is waiting for I/O for an unreasonable amount of time, investigate device I/O contention issues.
<i>Process State</i> is "Recv Sleep" or "Send Sleep" for an unacceptable amount of time	Investigate packet configuration and network activity.
<i>Process State</i> is "Lock Sleep" for an unacceptable amount of time	Investigate blocking locks.
<i>Page I/O: Physical Reads</i> high	Investigate data cache issues.
<i>Locks Denied</i> high or <i>Blocking</i> is Y or <i>Blocked</i> is Y	Investigate lock contention issues.

Process Current SQL Statement Monitor

This monitor shows information about the currently executing SQL statement for a process, including the currently executing query plan.

Use this monitor to:

- Tune long-running SQL statements
- Research process response-time problems

To start this monitor:

1. Start a Process Activity Monitor.
2. On the Process Activity Monitor, highlight the process whose SQL statement you want to examine.
3. On the File menu, select Monitors.
4. From the cascading menu, select Process Current SQL Statement.

► **Note**

The Process Current SQL Statement Monitor remains static, whereas the Process Activity Monitor refreshes at each sample interval. If both monitors are open when execution progresses to a new statement, the two monitors no longer display information about the same statement.

This monitor is not a trace tool. It displays the currently executing statement at the time of the request, and many short-running statements can be missed between requests. Use the Process Activity Monitor to see accumulated resource usage for all statements executed by a process. If you need to watch every SQL statement as it executes, use a debug or trace tool.

The following figure shows the Process Current SQL Statement Monitor. Data item explanations follow the figure. See *Table 26-12: Tuning using the Process Current SQL Statement Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Tuning SQL Statements” on page 26-7.

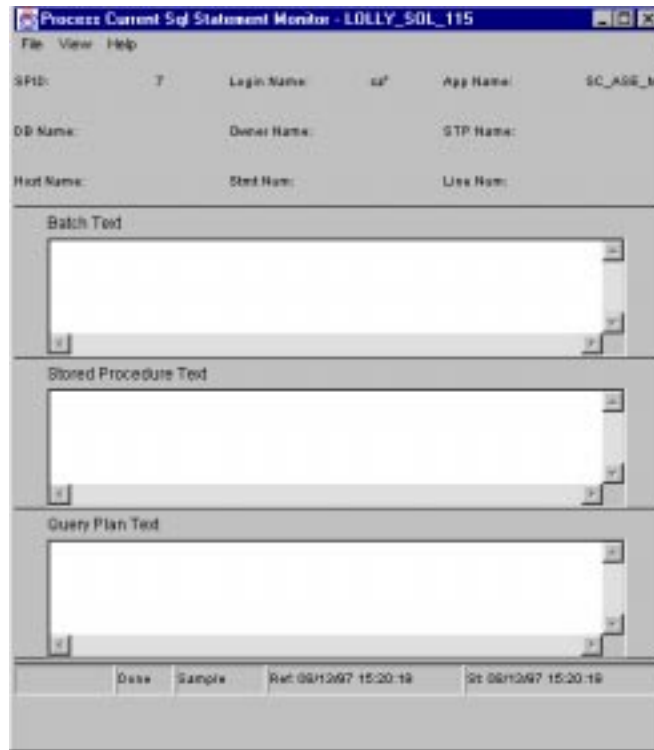


Figure 26-13: Process Current SQL Statement Monitor

The Process Current SQL Statement Monitor displays data about the SQL statement that a selected process is executing when the monitor is requested. The displayed data remains static.

The Process Current SQL Statement Monitor is intended for use only for long-running statements. Monitor Server begins collecting data about a statement only after you open the monitor. Since Monitor Server misses the beginning of statement execution, it can infer various information about the statement with much success, but it cannot guarantee that the information will always be correct. In particular:

- Batch text, stored procedure text, or the query plan could be mismatched; that is, they might not all correspond to the same statement. However, for long-running statements, this inaccuracy would be rare.

- The batch text could be the text of the wrong batch. However, for long-running statements, this inaccuracy would be rare.

This monitor shows a snapshot of data about the SQL statement that was executing when you selected the process in the Process Activity Monitor.

- *Start Time* – time the snapshot of data was taken, which is the same as the time the session started.
- *SPID* – the identifier of a process within Adaptive Server, assigned by Adaptive Server. When a process terminates, Adaptive Server may immediately reassign the SPID to a new process. Therefore, use the SPID only to distinguish among processes that are running simultaneously.
- *Login Name* – the login name of the user who is executing the SQL statement.
- *Application Name* – the name of an Open Client Library application written using the `ct_lib` library. Adaptive Server knows application names only if the application provides it by setting the `CS_APPNAME` property in the `ct_con_props` function.
- *Database Name* – when the current statement is within a stored procedure, this item identifies the database where the stored procedure resides. Otherwise, this item is blank.
- *Owner Name* – when the current statement is within a stored procedure, this item identifies the owner of the stored procedure. Otherwise, this item is blank.
- *Stored Procedure Name* – when the current statement is within a stored procedure, this item identifies the stored procedure name. Otherwise, this item is blank.
- *Host Name* – the user's machine name, if known. Adaptive Server knows the user's machine name only if the application provides it using an Open Client Library function.
- *SQL Statement Number* – the SQL statement in the batch or stored procedure that is currently executing. When both batch and stored procedure text are displayed, this item refers to the stored procedure text.
- *Line Number* – the text line within the batch or stored procedure containing the currently executing query. When both batch and stored procedure text are displayed, this item refers to the stored procedure text.

- *Batch Text* – the actual text of the currently executing batch, or the beginning of that text.

The amount of batch text available for viewing is controlled by the Adaptive Server configuration parameter `max SQL text monitored`. If this parameter is zero or too small, batch text is blank or incomplete. Changing this parameter requires rebooting Adaptive Server.

- *STP Text* – the text of the currently executing stored procedure. If a stack of nested stored procedures exists, this is the text of the most recently called stored procedure. The complete identity of the currently executing stored procedure (*database.owner.name*) is always available at the top of the monitor.

STP text could be blank if:

- no stored procedure is executing
 - the stored procedure text was deleted from the *syscomments* table or was encrypted. Third-party developers of Adaptive Server applications might delete or encrypt text to protect proprietary interests.
- *Query Plan* – the query plan of the currently executing statement.
Rarely, the query plan could be blank if the plan was removed from the Adaptive Server cache of query plans before Monitor could access it.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-12: Tuning using the Process Current SQL Statement Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Query Plan</i> is old (for example, statistics or arguments in the plan are wrong; many changes were recently made to table keys.)	Run <code>update_statistics</code> and <code>sp_recompile</code> commands. Take action as soon as practicable. Old query plans can affect performance severely.	
Statement is a heavy resource consumer	To determine the statement's area of resource overloading, look for spikes in the following areas while the statement executes: <ul style="list-style-type: none"> • CPU usage • I/O in the transaction log • Physical page I/O • Writes • Locks Granted After Waiting • Locks Denied Tune the statement and the objects it uses to reduce table scans, page locks, hot spots, and number of pages accessed. Perform maintenance on indexes.	<ul style="list-style-type: none"> • Process Activity Monitor • Stored Procedure Activity Monitor • Object Lock Status Monitor to identify hotspots. (If the same page address shows up repeatedly, the page might be a hot spot.)

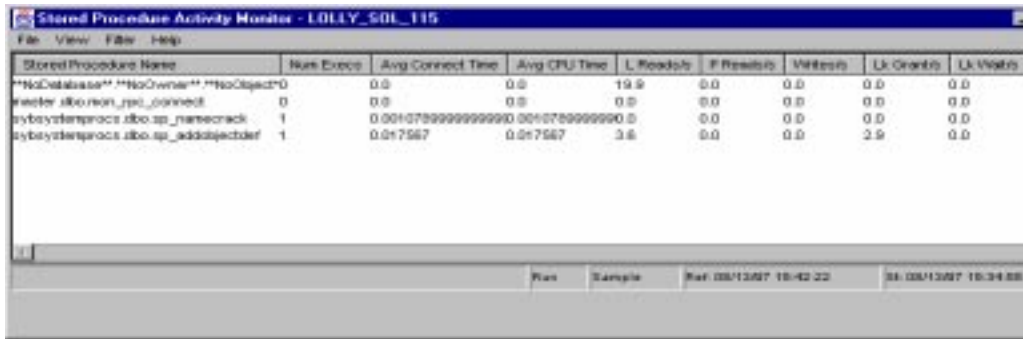
Stored Procedure Activity Monitor

This monitor shows metrics on execution, page I/O, and locks for all stored procedures and triggers that have executed during the sample or session. (A trigger is a special type of stored procedure.)

Use this monitor to:

- Identify the stored procedures that are executed most often and consume the most resources. Improving the performance of a few of those procedures can have a great impact on overall system performance.
- Compare current performance against benchmarked values.

The following figure shows the Stored Procedure Activity Monitor. Data item explanations follow the figure. See *Table 26-13: Tuning using the Stored Procedure Activity Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Tuning SQL Statements” on page 26-7.



Stored Procedure Name	Num Execs	Avg Connect Time	Avg CPU Time	L Reads/s	P Reads/s	Writes/s	Lk Grants	Lk Waits
NoDatabase **NoOwner** **NoObject**	0.0	0.0	0.0	19.9	0.0	0.0	0.0	0.0
master.dbo.ron_psi_connected	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
sybsystemprocs.dbo.sp_namecheck	1	0.0010789999999999990	0.0107899999999999990	0.0	0.0	0.0	0.0	0.0
sybsystemprocs.dbo.sp_addobjectdef	1	0.017567	0.017567	3.6	0.0	0.0	3.9	0.0

Figure 26-14: Stored Procedure Activity Monitor

The monitor shows:

- *Stored Procedure Name* – name of the stored procedure, in the format *Database.Owner.StoredProcedureName*.
- *Number of Times Executed* – number of times execution of this stored procedure completed during the sample interval (sample) or since the monitor opened (cumulative). Executions that started before the monitor opened are not counted.

- *Average Execution Time* – elapsed time, in seconds, since the stored procedure started executing. The displayed value is an average per execution of the stored procedure during the sample interval (sample) or since the monitor opened (cumulative). Executions that started before this monitor was opened are not reflected.
- *Average CPU Time per Execution* – amount of time, in seconds, that the executing process was in a running state. The displayed value is an average per execution of the stored procedure during the sample interval (sample) or since the monitor opened (cumulative). Executions that started before this monitor was opened are not reflected.
- *Logical Reads* (#/second for Sample or #/second for Cumulative) – number of data page reads (whether satisfied from cache or from the database device) made on behalf of this stored procedure. Counts are collected when the read occurs, regardless of when the stored procedure starts or ends.
- *Physical Reads* (#/second for Sample or #/second for Cumulative) – number of data page reads that could not be satisfied from the data cache made on behalf of this stored procedure. Counts are collected when the read occurs, regardless of when the stored procedure starts or ends.
- *Writes* (#/second for Sample or #/second for Cumulative) – number of data page writes made on behalf of this stored procedure. Counts are collected when the write occurs, regardless of when the stored procedure starts or ends.
- *Locks Granted* – number of shared, exclusive, or update lock requests from this stored procedure that were either:
 - Granted immediately, or
 - Not needed because the requester already held a sufficient lock
- *Locks Granted after Wait* – average number per execution of shared, exclusive, or update lock requests from this stored procedure that were granted after the requester waited.
- *Locks Denied* – average number per execution of shared, exclusive, or update lock requests from this stored procedure that were not granted. Reasons for a denied lock are:
 - Didn't wait – the lock was not available immediately but the requester did not want to wait for the lock request to be queued.
 - Deadlock – the requester was selected as a deadlock victim.

- Interrupted – the lock request was interrupted by an attention condition.
 - *Execution Class* – the execution class configured for this stored procedure. Because of override features, the configured execution class does not necessarily reflect the priority and engine group at runtime. For more information, see “More about Execution Classes” on page 26-12.
- System administrators can override the engine affinity implied by execution class.

► **Note**

In some cases, raising the execution class of a poorly performing stored procedure can make performance worse. It is important to evaluate how this procedure’s execution attributes interact with those of other concurrently running procedures, logins, and applications.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-13: Tuning using the Stored Procedure Activity Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Highest values in # of Times Executed and Average Execution Time</i>	Concentrate on improving execution time for these procedures.	<ul style="list-style-type: none"> • Process Activity Monitor to see execution classes of concurrently running processes. • Execution Class property sheet - Bindings tab.
<i>Physical Reads high</i>	Investigate data cache issues.	
<i>Locks Denied high</i>	Investigate lock contention issues.	<ul style="list-style-type: none"> • Object Lock Status Monitor

Transaction Activity Monitor

This monitor shows summary information about transaction activity handled by Adaptive Server.

Use this monitor to:

- Compare current transaction activity against benchmark values. Large deviations in transaction throughput rates can indicate developing bottlenecks.
- Determine if tuning SQL statements or indexes is making data modification operations more efficient.

The following figure shows the Transaction Activity Monitor. Data item explanations follow the figure. See *Table 26-14: Tuning using the Transaction Activity Monitor* for Adaptive Server tuning tips. For more background information, see “Information about Tuning Transactions” on page 26-7.

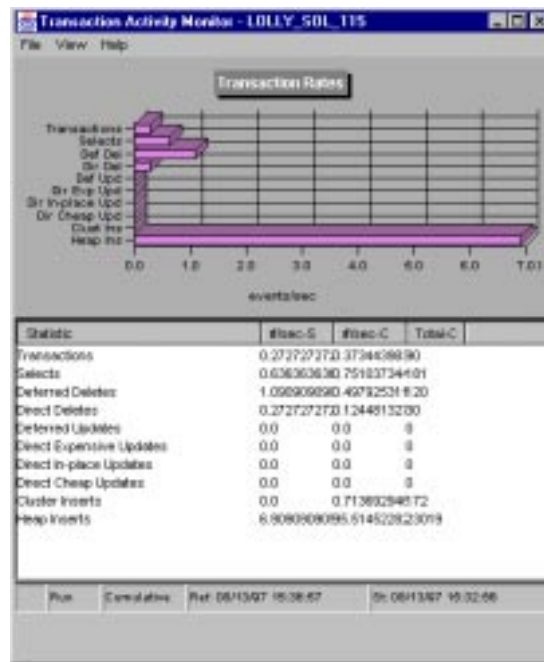


Figure 26-15: Transaction Activity Monitor

The monitor shows three transaction count metrics: #/second for Sample, #/second for Cumulative, and Cumulative Counts.

The Transaction Activity Monitor shows:

- *Transactions* – number of explicit, implicit, and ANSI standard transactions completed (committed) during the sample or session. Multidatabase transactions are counted per database affected. Rollbacks are not counted. Transactions from all applications are counted, including Adaptive Server Monitor applications using Monitor Client Library. (The Sybase Central monitors are a Monitor Client Library application.) If a transaction starts in one sample interval and completes in the following sample interval, it is counted in the second sample interval only.

Figure 26-16 illustrates implicit, explicit, and ANSI standard transactions.

The transaction count indicates total transaction throughput handled by Adaptive Server. In systems where activity level is predictable, you can compare benchmarked values to current activity to gauge the health of the system.

Implicit Transaction		
Definition	Example	Counts
A begin transaction statement is not supplied. Each data modification command, such as insert , update , or delete , is a transaction. An explicit commit transaction statement is not required.	<pre>1> insert . . . 2> go 1> insert . . . 2> go 1> insert . . . 2> go</pre>	In this example: Inserts: 3 Transactions: 3

Explicit Transaction		
Definition	Example	Counts
Data modification commands are enclosed within begin transaction and commit transaction statements. Each commit is a transaction.	<pre>1> begin trans 2> insert . . . 3> insert . . . 4> insert . . . 5> commit trans 6> go</pre>	In this example: Inserts: 3 Transactions: 1

ANSI Transaction Model		
Definition	Example	Counts
Any select or data modification command starts a transaction, but an explicit commit transaction statement must complete the transaction.	<pre>1> insert . . . 2> insert . . . 3> insert . . . 4> commit trans 5> go</pre>	In this example: Inserts: 3 Transactions: 1

Figure 26-16: Counting transactions and row inserts

- *Selects* – the number of times a select statement was executed.

- *Deferred Deletes* – number of rows deleted from database tables in deferred mode.

Low values are desirable.

- *Direct Deletes* – number of rows deleted from database tables in direct mode.
- *Deferred Updates* – number of deferred updates performed on database tables. This type of update takes place in two steps. First, the log records for deleting the existing entry and inserting the new entry are written to the log, but only the delete changes to the data pages and indexes take place. In the second step, the log pages are rescanned, and the insert operations are performed on the data pages and indexes. This type of update has the most impact on performance.

Low values are desirable. These updates increase lock contention and page I/O, and generally take more time to complete.

- *Direct Expensive Updates* – number of direct expensive updates performed on database tables. This type of update deletes a row from its original location, and inserts it in a new location. This type of update has more impact on performance than cheap updates, but not as much as deferred updates.
- *Direct In Place Updates* – number of direct in-place updates performed on database tables. This type of update does not cause data rows to move on the data page. This type of update has the least impact on performance.
- *Direct Cheap Updates* – number of direct cheap updates performed on database tables. This type of update is performed when the length of the data row changes. The changed data row remains on the same data page, but other rows on the page may move. This type of update has more impact on performance than in-place updates, but not as much as expensive updates.
- *Clustered Index Table Inserts* – number of row inserts performed by Adaptive Server on tables with clustered indexes.
- *Heap Table Inserts* – number of row inserts performed by Adaptive Server on heap tables. (Heap tables are tables that do not have clustered indexes.) The count includes inserts into unpartitioned and partitioned database tables, select into commands and inserts into work tables, and slow bulk copy inserts into heap tables. It does not include fast bulk copy inserts or transaction log writes.

The insert count indicates activity level. In general, **lower** values are preferable. High insert activity on heap tables can create lock contention and I/O bottlenecks.

► **Note**

Figure 26-16 on page 26-69 shows how inserts are counted inside transactions. The count includes inserts caused by all Adaptive Server applications, including Adaptive Server Monitor applications using Monitor Client Library.

The following table suggests performance tuning actions based on information from this monitor.

Table 26-14: Tuning using the Transaction Activity Monitor

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Page Splits</i> high	Reorganize indexes using fill factor. For tables with clustered indexes, use ascending insert mode.	
<i>Heap Table Inserts</i> high	Create a clustered index on highly used database tables. Partition tables or increase the number of existing partitions on tables with high activity.	<ul style="list-style-type: none"> Object Lock Status Monitor to identify hot spots. (If the same <i>Page Number</i> shows up repeatedly, the page might be a hot spot.) Object Page I/O Monitor shows high use tables.

Table 26-14: Tuning using the Transaction Activity Monitor (continued)

If This is True	Consider These Actions	For More Detailed Research, Use
<i>Deferred Updates</i> high	<p>Increase the chances that an appropriate index exists for a direct update by:</p> <ul style="list-style-type: none"> • Creating at least one unique index on a table. • Redesigning statements to use non-key columns in the <i>where</i> clause when updating a key. <p>Use the not null attribute for columns that do not need the null value.</p> <p>Increase the chances that enough room will exist on a page for a direct update to occur by:</p> <ul style="list-style-type: none"> • Using fixed length data types instead of variable length data types. • Using <i>max_rows_per_page</i> attribute on tables or indexes. 	<ul style="list-style-type: none"> • Process Current SQL Statement Monitor. The query plan shows whether deferred or direct updates will be used.

27

Monitoring Techniques

What's in This Chapter

This chapter contains the following topics:

- Creating Benchmarks
- Monitoring Strategies
- Tuning Strategy
- Other Useful Tools

Creating Benchmarks

This section discusses how to create a benchmark that represents expected activity at your site, with your applications, databases, and configuration.

A benchmark defines a baseline expectation of performance. If you have meaningful benchmark data (expectations), you can compare it against current performance data and make informed decisions concerning current performance.

Benchmark data can help you decide if:

- Current performance meets expectations for your site
- Resource bottlenecks exist or seem to be developing
- Recent configuration changes improved or hindered performance
- Recent introductions of new stored procedures or applications have impacted performance
- Resource usage is increasing over time, for planning purposes

To be useful in these ways, a benchmark should be:

- Representative of the workload and the resources of the production system against which you will use the benchmark to make decisions.
- Current. If you add new applications or stored procedures to your system, change configuration parameters, add database devices, or redesign data cache bindings, you need to create a new set of benchmark statistics.
- Easy to reproduce.

Designing a Benchmark

This section provides suggestions on the data to track when creating performance benchmarks for Adaptive Server.

Statistics to Collect in a Benchmark

The statistics that you include in a benchmark depend on how you intend to use the benchmark. The set of statistics that you benchmark to help predict developing resource bottlenecks might be different from the set that you use to determine unusual current activity.

The following table shows three levels of benchmarks and the monitors you can use to collect statistics for those levels.

Table 27-1: Suggested Statistics for Benchmarks

Level	Suggested Statistics
Overall performance of Adaptive Server	<p>All statistics on the Performance Summary Monitor are appropriate for benchmarks. A minimal list of benchmark statistics would include:</p> <ul style="list-style-type: none"> • <i>Server CPU Percent</i> • <i>Transaction Rate</i> (for transaction processing systems) • <i>Total I/O</i> (Device I/O) • <i>Locks Denied</i> and <i>Locks Granted After Wait</i> • <i>Data Cache Hit Rate</i> • <i>Procedure Cache Hit Rate</i> • <i>Network I/O Rate</i> <p>For decision support systems, also collect <i>Number of Selects</i> from the Transaction Activity Monitor.</p> <p>Also consider:</p> <ul style="list-style-type: none"> • Cache Monitor • Engine Activity Monitor • Transaction Activity Monitor

Table 27-1: Suggested Statistics for Benchmarks

Level	Suggested Statistics
Performance of a database	Use database filtering with the following monitors: <ul style="list-style-type: none"> • Object Page I/O Monitor • Object Lock Status Monitor • Stored Procedure Activity Monitor If data storage on devices and data cache bindings are aligned by database, you can also use statistics from: <ul style="list-style-type: none"> • Data Cache Monitor, filtering on data cache name • Device I/O Monitor, filtering on device name
Performance of an application or stored procedure	<ul style="list-style-type: none"> • Application Activity Monitor • Stored Procedure Activity Monitor • Process Activity Monitor

Statistics that Describe the Benchmark Environment

Various workload and configuration conditions can affect performance. A benchmark is not useful if you do not understand the conditions under which the statistics were collected. For example, if you make a benchmark with 25 concurrent users, it might not be useful to compare that benchmark against actual performance when there are 50 concurrent users.

Table 27-2 lists workload and configuration statistics that you should record to fully describe the benchmark environment.

Table 27-2: Describing the Benchmark Environment

Statistic	How to Collect
Workload for decision support systems (heavy querying on databases)	<ul style="list-style-type: none"> • <i>Number of Selects</i> on the Transaction Activity Monitor
Workload for transaction systems (heavy updates on databases)	<ul style="list-style-type: none"> • <i>Number of Transactions</i> on the Transaction Activity Monitor
Number of concurrent users	<ul style="list-style-type: none"> • Server level – Performance Summary Monitor displays <i>Number of Connections</i> at the end of each sample interval. Use these numbers, or collect <i>Number of Connections</i> on the Performance Trends Monitor, to estimate an average. • Application level – Application Activity Monitor displays <i>Number of Connections</i> at the end of each sample interval. • Stored procedure level – Stored Procedure Monitor displays sample and cumulative counts for <i>Number of Times Executed</i>.
Time of day and Session length	<ul style="list-style-type: none"> • <i>Session Start Time</i> on the monitor status bar. • Calculate session length using the <i>Session Start Time</i> and the <i>Refresh Time</i> on the monitor status bar.
Adaptive Server memory configuration	<ul style="list-style-type: none"> • Memory Utilization Monitor

Table 27-2: Describing the Benchmark Environment

Statistic	How to Collect
Other configuration parameters. Most Adaptive Server configuration parameters can affect performance.	<ul style="list-style-type: none"> • Save the configuration file in effect during the benchmark. The configuration file parameter specifies this file's location. • To view the current value for any configuration parameter: <ol style="list-style-type: none"> 1. From the Sybase Central main window, highlight Adaptive Server. 2. Choose File --> Configure.
Is database loading included in the benchmark?	Was the database loaded before the benchmark started or is loading part of the benchmark measurement?
Is data cache loading included in the benchmark?	Were applications and stored procedures experiencing enough activity to load the data caches before the benchmark started, or does the benchmark include the physical reads required to load data caches?

If you are creating a benchmark in a testing environment, you can control these conditions. When designing the benchmark, decide which of them should remain constant and which can vary. You might even collect several sets of benchmark statistics, varying these conditions for each one.

If you are collecting the benchmark data in a production environment, you should include these statistics in the benchmark. When you are measuring actual performance to compare against a benchmark, these statistics defining the environment are important to include.

Running the Benchmark

The best place to create a benchmark is in a testing environment. Using an isolated testing environment, you can create a situation that is easy to reproduce. You can manipulate workloads and resource allocations to represent production workloads.

If necessary, you can monitor a production environment over a period of time to create a set of statistics that you know to be representative of activity at your site.

Benchmark statistics should be cumulative rates of activity. Set the statistic type to cumulative. In general, rates are easier to compare than raw totals. Since all rates in the monitors are counts per second, you can compare rates even when sample interval or session lengths are different.

Choose any convenient sample interval, since sample interval does not affect the cumulative rates.

A suggested session length is one or several hours. A session starts when you open the monitor in Sybase Central. The start time is available in the status bar on all of the monitors. Pause the data collection when you want to record values at the end of the desired time period. A session ends when you close the monitor.

You can also use Historical Server for recording benchmarks. Historical Server views are stored in a control file, which makes any session easily repeatable.

Monitoring Strategies

This section describes some ways that you can use the Sybase Central monitors. Topics are:

- Daily Monitoring 27-7
- Researching a Stopped Process 27-8
- Capacity Planning 27-9
- Tuning Applications and Stored Procedures 27-9
- Monitoring tempdb and Transaction Logs 27-10

Daily Monitoring

To use monitors to warn you of system throughput or other developing problems, constantly monitor a few basic statistics to alert you to system throughput problems. When an obvious slowdown in throughput occurs, use other more detailed statistics to research the cause of the problem.

The Performance Summary Monitor provides a summary of overall performance. The Transaction Activity Monitor provides a summary of activity levels. If you keep these two monitors open constantly, you can check Adaptive Server performance at a glance.

Another option is to create your own performance summary by using the Performance Trends Monitor to track the statistics of your choice. The Performance Trends Monitor shows values for the last 60 sample intervals.

Statistics to Monitor

To quickly gauge system performance, check whether:

- *Server CPU Percent* is higher than benchmark values.
- *Transaction Rate* or *Number of Selects* is lower than the benchmark value when *Number of Connections* is normal.
- *Cache, Procedure, or Device Hit Percentage* is lower than benchmark values when applications have been running for a while.
- *Number Sleeping on Locks* or *Locks Denied Rate* is higher than benchmark values when *Number of Connections* is normal.
- *Network I/O Rate* is lower than the benchmark value when *Number of Connections* and *Transaction Rate* are normal.
- *Cache Dirty Grabs* is non-zero. Always investigate this.

Recognizing a Problem

Before you consider an unexpected value to be a problem, consider the following:

- Duration of occurrence of the unexpected value. For any monitored statistic, the duration is more important than any single occurrence.

- Pattern of occurrence of the unexpected value. If the unexpected value occurs only during particularly busy times or only during a database load, your resolution might be different than if it happens during more normal activity.
- Special circumstances. Consider whether a new application or stored procedure was introduced that could be dominating resources. Determine whether resource-intensive utilities such as bulk copy or rebuilding indexes are running.
- Workload and number of connections. If there is no activity on Adaptive Server, various rates can be very different from benchmark values, but this is not an indication of a performance problem.
- Configuration changes. If you have changed data cache or buffer pool configuration, memory configuration, added or subtracted engines, or redesigned databases, your benchmarks are out of date.

See Chapter 26, “Interpreting Monitor Data” for help in investigating the reasons for a sudden decrease in system-wide performance. *Table 26-10: Tuning using the Performance Summary Monitor* on page 26-50 is a good place to start.

Researching a Stopped Process

To research a stopped process, use the Process Activity Monitor. From the information on that monitor, decide whether to investigate:

- The statement being executed, using the Process Current SQL Statement Monitor.
- Locking issues, using the Object Lock Status Monitor.
- Execution class assignments, using Sybase Central.

See *Table 26-11: Tuning using the Process Activity Monitor* on page 26-58 for more information.

Capacity Planning

You can use monitors to observe trends in increased resource usage. To monitor for capacity planning, such as anticipating the need for more memory, devices, or CPU time:

1. Research and record the resource limits of your system. For example, you should be aware of the following limits:
 - Maximum device I/O rates. Remember that different devices have different capacities.
 - Maximum I/O capacities for your system related to your network
 - Available memory.
2. Create benchmarks of resource usage.
3. Rerun the same benchmarks at regular intervals, perhaps monthly, during normal activity levels. Keep the environment as close as possible to the original benchmark environment. See *Table 27-2: Describing the Benchmark Environment* for a list of factors that can affect performance measurements.
4. Record and save results in a way that allows you to compare the trends over time. Note when the normal resource usage seems to be approaching your system limits.

Tuning Applications and Stored Procedures

The Process Current SQL Statement Monitor can be useful in optimizing new applications and stored procedures at the statement level. Developers can use the monitor in a testing environment to identify long running SQL batches, identify which statements are involved, measure the resources used, and try to improve performance. For example, they can change queries, indexing schemes, or table organization, and use the Process Current SQL Statement Monitor to observe resource usage after each change.

Monitoring *tempdb* and Transaction Logs

If *tempdb* and log files are separated onto separate data caches, you can watch the hit percentages for those data caches. Using the Object Page I/O Monitor with appropriate filters, you can watch activity on those data caches. If they are separated onto their own database device, you can watch the device hit rates as well. See the *Performance and Tuning Guide* for discussions of how to reduce bottlenecks caused by access to these files.

Tuning Strategy

You can seriously degrade Adaptive Server performance by trying to tune too many performance factors at once or by changing a configuration parameter to an inappropriate value. To avoid this problem, Sybase recommends the following tuning strategy:

1. Develop a hypothesis about what could be causing the problem.
2. Change one configuration detail.
3. Monitor activity to see if conditions improve.
4. If the problem is not resolved or resource usage has not improved, change that item back and try something else.
5. When you resolve the problem, monitor activity against benchmarks to ensure that some other problem has not developed as a result of your change.

Other Useful Tools

In addition to using Adaptive Server monitors in Sybase Central, the following tools might be useful for creating benchmarks and monitoring current performance:

- **Historical Server.** See *Sybase Adaptive Server Monitor Historical Server User's Guide*.
- **The `sp_sysmon` stored procedure.** See *Sybase Adaptive Server Enterprise Performance and Tuning Guide*. The data collected by `sp_sysmon` is server-level data only.

➤ **Note**

The `sp_sysmon` stored procedure and Monitor Server use the same internal counters for gathering performance statistics on Adaptive Server. Whenever `sp_sysmon` is invoked, it resets these counters to zero, which causes erroneous output for any Monitor Server connections in progress. Do not expect to use `sp_sysmon` and Monitor Server clients (such as the monitors in Sybase Central or Historical Server) simultaneously.

A

Adaptive Server Configuration Parameters

What's in This Appendix

This appendix contains information about Adaptive Server configuration parameters. Information is in the following forms:

- An alphabetical list of Adaptive Server parameters, along with a brief description of each.
- A list of Adaptive Server parameters grouped according to the area of Adaptive Server behavior they affect.

For more detailed information about each configuration parameter, see *System Administration Guide*.

Adaptive Server Parameter List

The following is a list of Adaptive Server configuration parameters:

- **SQL Perfmon Integration** enables Adaptive Server performance monitoring with Windows NT Perfmon.
- **additional network memory** allocates additional memory for clients which request packet sizes that are larger than the default packet size for the server.
- **address lock spinlock ratio** specifies the number of rows in the address locks hash table protected by one spinlock (rows per spinlock).
- **allow backward scans** enables or disables bidirectional index scans.
- **allow nested triggers** is a toggle that controls the use of nested triggers. When the value is set to 1, data modifications made by triggers can fire other triggers.
- **allow procedure grouping** enables or disables the grouping of stored procedures. The default is 1 (procedure grouping allowed).
- **allow remote access** determines whether users from remote servers can access Adaptive Server. The default is 1, to allow Adaptive Server to communicate with Backup Server.
- **allow resource limits** allows the Resource Governor feature to be used by System Administrators. Default is 0 (off); static.
- **allow sendmsg** enables or disables the ability to send messages to internet addresses. The default is 0 (sendmsg not allowed).

- **allow sql server async i/o** is a toggle that enables Adaptive Server to run with asynchronous disk I/O.
- **allow updates to system tables** allows system tables to be updated directly. The default is 0 (off).
- **auditing** enables or disables auditing for the server. The default is 0 (auditing disabled).
- **audit queue size** determines the number of audit records that the audit queue can hold. The default is 100.
- **cis bulk insert batch size** specifies the number of rows in a bulk copy batch when performing a SELECT INTO operation.
- **cis connect timeout** specifies the number of seconds to wait for a connect request to a remote server to complete.
- **cis cursor rows** specifies the maximum number of rows that will be returned in response to a single fetch operation.
- **cis packet size** specifies the maximum size of a network packet.
- **cis rpc handling** specifies whether the Component Integration Services DLL or a site handler is used to send RPCs. The default is 0 (a site handler is used).
- **configuration file** specifies the location of the configuration file you want to use.
- **cpu accounting flush interval** specifies how many machine clock ticks to accumulate before adding cpu usage data to *syslogins* for use in chargeback accounting statistics.
- **cpu grace time** specifies the maximum amount of time (in milliseconds) a user process can run without yielding the CPU before Adaptive Server infects it.
- **curread change w/ open cursors** determines whether users can modify their curread when a cursor is open. The default is 1 (allow users to modify their curread when a cursor is open).
- **current audit table** establishes where Adaptive Server writes audit records. Values can be 1 to 8, representing *sysaudits_01* through *sysaudits_08*; 0 means to set the current audit table to the next in sequence. The default is 1.
- **deadlock checking period** specifies the minimum amount of time (in milliseconds) a process must wait for a lock before Adaptive Server initiates a deadlock check.
- **deadlock retries** specifies the number of times a transaction will retry to acquire a lock after it has become a deadlock victim.

- **default character set id** is the number of the default character set used by the server.
- **default database size** sets the default number of megabytes allocated to each new user database. The default run value is 2 (megabytes). **default fill factor percent** determines how full Adaptive Server makes each page when it is creating a new index on existing data (unless the user specifies some other value in the create index statement). The default run value is 0.
- **default language id** is the number of the language that is used to display system messages unless a user has chosen another language from those available on the server.
- **default network packet size** sets the default size of network packets for all users on Adaptive Server.
- **default sortorder id** is the number of the sort order that is the current default on Adaptive Server. **Do not change this parameter.** See *System Administration Guide* for more information about changing the sort order.
- **deadlock checking period** specifies the minimum amount of time (in milliseconds) before Adaptive Server initiates a deadlock check for a process waiting on a lock to be released.
- **disable character set conversions** enables or disables character conversion before sending to or receiving from the client. The default is 0 (Adaptive Server converts characters before sending to or received from the client).
- **disk i/o structures** specifies the initial number of disk I/O control blocks Adaptive Server allocates on start-up.
- **enable cis** specifies whether the Component Integration Services DLL is loaded when Adaptive Server is started. The default is 0 (do not load).
- **enable rep agent threads** enables or disables use of the integrated Rep Agent Thread feature. The default is 0(off).
- **esp execution priority** specifies the extended stored procedure execution priority on the XP Server. The values range from 0 to 15. The default value is 8.
- **esp execution stack size** specifies the extended stored procedure stack size on the XP Server. The values range and the default value is platform dependent.
- **esp unload dll** enables unloading of a DLL after each execution of extended stored procedure.

- **event buffers per engine** specifies the number of events per Adaptive Server engine that can be simultaneously monitored. Events are used in conjunction with Monitor Server and a client tool for observing Adaptive Server performance.
- **event log computer name** specifies the computer name for centralized error logging with Windows NT Event Viewer.
- **event logging** enables Adaptive Server Error logging with Windows NT Event Viewer.
- **executable code size + overhead** reports the size of the Adaptive Server executable.
- **freelock transfer block size** specifies the number of locks moved between the engine freelock cache and the global freelock list.
- **global async prefetch limit** sets the percentage of space allotted for all pools on Adaptive Server. The default value is 10 percent. Valid values range from 0 (no prefetch) to 100.
- **housekeeper free write percent** determines the maximum percentage by which database writes can increase as a result of free writes initiated by the housekeeper process during the server's idle cycles. Values can range from 0 through 100.

Setting this parameter to 0 disables the housekeeper process. Setting it to 100 allows the housekeeper process to work continuously during the server's idle cycles. The default value, 10, allows the housekeeper process to continue moving buffers into the buffer wash region during the server's idle cycles as long as database writes do not increase by more than 10%.

- **i/o accounting flush interval** specifies how many disk I/Os to accumulate before flushing the data to *syslogins* for use in chargeback accounting.
- **i/o polling process count** specifies the number of tasks the scheduler will run before checking for disk and/or network I/O completions.
- **identity burning set factor** determines the percentage of potential IDENTITY column values that is made available in each block. The default value, 5000, releases .05 percent of the potential IDENTITY column values for use at a time.
- **identity grab size** allows each Adaptive Server process to reserve a block of IDENTITY column values for inserts into tables that have an IDENTITY column.

- **lock shared memory** disallows swapping of Adaptive Server pages to disk, and allowing the operating system kernel to avoid the server's internal page locking code.
- **lock promotion HWM** sets the maximum number of page locks allowed before Adaptive Server escalates to a table lock. The default value is 200.
- **lock promotion LWM** sets the minimum number of page locks allowed before Adaptive Server escalates to a table lock. The default value is 200.
- **lock promotion PCT** sets the percentage of page locks allowed before Adaptive Server escalates to a table lock. The default value is 100.
- **log audit logon failure** enables Adaptive Server audit (logon failure) logging with Windows NT Event Viewer.
- **log audit logon success** enables Adaptive Server audit (logon success) logging with Windows NT Event Viewer.
- **max SQL text monitored** specifies the amount of memory allocated per user connection for saving SQL text.
- **max async i/o's per engine** specifies the maximum number of asynchronous disk I/O requests that can be outstanding for a single engine at one time.
- **max async i/o's per server** specifies the maximum number of asynchronous disk I/O requests that can be outstanding for Adaptive Server at one time.
- **max cis remote connection** specifies the maximum number of connections to remote servers that can be active at one time.
- **max cis remote servers** specifies the maximum number of remote servers that can be accessed at one time.
- **max engine freelocks** specifies the maximum number of locks available in an engine freelock cache.
- **max online engines** controls the number of engines in a symmetric multiprocessor environment.
- **max network packet size** sets the maximum network packet size that a client program can request.
- **max number of network listeners** specifies the maximum number of network listeners that can be open at one time.
- **max parallel degree** specifies the maximum number of worker processes that can be used by a given query at any one time.

- **max roles enabled per user** specifies the maximum number of roles that can be enabled by any user in a session.
- **max scan parallel degree** specifies the maximum number of worker processes that can be used for a parallel scan against a non-partitioned object.
- **memory alignment boundary** determines on which boundary buffer caches are aligned.
- **msg confidentiality reqd** requires all messages into and out of Adaptive Server, to be encrypted. If this parameter is 0 (the default), message confidentiality is not required but may be established by the client.
- **msg integrity reqd** specifies that all messages be checked for tampering. If this parameter is 0 (the default), message integrity is not required but may be established by the client.
- **msg origin checks reqd** requires that the origin of all messages be checked. If this parameter is 0 (the default), origin checks are not required but may be established by the client.
- **msg out-of-seq checks reqd** specifies that all messages must be checked to make sure the sequence has not changed. If this parameter is 0 (the default), sequence checking is not required, but may be established by the client.
- **msg replay detection reqd** requires that all messages be checked to ensure they have not been replayed or intercepted. If this parameter is 0 (the default), replay detection is not required but may be established by the client.
- **number of alarms** specifies the number of alarms allocated by Adaptive Server. Alarms are used with the Transact-SQL `waitfor` command.
- **number of aux scan descriptors** sets the maximum number of auxiliary scan descriptors that can be shared among all connections at one time. The default run value is 200.
- **number of devices** controls the number of database devices Adaptive Server can use. It does not include devices used for database dumps.
- **number of index trips** specifies the number of times an aged index page recycles itself onto the MRU chain.
- **number of languages in cache** specifies the maximum number of languages that can simultaneously be held in the language cache. The default is 3.

- **number of locks** sets the number of available locks. The default run value is 5000.
- **number of mailboxes** specifies the number of mailbox structures Adaptive Server allocates on start-up. Mailboxes are used for process-to-process communication and synchronization.
- **number of messages** specifies the number of message structures allocated by Adaptive Server at start-up time. Messages are used in conjunction with mailboxes for process-to-process communication and synchronization.
- **number of oam trips** specifies the number of times an aged Object Allocation Map (OAM) page recycles itself onto the MRU chain.
- **number of open databases** sets the maximum number of databases that can be open at one time on Adaptive Server. The default run value is 12.
- **number of open objects** sets the maximum number of database objects that can be open at one time on Adaptive Server. The default run value is 500.
- **number of open indexes** sets the maximum number of indexes that can be open at one time on Adaptive Server. The default value is 500.
- **number of preallocated extents** specifies the number of extent structures allocated in a single trip to the page manager.
- **number of remote connections** controls the limit on active connections initiated to and from Adaptive Server. The default is 20.
- **number of remote logins** controls the number of active user connections from Adaptive Server to remote servers. The default is 20.
- **number of remote sites** controls the number of simultaneous remote sites that can access Adaptive Server. The default is 10.
- **number of sort buffers** specifies the number of buffers used to hold pages read from input tables.
- **number of user connections** sets the maximum number of user connections that can be connected to Adaptive Server at the same time. The maximum value for your system is stored in the global variable @@max_connections, and varies according to platform and operating system.
- **number of worker processes** specifies the maximum number of worker processes that can be in use Server-wide at any one time.

- **o/s file descriptors** indicates how many file descriptors have been allocated to Adaptive Server. It is not user-configurable.
- **open index hash spinlock ratio** specifies the number of internal hash tables for open indexes protected by one spinlock. The default value is 100.
- **open index spinlock ratio** specifies the number of open indexes protected by one spinlock. The default value is 100.
- **open object spinlock ratio** specifies the number of open objects protected by one spinlock. The default value is 100.
- **page lock spinlock ratio** specifies the ratio of **spinlocks** protecting the internal page locks hash table.
- **page utilization percent** controls when Adaptive Server performs an OAM (Object Allocation Map) scan to find unused pages. The default run value is 95.
- **partition groups** specifies how many partition groups to allocate for the server. Partition groups are internal structures that Adaptive Server uses to control access to individual partitions of a table. Adaptive Server allocates partition groups to a table when you partition the table or when you access it for the first time after restarting the server.

A partition group is composed of 16 partition caches, each of which stores information about a single partition. All caches in a partition group are used to store information about the same partitioned table. The default value, 64, allows for a maximum of 64 open partitioned tables and 1024 (64 times 16) open partitions.

- **partition spinlock ratio** specifies the number of partition caches that each spinlock protects. A partition spinlock prevents a process from accessing a partition cache currently used by another process.

The default value of 32 (1 spinlock for every 32 partition caches) is correct for most servers. Increasing or decreasing it may have little impact on performance. The suggested number of available spinlocks is 10 percent of the total number of partitions in use at any one time.

- **perform disk i/o on engine 0** is used on multiprocessor machines to tie disk I/O to Adaptive Server engine 0.
- **permission cache entries** determines the number of cache protectors per task.

- **print deadlock information** enables printing of deadlock information to the error log.
- **print recovery information** sets a toggle that determines what information Adaptive Server displays on the console during recovery. The default run value is 0, which means that Adaptive Server displays only the database name and a message saying that recovery is in progress.
- **procedure cache percent** specifies the amount of memory allocated to the procedure cache after Adaptive Server's memory needs are met. The default run value is 20.
- **recovery interval in minutes** sets the maximum number of minutes per database that Adaptive Server should use to complete its recovery procedures in case of a system failure. The default is 5 (minutes per database).
- **remote server pre-read packets** controls the number of packets that a site handler will pre-read in connections with remote servers. The default is 3.
- **runnable process search count** specifies the number of times an engine will loop looking for a runnable task before relinquishing the CPU.
- **secure default login** establishes a login for users who are already authenticated by a security mechanism but who do not have their own login names in *master..syslogins*.
- **select on syscomments text** enables or disables granting select to public to the text column of the *syscomments* table. The default is 1 (grant select on *syscomments.text* to *public*).
- **shared memory starting address** determines the virtual address at which Adaptive Server starts its shared memory region.
- **size of auto identity column** sets the precision of IDENTITY columns automatically created with the *sp_dboption* "auto identity" option.
- **sql server clock tick length** specifies the duration of the server's clock tick, in microseconds.
- **stack guard size** specifies the size of the stack guard area.
- **stack size** sets the size of the Adaptive Server execution stack.
- **start mail session** enables start of Sybmail session with Windows NT mail subsystem with Adaptive Server startup.

- **suspend audit when device full** controls what happens when the audit table becomes full. Values are (1) suspend the audit process and all auditable user processes (default), and (0) truncate the next audit table and start using it as the current table.
- **syb_sendmsg port number** indicates the default port number for sending internet messages.
- **systemwide password expiration** specifies the number of days that passwords remain in effect after they are changed. The default is 0 (passwords do not expire).
- **table lock spinlock ratio** specifies the number of spinlocks protecting the table locks hash table.
- **tape retention in days** sets the number of days that you expect to retain each tape after it has been used for a database or transaction log dump. The default run value is 0.
- **tcp no delay** disables TCP packet batching.
- **time slice** sets the number of milliseconds that the Adaptive Server scheduler allows a user process to run. The default run value is 100 milliseconds.
- **total data cache size** represents the amount of memory that is currently available for use as a data cache. It is a calculated value that is not directly user-configurable.
- **total memory** sets the size of memory, in 2K units, that Adaptive Server allocates from the operating system.
- **unified login required** specifies that all users must be authenticated by a security mechanism before logging in to Adaptive Server. A value of 1 requires unified login, while a value of 0 causes Adaptive Server to accept traditional login names and passwords, as well as already authenticated certificates.
- **upgrade version** is changed by the upgrade program provided with new releases.
- **use security services** enables network-based security. If this parameter is 0 (the default), network-based security services such as unified login or message confidentiality are not available. This parameter does not take effect until the server is rebooted.
- **user log cache size** specifies the size (in bytes) for each user's log cache.
- **user log cache spinlock ratio** specifies the number of user log caches per user log cache spinlock.

- **xp_cmdshell context restricts the client execution privileges of the OS command specified with xp_cmdshell. Consult documentation for further details.**

Parameters by Functional Group

Configuration parameters are grouped by functional area. While each parameter has a primary group to which it belongs, many are also present in secondary groups. For example, the `number of remote connections` parameter belongs primarily to the Network Communications group but also to the Adaptive Server Administration group and the Memory Use group. The parameter lists in this section include parameters in all groups to which they belong.

The groups are:

- Backup/Recovery
- Cache Manager
- Disk I/O
- General Information
- Languages
- Lock Manager
- Memory Use
- Network Communications
- Operating System Resources
- Physical Memory
- Processors
- Adaptive Server Administration
- User Environment

Adaptive Server Administration Parameters

- `SQL Perfmon Integration`
- `allow backward scans`
- `allow nested triggers`
- `allow resource limits`
- `allow updates to system tables`
- `cpu accounting flush interval`
- `cpu grace time`
- `deadlock retries`

- default database size
- default fill factor percent
- event buffers per engine
- housekeeper free write percent
- identity burning set factor
- identity grab size
- i/o accounting flush interval
- i/o polling process count
- lock promotion HWM
- lock promotion LWM
- lock promotion PCT
- number of alarms
- number of aux scan descriptors
- number of mailboxes
- number of messages
- number of pre-allocated extents
- number of sort buffers
- partition groups
- partition spinlock ratio
- print deadlock information
- runnable process search count
- size of auto identity column
- sql server clock tick length
- time slice
- upgrade version

Backup/Recovery Parameters

- print recovery information
- recovery interval in minutes
- tape retention in days

Cache Manager Parameters

- global asynch prefetch limit
- memory alignment boundary
- number of index trips
- number of oam trips
- procedure cache percent
- total data cache size

Component Integration Services

- enable cis
- max cis remote connections
- max cis remote servers
- cis bulk insert batch size
- cis connect timeout
- cis cursor rows
- cis packet size
- cis rpc handling

Disk I/O Parameters

- allow sql server async i/o
- disk i/o structures
- number of devices
- page utilization percent

Error Log

- event log computer name
- event logging
- log audit logon failure
- log audit logon success

Extended Stored Procedures

- esp execution priority
- esp execution stacksize
- esp unload dll
- start mail session
- xp_cmdshell context

General Information Parameter

- configuration file

Languages Parameters

- default character set id
- default language id
- default sortorder id
- disable character set conversions
- number of languages in cache

Lock Manager Parameters

- address lock spinlock ratio
- deadlock checking period
- freelock transfer block size
- max engine freelocks
- number of locks
- page lock spinlock ratio
- table lock spinlock ratio

Memory Use Parameter

- executable code size + overhead

Meta-data Caches

- number of open databases
- number of open indexes
- number of open objects
- open index hash spinlock ratio
- open index spinlock ratio
- open object spinlock ratio

Network Communications Parameters

- allow remote access
- allow sendmsg
- default network packet size
- max network packet size
- max number of network listeners
- number of remote connections
- number of remote logins
- number of remote sites
- remote server pre-read packets
- syb_sendmsg port number
- tcp no delay

Operating System Resources Parameters

- max async i/o's per engine
- max async i/o's per server
- o/s file descriptors

Parallel Query

- max parallel degree
- max scan parallel degree
- memory per worker process
- number of worker processes

Physical Memory Parameters

- additional network memory
- lock shared memory
- max SQL text monitored
- shared memory starting address
- total memory

Processors Parameters

- max online engines
- min online engines

Rep Agent Thread Administration

- enable rep agent threads

Security Related

- allow procedure grouping
- audit queue size
- auditing
- curread change w/ open cursors
- current audit table
- evaluated configuration
- max roles enabled per user
- msg confidentiality reqd
- msg integrity reqd
- msg origin checks reqd
- msg out-of-seq checks reqd
- msg replay detection reqd
- secure default login
- select on syscomments.text
- suspend audit when device full
- systemwide password expiration
- unified login required
- use security services

User Environment Parameters

- number of user connections
- permission cache entries
- stack guard size
- stack size
- user log cache size
- user log cache spinlock ratio

Glossary

access

In the audit system, the use of the `select`, `insert`, `update`, or `delete` command on a table or view.

account locking

An Adaptive Server facility that prevents a user from logging in to Adaptive Server, but permits the user's account to own databases and objects. The user can be denied access without disrupting the permissions the user may have granted to other users or the availability of database objects owned by the user. Also known as login locking.

Adaptive Server

The server in the Sybase Client/Server architecture, Adaptive Server 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.

Adaptive Server login

The name a user uses to log in to Adaptive Server. A login is valid if Adaptive Server has an entry for that user in the system table *syslogins*.

alias

A defined user login that can be used as an alternate name. Using an alias, multiple users can have the same privileges in a database.

allocation unit

An allocation unit is a logical unit of Adaptive Server storage equal to 256 2KB data pages (1/2 megabyte).

application

A client program that interacts with Adaptive Server. Specifically, an application name refers to the "program_name" column in a *sysprocesses* table entry. The client libraries provide a mechanism to specify the "program_name" of an application.

authorization

See **role**.

backup

A copy of a database or transaction log, used to restore data after a media failure.

Backup Server

Backup Server performs local or remote backups (dumps) and restores (loads) on selected databases and transaction logs on behalf of Adaptive Server. A Backup Server must be running on the same system as Adaptive Server.

base priority

The priority of a task or thread when it is created. At any time, the execution priority of a thread may be greater than or equal to its base priority, but the execution priority of a thread will never be lower than its base priority.

batch

One or more Transact-SQL statements submitted as a group to Adaptive Server for processing and terminated by an end-of-batch signal.

binding

An association between a default or a rule and a table column or a user datatype. When a rule or a default is bound to a table column or a specified user datatype, it affects the data that can be entered in the column. Binding can also refer to an association between a named data cache and a database, table, or index. When a database, table, or index is bound to a named data cache, reads from the database, table, or index go into the named cache.

bulk copy

The utility for copying data in and out of databases, called `bcp`.

cache

A portion of Adaptive Server memory used to store the most recently used database information. Data cache holds data and transaction log information, and procedure cache holds query execution plans for stored procedures. Users can partition the default data cache into named caches.

cascading menu

A submenu that appears to the right of a selected pull-down menu item. An arrow (→) next to an item on a pull-down menu indicates the existence of a cascading menu.

character set

A set of specific characters with an encoding scheme that uniquely defines each character. ASCII is a common character set.

check constraint

A method of restricting the data inserted into a table column. When an insert on the column occurs, the value of the data being inserted must pass the criteria specified in the check constraint before Adaptive Server updates the table.

checkpoint

The point at which all data pages that have changed are guaranteed to have been written to the database device.

clear

To deselect an option on a dialog box.

clustered index

An index in which the physical order and the logical (indexed) order are the same. The leaf level of a clustered index represents the data pages themselves. You can have only one clustered index per table.

command

A statement that instructs the computer to perform an operation. The command begins with a word, often a verb, that names the operation. In addition, the command can include one or more keywords with or without variable values that tailor the command.

command dialog box

A dialog box that opens when you execute a command that requires information on how it should execute, such as a dialog box for binding a rule.

command permissions

Permissions to create a specific object, such as `create table` or `create procedure`. Command permissions are granted and revoked by a Database Owner.

Component Integration Services

Component Integration Services allows Adaptive Server to present a uniform view of enterprise data to client applications and provides location transparency to enterprise-wide data sources. It allows users to access both Sybase and non-Sybase databases on different servers. These external data sources include tables, views, files, and RPCs (remote procedure calls) in database systems such as Adaptive Server, Oracle, and DB2.

constraints

Mechanisms for restricting the data that can be inserted in a table. See **check constraint**, **default**, **referential constraint**.

data definition

The process of setting up databases and creating database objects such as tables, indexes, rules, defaults, procedures, triggers, and views.

data definition language (DDL)

Transact-SQL scripts that capture the state of a database. These scripts can recreate the database in the same location or on a different server.

database

A set of related tables of data and other database objects that are organized and presented to serve a specific purpose.

database device

A device dedicated to the storage of the objects that make up databases. It can be any piece of a disk or a file in the file system that is used to store databases and database objects.

database object

A database object is one of the components of a database: user, group, login, table, view, index, stored procedure, trigger, column, default, or rule.

database owner

The owner of a database. The System Administrator creates a database and is the original database owner. The System Administrator can designate another user as the owner of a database. A database owner controls all the database objects in that database and can grant object and command permissions to other users. The user name for the database owner within his or her own database is *dbo*.

datatype

Specifies what kind of information each column holds and how the data is stored. Datatypes include *char*, *int*, *money*, and so on. Users can construct their own datatypes based on system datatypes.

dbcc commands

Instructions to the Database Consistency Checker (dbcc), which checks the logical and physical consistency of a database.

dbo

In a user's own database, Adaptive Server recognizes the user as "dbo." A database owner (dbo) logs into Adaptive Server using his or her assigned login name and password.

DDL

See **data definition language**.

default

1. The user-specified value that the server inserts for a column when no value is provided.
2. The option chosen by the system when no other option is specified.

default database

The database that users get by default when they log in to Adaptive Server.

dependency

A relationship between objects that occurs when one object refers to another, such as a stored procedure that refers to a table. If you delete the table without changing the stored procedure, errors occur.

disk allocation pieces

Disk allocation pieces are the groups of allocation units from which Adaptive Server constructs a new database file. The minimum size for a disk allocation piece is one allocation unit (256 2Kb pages).

disk initialization

The process of preparing a database device or file for Adaptive Server use. Once the device is initialized, it can be used for storing databases and database objects.

disk mirror

A duplicate Adaptive Server database device. All writes to the device being mirrored are copied to a second physical device. If one device fails, the other contains an up-to-date copy of all transactions.

Dynamic Link Library (DLL)

Dynamic linking provides a mechanism for linking applications to libraries of functions at run time. A Dynamic Link Library resides in its own file and is not linked into an application's executable file, as would be the case with static linking. The DLL is loaded by the application during the execution of the application process.

drag and drop

A quick way to perform operations involving two Adaptive Server or database objects. You can drag the icon for one object and drop it on the icon for another object to perform a variety of operations, including backing up the database, creating a user, binding rules, modifying permissions, and so on.

dump

A backup copy of a database or a transaction log, or the process of creating a backup.

dump device

A single tape, partition, or file used for a database or transaction dump. A dump can span many devices, or many dumps can be made to a single tape volume.

dump file

The name of a dump file used to identify a specific backup on the backup media. The name cannot exceed 17 characters and must conform to operating system conventions. If you do not enter a name, Backup Server generates a default name based on:

- The last 7 characters of the database name
- The two-digit year
- The three-digit day of the year (1 through 366)
- Hexadecimal-encoded time at which the dump file was created

For example, the file *cations93059E100* contains a copy of the Publications database created on the fifty-ninth day of 1993.

dynamic configuration parameter

An Adaptive Server configuration parameter that is updated immediately when you reset it. A restart of Adaptive Server is not needed.

engine

An operating system process executing an Adaptive Server image. A uniprocessor platform has a single engine. Multiprocessor systems can have more than one engine.

engine group

A set of one or more engines.

error message

A message that Adaptive Server issues, usually to the user's terminal, when it detects an error condition.

error state number

The number attached to an Adaptive Server error message that allows unique identification of the line of Adaptive Server code at which the error was raised.

execution class

A classification of applications, logins, and stored procedures based on the performance requirements of these entities. An execution class is defined by a set of execution attributes associated with the objects belonging to that class.

execution priority

The absolute priority of a task. Used interchangeably with Current Priority.

extended stored procedure

A stored procedure implemented using a DLL.

extent

When a table or index requires space, Adaptive Server allocates a block of eight 2K pages. This block is called an extent. Each 256-page allocation unit contains 32 extents.

filter

An operation that allows you to specify which objects to hide or display in a window or dialog.

folder icon

An icon that represents all objects of a specific type that are contained in a database or Adaptive Server. For example, the Users folder represents all users in a database.

for load

Specifies that a database will be created for restoration from tape.

free-space threshold

A user-specified threshold that specifies the amount of space on a segment and the action to be taken when the amount of space available on that segment is less than the specified space.

guest

If the user name “guest” exists in the *sysusers* table of a database, any user with a valid Adaptive Server login can use that database, with limited privileges.

heterogeneous server

A remote server that is not an Adaptive Server installation, for example, an Informix or DB2 database.

hysteresis

A value used to control the spacing of thresholds on a segment and to prevent the stored procedure associated with a threshold from being triggered too frequently.

icon

A graphical representation of an object or an action.

identifier

A string of characters used to identify a database object, such as a table name or column name.

identity column

A table column containing a system-generated value that uniquely identifies each row in the table. A table can have only one identity column.

index

An index is created on one or more columns of a table to speed up data retrieval by pointing to the place where the column's data is stored. See also **unique index**, **clustered index**, and **nonclustered index**.

initialization

See **disk initialization**.

interfaces file

The interfaces file is an operating system file that must be available on each machine from which connections to Adaptive Server are made. By default, this file is located in the directory that is specified in the SYBASE environment variable. The name of the file varies, but is usually *interfaces* on a UNIX system and *sql.ini* on a PC.

Each entry in the interfaces file tells the host machine how to connect to Adaptive Server. An interfaces file entry contains the name of an Adaptive Server installation and a list of services provided by the server.

isql

A command line interface to Adaptive Server.

keyword

A word or phrase that is reserved for exclusive use by Transact-SQL. Also known as **reserved word**.

last-chance threshold

A threshold created by Adaptive Server on a segment reserved for transaction log activity. Adaptive Server automatically adjusts its placement on the segment to ensure that enough space remains to dump the transaction log.

load

1. A copy of a database, used to recover from a media failure.
2. The process of copying a database backup to recover a database.

locking

The process of restricting access to resources in a multi-user environment to maintain security and prevent concurrent access problems. Adaptive Server automatically applies locks to tables or pages. The System Security Officer can lock a user's login to restrict access without the time and complications involved with deleting a user or a login.

login

The name a user uses to log in to Adaptive Server. A login is valid if Adaptive Server has an entry for that user in the system table *syslogins*.

master database

Controls the user databases and the operation of Adaptive Server as a whole. Known as *master*, it keeps track of such things as user accounts, ongoing processes, and system error messages.

message number

The number that uniquely identifies an error message.

mirror

See disk mirror.

mirror device

A duplicate Adaptive Server database device. All writes to the primary device are copied (mirrored) to a second physical device. Writes can be either serial (consecutive) or parallel (simultaneous). If one device fails, the other contains an up-to-date copy of all transactions.

model database

A template for new user databases. Each time a database is created, Adaptive Server makes a copy of *model* and extends it to the size requested, if necessary.

nonclustered index

An index that stores key values and pointers to data. The leaf level points to data pages rather than containing the data itself.

nondynamic configuration parameter

An Adaptive Server configuration parameter that requires you to restart Adaptive Server after resetting the value.

null

Having no explicitly assigned value. NULL is not equivalent to zero or to blank. A value of NULL is not considered to be greater than, less than, or equivalent to any other value, including another value of NULL.

object

See **database object**.

object icon

See **icon**.

object owner

User who owns an object either by having created the object or by being given ownership. For example, a System Administrator can designate a user as a database owner, a database owner can designate a user as a table owner or give a user permission to create a table.

object permissions

Permissions to access and modify tables, views, or procedures, such as select, insert, execute, and so on. Object permissions are granted and revoked by an object owner.

operating system

A group of programs that translates your commands to the computer, helping you perform such tasks as creating files, running programs, and printing documents.

Operator

An Adaptive Server user in charge of performing server-wide database operations such as backing up and restoring databases. The System Security Officer can assign the role of Operator to a user.

See also **role**.

parameter

1. A variable value used in conjunction with a command or a stored procedure.
2. A keyword and value that define a SQL statement.

password encryption

The process of storing a password in nondecipherable encrypted form.

permission

The authority to perform actions on certain database objects, such as *execute* a procedure or *select* from a table, or to run certain commands, such as *create table*.

See also **privilege**, **command permissions** and **object permissions**.

plug-in

A file that adds support for a database product to Sybase Central.

pop-up menu

A menu opened by clicking the second mouse button over an icon. Also called a shortcut menu.

precision

The number of significant digits in a *numeric* or *decimal* datatype that can be stored in a column. For *float* datatypes, precision is the number of significant binary digits in the mantissa.

privilege

The authority to access database objects and use database commands that is implied with a user's role.

See also **permission**.

procedure

A collection of SQL statements and optional control-of-flow statements stored under a name. Adaptive Server-supplied procedures are called system procedures.

proxy table

A table created using Component Integration Services. This table does not contain any data. It references a remote table, view, file, or RPC and allows Adaptive Server to access the data in the remote object.

public

All registered users of a database are members of the group, “public.” Users at this level of authority can create a temporary table and have access to objects whose owners have granted permissions to “public.”

pulldown menu

A menu opened from a main menu selection. The pulldown menu appears below the main menu selection.

query

1. A request for the retrieval of data.
2. Any SQL statement that manipulates data.

recovery

The process of rebuilding one or more databases from database dumps or log dumps or both.

referential constraint

A type of constraint used to ensure that data being inserted in specified columns of one table matches the data in specified columns of another table.

remote procedure call (RPC)

A procedure executed on a remote server.

role

An user’s authorization level. System-defined roles are: System Administrator, System Security Officer, and Operator. The System Security Officer can define additional roles. Roles confer permission to use commands and access and modify database objects. A user can be assigned more than one role and more than one user can have the same role.

rule

The domain of acceptable values for a table column or a user datatype.

runserver file

The file used as a reference in restarting Adaptive Server or Backup Server. By default, the runserver file is named *RUN_servername* and is created when you install Adaptive Server. Runserver files are created in the *\$SYBASE/install* directory.

sa

See **System Administrator**.

scale

The number of digits to the right of the decimal point in a *numeric* or *decimal* datatype.

schema

A collection of objects owned by a single user and created in one transaction. The schema can include tables, permissions, and all database objects.

segment

A named subset of database devices available to a particular database. It is a label that points to one or more database devices. Segments can be used to control the placement of tables and indexes on specific database devices.

server user ID

The ID number by which a user is known to Adaptive Server.

severity level number

The severity of an error condition: errors with severity levels of 19 and above are fatal errors.

shortcut menu

A menu opened by clicking the second mouse button over an icon. The shortcut menu contains the same commands as the object-specific menu.

single-user mode

Starting Adaptive Server in single-user mode allows only one System Administrator to log in, and turns on the allow updates configuration variable. Use this mode to restore the master database. This option creates a *m_RUN_servername* file and overwrites any existing *m_RUN_servername* file.

site handler

The default method of transmitting RPCs from a local server to a remote server. Site handler establishes a single physical connection between local and remote servers and multiple logical connections as required by RPCs.

sql.ini file

The interfaces file containing definitions for each Adaptive Server to which your workstation can connect. The file must be on each machine from which clients connect to Adaptive Server. Each *sql.ini* file entry tells a client or host machine how to connect to a specific Adaptive Server. The file contains the name of Adaptive Server, a list of services provided by Adaptive Server, and the port to use for connecting to Adaptive Server for each service.

sso

The system security officer role, `sso_role`, in charge of security-sensitive tasks in Adaptive Server, such as creating, dropping, and locking user accounts and changing user passwords.

status bar

A horizontal bar at the bottom of the Sybase Central window that displays information about the current action or object.

See also **toolbar**.

stored procedure

See **procedure**.

suid

See **server user ID**.

Sybase Central

A graphical, user interface that allows you to manage Sybase products using plug-ins specific to each product.

System Administrator

The user in charge of administrative tasks including managing disk storage, creating databases, creating user accounts, assigning permissions, and running diagnostic and repair functions.

See also **role**.

system databases

The three databases on a newly installed Adaptive Server: the *master* database, which controls user databases and the operation of the Adaptive Server; the temporary database (*tempdb*), used for temporary tables; and the *model* database, which is used as a template to create new user databases.

system function

A function that returns special information from the database, particularly from the system tables.

system procedures

Stored procedures that Adaptive Server supplies for use in system administration. These procedures are shortcuts for retrieving information from the system tables, as well as mechanisms for accomplishing database administration and other tasks that involve updating system tables.

System Security Officer

The user in charge of security-sensitive tasks, such as creating, dropping, and locking user accounts, and changing passwords of other users.

See also **role**.

system table

One of the data dictionary tables. The system tables keep track of information about the Adaptive Server as a whole and about each user database. The *master* database contains some system tables that are not in user databases.

tab

A panel in a dialog box. Tabs display information about an object and allow you to modify the object.

temporary database

The temporary database in Adaptive Server, *tempdb*, that provides a storage area for temporary tables and other temporary working storage needs.

threshold

A space usage value of a database or log segment. When free space in a segment falls below a threshold value, a system procedure executes.

toolbar

A horizontal in a window that contains buttons. The buttons give you alternate ways to execute menu commands or let you show and hide objects.

See also **status bar**.

Tooltip

A small window that appears when the mouse cursor rests over a command button for one second or more. The window contains brief help on the command button.

transaction

A mechanism for ensuring that a set of actions is treated as a single unit.

transaction log

A system table (*syslogs*) in which all changes to the database are recorded.

transaction log options

You have the following options for the transaction log:

- Backup, truncate, and log—back up the transaction log, remove the inactive part of the log, and create a new transaction log entry recording the backup.
- Backup and log—back up the transaction log and create a new transaction log entry recording the backup. This option retains the transaction log entries.
- Truncate and log—remove the inactive part of the log without backing it up and create a new transaction log entry recording the dump.
- Truncate only—remove the inactive part of the log without backing it up and without creating a new transaction log entry to record the dump.

trigger

A special form of stored procedure attached to a column that goes into effect when a user gives a change command such as insert, delete, or update to a specified table or column. Triggers are often used to enforce referential integrity.

unique index

An index in which no two rows can have the same key value. You cannot create a unique index on a column that includes duplicate values or more than one null value.

user ID

The ID number by which a user is known in a specific database. Distinct from **server user ID**.

user table

A database table that stores user data.

view

An alternative way of looking at the data in one or more tables. Usually created as a subset of columns from one or more tables.

wash area

A portion of each memory pool within a data cache. The wash area is used to ensure that queries that need clean pages in a data cache can find them. When the number of dirty pages (pages that have been changed in cache) fills the rest of the buffers and begins to enter the wash area, Adaptive Server writes the data in the wash area to disk. When this write completes, the wash area is marked clean and is available for queries needing clean pages.

wildcard

A special character used to represent one or more characters in a pattern-matching string. Any character or set of characters can replace a wildcard character. The -**wildcard** option is available for some commands to allow using wildcards for specifying names.

with cascade option

An option, when revoking permission, that revokes permission from a specified user or group and also from anyone to whom the specified user or group has granted permission.

with grant option

An option, when granting permission, that allows a specified user or group to grant the same permission to other users or groups.

XP Server

An Open Server application program that executes extended stored procedures when Adaptive Server requests these services via remote procedure calls (RPCs).

Index

Symbols

@@thresh_hysteresis system variable 11-7

A

Abort Transaction If Log Full database option 9-14

Access

- Adaptive Server to 12-1
- aliases 12-1
- command permissions 12-1
- dialog boxes 12-13
- granting 12-13
- object permissions 12-1
- revoking 12-13
- roles 12-2, 12-13
- shared logins 12-1
- tabs 12-13

Activating

- roles 13-9, 15-7, 15-8

Activation exclusivity 15-9

Active status 7-2

Adaptive Server

- configuring 4-10, 4-11
- connecting to 4-2, 4-4
- disconnecting from 4-6
- killing processes 4-15
- name in monitors 5-28
- pausing 4-7
- property sheet 4-9
- roles 12-3
- security 6-7, 12-1
- shutting down 4-8
- starting 4-7

Adaptive Server Monitor

- architecture 5-2
- components 5-1
- definition of 5-1

Add Alias dialog box 13-10

Add Database Device to Segment dialog box 11-5

Add Identity Column in Non-Unique Indexes database option 9-15

Adding

See also Creating

- aliases to a login 13-10
- columns 21-1
- columns to table 22-13
- disk mirroring 8-9
- engine to engine group 25-13
- remote servers 6-4
- users to group 14-14

additional network memory configuration parameter A-1

Add Memory Pool dialog box 7-8

Add New Threshold dialog box 11-9

address lock spinlock ratio configuration parameter A-1

Add Specific Mapping dialog box 6-14

Add User to Group dialog box 14-15

Administration tasks 1-2

Affinity 25-3

- changing 25-17

Aliases

- access 12-1
- adding to a login 13-10
- deleting 13-11
- sharing logins 12-2, 13-9
- user to login 14-9

All scope 25-18

Allocating

- index pages 24-9
- space for a database 9-8
- space for logs 9-2
- storage space 9-2, 9-7
- table pages 22-28

Allocation

- errors 22-28
- fixing 9-20
- maps 22-28

- pages 22-28
- table 22-28
- Allocations
 - displaying 9-7
- allow backward scans configuration
 - parameter A-1
- allow nested triggers configuration
 - parameter A-1
- Allow Nulls 20-3, 22-7
- Allow Nulls by Default database
 - option 9-14
- allow procedure grouping configuration
 - parameter A-1
- allow remote access configuration
 - parameter A-1
- allow resource limits configuration
 - parameter A-1
- Allow Select Into and Bulk Copy
 - database option 9-13
- allow sendmsg configuration
 - parameter A-1
- allow sql server async i/o configuration
 - parameter A-2
- allow updates 4-10
- allow updates to system tables configuration
 - parameter A-2
- ANYENGINE 25-10
- Append 10-6
- Application
 - scope 25-19
- Application Activity Monitor 26-9
- Application programming interface 5-2
- Applications
 - tuning 25-1
 - tuning with monitors 27-9
- Architecture
 - Adaptive Server Monitor 5-2
- Assistance 3-5
- auditing configuration parameter A-2
- audit queue size 4-10
- audit queue size configuration
 - parameter A-2
- Auto Identity database option 9-15
- Automatic connections 4-5

- Automatic role activation 13-9, 15-8

B

- Backing up
 - configuration parameters A-13
 - dump devices 10-2
 - guidelines 10-5
 - identifying Backup Server 10-2
 - master database 9-4, 10-1
 - on multiple volumes 10-6
 - planning 10-1
 - scheduling 10-5
 - transaction logs 9-13
 - using dbcc 9-18
- Backup
 - baseline 10-5
 - name 10-6
 - stopping 10-7
- Backup/Recovery configuration
 - parameters A-13
- Backup name 10-9
- Backup Server 10-2
 - messages 10-8, 10-10
- Bars
 - displaying 2-14
 - hiding 2-14
- Baseline backup 10-5
- bcp utility 9-13
- Benchmarks
 - creating 27-1
 - Historical Server and 27-6
 - sample interval in 27-6
 - session length and 27-4, 27-6
 - statistic type and 27-6
- Binding
 - See also* Unbinding
 - default to column 16-5, 22-32
 - default to column vs. user-defined
 - datatype 16-5
 - default to user-defined datatype 16-7
 - objects to a named cache 7-10
 - object to execution class 25-2, 25-17
 - rules 18-6

- rules to user-defined datatypes 18-7
- rule to column 18-6, 22-31
- scope and precedence rules 25-6
- Buffer pools**
 - changing 7-8
 - creating 7-5, 7-7
 - deleting 7-9
 - wash area 7-9
- Buffer prefetch** 24-5
- Building indexes** 24-4
- Buttons**
 - Consistency 24-9
 - tab 3-3
 - toolbar 2-14
- C**
- Cache** 24-2
 - See also* Named caches
 - and large I/Os 24-5
 - configuration parameters A-14
 - creating buffer pool 7-7
 - data 7-1
 - management of 7-1
 - procedure 7-1
 - property sheet 7-3
 - replacement strategy, index 24-5
 - replacement strategy, table 22-14
 - size 7-2
 - strategy 24-3, 24-5
 - setting 22-14
- Cache Bindings tab** 7-11
- Cache Dirty Grabs** 27-7
- Cache Hit Percentage** 27-7
- Cache Monitor** 26-14
- Canceling filters** 5-14
- Capacity** 10-3
- Capacity planning** 27-9
- Cascade** 13-2
- Change Password dialog box** 13-6
- Changing**
 - buffer pools 7-8
 - cache properties 7-3
 - configuration parameters 4-11
 - connection profile 4-6
 - database defaults 9-12
 - default devices 8-12
 - engine affinity 25-17
 - engines in an engine group 25-12
 - log buffer size 9-11
 - object properties 3-3
 - passwords 13-6
 - procedures 17-3
 - size of window pane 2-3
 - window display 2-6
- Character string filters** 5-16
- checkalloc** 9-20
- checkcatalog** 9-21
- Check constraints** 22-2
 - column 22-2, 22-8
 - creating 22-19
 - deleting 22-32
 - format of 22-9
 - ways of creating 22-2
- checkdb** 9-19
- Checking**
 - database consistency 9-18
 - index consistency 24-9
 - overall table consistency 22-27
- Checkpoint On Recovery** 10-9
- Checkpoints** 9-17
 - issuing automatically 9-17
 - transaction logs 9-17
- checktable** 22-27
- cis bulk insert batch size configuration**
 - parameter A-2
- cis connect timeout configuration**
 - parameter A-2
- cis cursor rows configuration**
 - parameter A-2
- cis packet size configuration**
 - parameter A-2
- CIS RPC handling** 6-2
- cis rpc handling** 23-3
- cis rpc handling configuration**
 - parameter A-2
- Class**
 - execution 25-2

- Closing
 - monitors 5-7
 - property sheet 3-3
- Clustered index 22-24
- Clustered indexes 24-1, 24-2
- Code
 - editing trigger 19-3
- Code editor 2-16
 - executing code 2-18
 - opening 2-17
 - syntax highlighting 2-18
 - toolbar 2-18
- Columns
 - adding 22-7, 22-13
 - binding defaults to 16-5, 22-32
 - binding rules 18-6
 - binding rules to 22-31
 - check constraint 22-2
 - constraints 22-8
 - default values 16-1
 - deleting check constraints 22-32
 - identity 20-3
 - match] 22-22
 - properties 22-30
 - property sheet 22-30
 - to match 22-22
 - unbinding defaults from 22-32
 - unbinding rules 18-9
 - unbinding rules from 22-31
 - views 21-1
- Column width, on monitors 5-8
- Command permissions
 - granting and revoking 14-8
 - group 14-18
 - roles 15-17
- Commands
 - access 12-2
 - create 12-6
 - executing triggers 19-1
 - permissions 12-2, 12-6, **12-9**
 - Transact-SQL 3-5
 - Verify Table 23-7
- Component Integration Services
 - See also* proxy table
- configuration parameters A-14
- logins 6-14
- options 23-2
- server class 6-6
- configuration file configuration
 - parameter 27-5, A-2
- Configuration parameters
 - Adaptive Server A-1
 - Adaptive Server Administration A-12
 - backup/recovery A-13
 - cache A-14
 - cis rpc handling 23-3
 - component integration services A-14
 - disk I/O A-14
 - enable cis 23-3
 - Error Log A-14
 - extended stored procedures A-15
 - functional groups of A-12 to A-18
 - general information A-15
 - languages A-15
 - lock manager A-15
 - meta-data caches A-16
 - processors A-17
 - rep agent thread administration A-17
 - security related A-18
 - setting 4-11
 - user environment A-18
- Configure command 4-11
- Configuring
 - Adaptive Server 4-10, 4-11
 - roles required 4-10
 - extended stored procedures 17-12
 - immediate and restart 4-10
 - monitors 5-3
 - Monitors folder 5-5
 - RPC handling 6-1
 - System Administrators and 4-10
- Connect dialog box 4-2
- Connecting
 - Adaptive Server 4-2
 - Adaptive Server to 4-4
 - at startup 4-5
 - checking environment 4-6
 - diagnosis of 4-6

- Connection profile 4-4
 - changing 4-6
 - creating 4-4
- Connections
 - diagnosing unsuccessful 4-6
- Connect toolbar button 2-15
- Considerations
 - binding rules 18-6
- Consistency
 - button 24-9
 - checking table 22-26
 - fix text for tables 22-29
 - index 22-27
 - reindexing tables 22-29
 - system tables 9-21
- Consistency checking
 - database 9-18
 - dbcc utility 9-18
 - index 24-9
- Constraints
 - check 22-2, 22-9
 - column 22-8
 - creating 22-2
 - creating table 22-19
 - default 22-2
 - description of 22-2
 - referential 22-2, 22-3, 22-21
 - rules 22-2
 - summary of 22-3
 - unique 22-24
- Container icons 2-3
- Context-sensitive help 3-5
- Context-sensitive menus
 - pop-up 2-13
 - shortcut 2-13
- Copying device contents 8-7
- Corrupted databases *See* Damaged Databases
- Counts, of activity in monitors 26-4
- cpu accounting flush interval configuration
 - parameter A-2
- cpu grace time configuration
 - parameter A-2
- CPU usage
 - configuration parameters
 - affecting A-2
- Creating
 - benchmarks 27-1
 - buffer pools 7-7
 - check constraints 22-9
 - column constraints 22-8
 - connection profile 4-4
 - constraints 22-2
 - database devices 8-3
 - databases 9-2
 - defaults 16-1
 - devices 8-3
 - dump devices 10-2
 - engine group 25-9
 - execution class 25-15
 - extended stored procedure 17-8
 - foreign key 22-22
 - groups 14-12
 - guest users 9-3
 - indexes 24-2
 - logins 13-1
 - named cache binding 7-11
 - named cache buffer pools 7-5
 - named caches 7-2
 - objects 3-1
 - primary key 22-10
 - procedures 17-1
 - proxy table 23-3
 - referential constraints 22-21
 - remote servers 6-4
 - roles 15-2
 - rules 18-1
 - segments 11-2
 - segment thresholds 11-8
 - table columns 22-7, 22-13
 - tables 22-2, 22-6
 - toolbars and 1-3
 - triggers 19-2
 - unique constraints 22-24
 - user-defined datatypes 20-1
 - user-defined roles 15-2
 - users 14-2
 - views 21-1

- Cumulative counts 26-4
- Cumulative data on monitors 5-11
- Cumulative rates 26-3
- Cumulative rates in processes 26-4
- curread change w/ open cursors configuration
 - parameter A-2
- current audit table configuration
 - parameter A-2

- D**
- d_master* device 8-2
- Daily monitoring 27-7
- Damaged databases, deleting 9-4
- Data
 - retrieving 24-1
 - sorting 24-4
- Data and log cache 7-5
- Database Allocations dialog box 9-9
- Database definition language
 - generating 3-5
- Database Device Databases tab 8-6
- Database Device property sheet 8-5
- Database devices
 - See Devices
- Database Device Segments tab 8-7
- Database object filter 5-17
- Database objects
 - creating 3-1
- Database options
 - read only 10-9
- Database Options dialog box 9-12
- Database Owner 12-4
- Database property sheet 9-5
- Databases
 - backing up 9-4
 - backup guidelines 10-5
 - binding to a named cache 7-10
 - changing log buffer size 9-11
 - checking consistency 9-18
 - checkpoints 9-17
 - creating 9-2
 - default options 9-12
 - deleting 9-4
 - deleting damaged 9-4
 - devices 8-3, 8-6
 - displaying owners 13-7
 - displaying properties 9-4
 - displaying storage allocation 9-7
 - estimating size 9-1
 - expanding 9-8
 - fixing allocation errors 9-20
 - generating DDL 9-15
 - guest users 9-3
 - increasing storage 9-7
 - master* 8-3
 - model* 8-2
 - permissions 12-4
 - properties 9-5
 - recovering 10-8
 - reducing size 9-7
 - repairing and dropping 9-22
 - restoring 10-9
 - segments 8-6, 11-1
 - setting online 9-22
 - storing 8-2, 8-6
 - sybssystemprocs* 8-2
 - tempdb* 8-2
 - transferring ownership 9-6
- database size configuration variable 9-3
- Data cache 7-1
- Data Cache Hit Percentage* 27-10
- Data Cache Monitor 26-17
- Data cache name
 - filter on 5-19
- Data Definition Language in Transaction
 - database option 9-14
- Data logs 8-13
- Data pages 9-17
- Datatypes 22-7
 - See also User-defined datatypes 20-1
- dbcc**
 - index consistency checking 24-9
 - options 9-19
 - table consistency checking 22-26
 - utility 9-18, 22-26
- dbcc checkalloc** 9-20
- dbcc checkcatalog** 9-21

- dbcc checkdb 9-19
- dbcc checktable 22-27
- dbcc fix_text 22-29
- dbcc indexalloc 24-9
- dbcc reindex 22-29
- dbcc tablealloc 22-28
- dbo 12-4
- DDL
 - generating database 9-15
 - generating scripts 3-5
- Deactivating
 - roles 15-7, 15-8
- deadlock checking period configuration
 - parameter A-2, A-3
- deadlock retries configuration
 - parameter A-2
- Default
 - database 13-1, 13-3
 - device 8-12
 - execution attributes 4-14
 - language 13-1, 13-3
 - login mapping 6-12
 - multiple devices 8-12
 - role activation 13-9, 15-8
 - storage devices 8-2
- default character set id configuration
 - parameter A-3
- default database size configuration
 - parameter A-3
- default fill factor percent configuration
 - parameter A-3
- default language id configuration
 - parameter A-3
- Default Login Mapping tab 6-12
- default network packet size configuration
 - parameter A-3
- Defaults
 - as constraints 22-2
 - binding to column 16-5
 - binding to columns 22-32
 - changing device 8-12
 - creating 16-1
 - deleting 16-3
 - dependencies 16-4
 - displaying properties 16-3
 - options 9-12
 - permissions 12-9
 - property sheet 16-3
 - unbinding from column 16-8
 - unbinding from columns 22-32
 - unbinding from user-defined datatype 16-8
- default* segment 11-1
- default sortorder id configuration
 - parameter A-3
- Defining
 - columns 22-7
- Delete permission 12-7, 14-7, 14-17, 15-16, 21-7, 22-17
- Delete toolbar button 2-15
- Deleting
 - aliases 13-11
 - buffer pools 7-9
 - check constraints 22-32
 - database 9-4
 - database devices 8-4
 - databases, damaged 9-4
 - default 16-3
 - dump devices 10-3
 - engine from engine group 25-13
 - engine groups 25-10
 - execution class 25-15
 - extended stored procedures 17-9
 - groups 14-12
 - index 24-6
 - logins 13-2, 13-4
 - named cache binding 7-13
 - named cache buffer pools 7-9
 - named caches 7-3
 - objects 3-4
 - proxy table 23-6
 - remote servers 6-5
 - role assignment 15-6
 - rules 18-3
 - segments 11-2
 - segment thresholds 11-11
 - stored procedures 17-2
 - tables 22-11

- triggers 19-2
- user-defined datatypes 20-3
- user-defined roles 15-2
- users 14-3
- view
 - Views
 - deleting 21-2
- Dependencies
 - default 16-4
 - displaying procedure 17-5
 - displaying rule 18-5
 - displaying table 22-14
 - displaying trigger 19-5
 - displaying user datatype 20-5
 - displaying view 21-5
 - extended stored procedure 17-11
 - table 22-14
- Details
 - display format 2-6
 - toolbar button 2-15
- Device Hit Percentage* 27-7
- Device I/O Monitor 26-23
- Device name
 - filter on 5-24
- Devices
 - allocating space on 9-8
 - changing default 8-12
 - copying contents 8-7
 - creating 8-3
 - d_master* 8-2
 - databases 8-3, 8-6
 - default storage 8-2
 - deleting 8-4
 - disks 8-2
 - displaying properties 8-5
 - dump 10-2
 - extending segments 11-4
 - master 8-12
 - multiple default 8-12
 - name 10-6, 10-9
 - number 8-4
 - segments 8-6
 - segments pointing to 11-1
 - separating usage 8-14
 - storage space on 8-6
- Device Size dialog box 9-9
- Diagnosing
 - connections 4-6
- Dialog boxes
 - access 12-13
 - Add Alias 13-10
 - Add Database Device to Segment 11-5
 - Add Memory Pool 7-8
 - Add New Threshold 11-9
 - Add Specific Mapping 6-14
 - Add User to Group 14-15
 - Change Password 13-6
 - Device Size 9-9
 - Edit Check Constraints 22-8
 - Edit Primary-key Columns 22-10
 - granting permissions 12-13
 - Logon 4-2
 - New Check Constraint Properties 22-20
 - New Database Owner 9-6
 - New Default Properties 16-2
 - New User-defined Datatype 20-2
 - object 3-2
 - refreshing 2-16
 - Select a Creator 18-2
 - Server Configuration 4-11
- Direct updates to system tables A-2
- disable character set conversions
 - configuration parameter A-3
- Disabling disk mirroring 8-10
- Disconnect command 4-6
- Disconnecting
 - from Adaptive Server 4-6
- Disconnect on timeout 6-7
- Disconnect toolbar button 2-15
- Disk I/O configuration
 - parameters A-14
- disk i/o structures configuration
 - parameter A-3
- Disks
 - adding mirroring 8-9
 - devices 8-2

- disabling mirroring 8-10
 - mirroring 8-3, 8-7
 - re-enabling mirroring 8-11
 - Display
 - formats, selecting 2-6
 - synchronizing windows 2-16
 - updating 2-16
 - Displaying
 - Adaptive Server processes 4-12
 - Adaptive Server properties 4-9
 - bars 2-14
 - cache properties 7-3
 - column properties 22-30
 - database device properties 8-5
 - database properties 9-4
 - databases owned by login 13-7
 - data on monitors 5-11
 - default dependencies 16-4
 - dependencies for extended stored
 - procedures 17-11
 - dump device properties 10-4
 - engine group properties 25-10
 - groups 14-13
 - icons 2-14
 - index properties 24-6
 - logins 13-2
 - named caches 7-3
 - objects on segments 11-6
 - procedure dependencies 17-5
 - procedure properties 17-4
 - remote server properties 6-5
 - role properties 15-3
 - rule dependencies 18-5
 - rule properties 18-4
 - segment properties 11-3
 - storage allocation 9-7
 - table dependencies 22-14
 - toolbar and status bar 2-8
 - trigger dependencies 19-5
 - trigger properties 19-4
 - unique constraint properties 22-25
 - user datatype dependencies 20-5
 - user-defined datatype properties 20-4
 - user processes 4-12
 - users 14-3
 - view dependencies 21-5
 - view permissions 21-6
 - views 21-3
 - Drag-and-drop
 - activities summarized 2-9
 - icons 2-8
 - tables 22-17
 - users 22-17
 - Dragging and dropping icons 2-8
 - Drop-down list toolbar control 2-15
 - Dropping
 - databases 9-22
 - objects 3-4
 - dump database command 9-13
 - Dump devices
 - backing up 10-2
 - creating 10-2
 - deleting 10-3
 - displaying properties 10-4
 - Dump devices property sheet 10-4
 - Dumping transaction logs 9-13
 - Duplicate
 - key 24-2
 - row 24-3
 - Duplicating. *See* Mirroring and Copying
- E**
- EC1 25-2
 - EC2 25-2
 - EC3 25-2
 - Edit Check Constraints dialog box 22-8
 - Editing
 - connection profile 4-6
 - procedure 17-3
 - trigger code 19-3
 - Edit menu 2-12
 - Editor
 - code 2-16
 - table 22-4
 - Edit Primary-key Columns dialog
 - box 22-10
 - enable cis 23-3

- enable cis configuration parameter A-3
- enable rep agent threads configuration parameter A-3
- Encrypt network password 6-7
- Enforcing referential integrity 19-1, 22-21
- Engine Activity Monitor 26-27
- Engine affinity 25-1, 25-3
 - changing 25-17
- Engine group
 - adding engines 25-13
 - changing engines 25-12
 - creating 25-9
 - definition of 25-9
 - deleting 25-10
 - deleting engine from 25-13
 - properties 25-10
 - setting for process 4-14
- Engine number
 - filter on 5-25
- Engines
 - multiple 25-3
 - processing objects 25-3
- Error log configuration
 - parameters A-14
- esp execution priority configuration parameter A-3
- esp execution stack size configuration parameter A-3
- esp unload dll configuration parameter A-3
- Estimating
 - database size 9-1
 - running sp_estspace 9-2
 - space requirements 9-2
- event buffers per engine configuration parameter A-4
- event log computer name configuration parameter A-4
- event logging configuration parameter A-4
- Examples
 - check constraints 22-9
- Exclusivity
 - role 15-9
- executable code size configuration parameter A-4
- Execute permission 12-7
- execute permission 17-6
- Executing
 - code in code editor 2-18
 - procedures 17-6
 - triggers 19-1
- Execution
 - attributes
 - setting for process 4-14
 - objects 25-1
 - preference 25-2
- Execution attributes
 - default 4-14
- Execution Attributes tab 4-14
- Execution class
 - application scope 25-19
 - binding objects to 25-2
 - binding object to 25-17
 - changing engine affinity 25-17
 - creating 25-15
 - definition 25-2
 - deleting 25-15
 - login scope 25-18
 - predefined 25-2
 - priority of 25-4
 - properties 25-16
 - resolving precedence conflicts 25-8
 - scope, specifying 25-18
 - scope and precedence 25-6
 - unbinding objects 25-19
 - user defined 25-3
- Expanding
 - databases 9-8
 - object tree 2-4
- Explicit login mapping 6-13
- Expression, rule 18-3
- Extended stored procedures
 - configuration parameters 17-12, A-15
 - creating 17-8
 - definition of 17-8
 - deleting 17-9

- dependencies 17-11
- parameters 17-11
- setting permissions 17-11
- Extending segments 11-4

F

- File menu 2-12
- Fill factor 22-24, 24-3, 24-4, 24-5
- Filters
 - canceling 5-14
 - character string 5-16
 - database object 5-17
 - data cache name 5-19
 - device name 5-24
 - engine number 5-25
 - lock status 5-20
 - lock type 5-21
 - numeric data 5-14
 - process ID 5-22
 - process state 5-23
 - setting 5-13
 - threshold 5-15
 - top n 5-14
 - wildcards in 5-18
- fix_text 22-29
- Fixing
 - allocation errors for a database 9-20
- Fixing text 22-29
- Fonts
 - on monitors 5-8
- Foreign key
 - creating 22-22
 - properties 22-23
- For load 9-3
- freelock transfer block size configuration
 - parameter A-4
- Free space 11-6
 - changing on threshold 11-10
- Future Only 16-7, 18-8

G

- General information configuration
 - parameters A-15
- Generating DDL
 - database 9-15
- global async prefetch limit configuration
 - parameter A-4
- Granting
 - command permissions for
 - group 14-18
 - command permissions for users 14-8
 - object permissions 14-6
 - object permissions for a group 14-16
 - object permissions for users 14-5
 - permissions for roles 15-15, 15-17
 - procedure permissions 17-6
 - table permissions 22-16
 - view permissions 21-6
 - with grant 12-6
- Graphical user interface 5-1
- Group
 - property sheet 14-13
- Groups
 - adding users 14-14
 - changing membership 14-4
 - creating 14-12
 - deleting 14-12
 - displaying properties 14-13
 - granting and revoking command
 - permissions 14-18
 - object permissions 14-16
 - permissions 12-5
 - procedure permissions 17-6
 - Public 14-2
 - removing users 14-15
 - sharing permissions 14-4, 14-12
 - table permissions 22-16
 - view permissions 21-6
- Guest user 12-1
- Guest users
 - creating 9-3
 - databases 9-3
- Guidelines, backing up 10-5

H**Help**

- how to use xxxvi
- menu 2-13
- toolbar buttons 2-16

Help, online 3-5**Hiding**

- bars 2-14
- icons 2-14
- monitor status bar 5-9
- toolbar and status bar 2-8

Hierarchical list 2-4**Hierarchy**

- example of 15-13
- roles 12-3, 15-11

Historical Server 27-6

- definition of 5-1
- features 5-3
- Monitor Client Library and 5-3
- playback and 5-3

housekeeper free write percent configuration

- parameter A-4

Hysteresis value 11-7**I****i/o accounting flush interval configuration**

- parameter A-4

i/o polling process count configuration

- parameter A-4

I/Os, large 7-5, 24-5**Icons** 1-2

- container 2-3
- displaying 2-14
- dragging and dropping 2-9
- hiding 2-14
- object 2-3

identity burning set factor configuration

- parameter A-4

IDENTITY column 9-15**Identity columns** 20-3**identity grab size configuration**

- parameter A-4

Inactivating**roles** 15-7**Increasing**

- database storage space 9-7

indexalloc 24-9**Indexes**

- allocating pages 24-9
- binding to a named cache 7-10
- building 24-4
- cache replacement strategy 24-5
- checking consistency 22-27, 24-9
- clustered 24-1, 24-2
- creating 24-2
- deleting 24-6
- estimating space requirements 9-2
- keys 24-2
- maximum rows per page 24-5
- nonclustered 22-18, 24-1
- nonunique 24-1
- on segments 11-6
- properties 24-6
- property sheet 24-6
- sorting data 24-4
- statistics 24-11
- unique 24-1, 24-2
- unique constraints 22-24
- values 24-2

In Future Only 16-7, 18-8**Insert permission** 12-7**insert permission** 14-7, 14-17, 15-15, 21-7, 22-17**Integrity**

- checks 24-9
- referential 19-1

Interfaces file 4-1

- remote servers and 6-2, 6-3

Issuing checkpoints 9-17**J****Java virtual machine** 5-4**K****Keyboard shortcuts** 2-11

Keys

- See also* Primary key
- creating foreign 22-22
- index 24-2
- primary 22-9, 22-10

Killing

- user processes 4-15

L**Languages configuration**

- parameters A-15

Large buffer prefetch 24-5**Large I/Os** 7-5, 24-5**Large icon**

- display format 2-6
- toolbar button 2-15

Last-chance threshold 11-7**LASTONLINE** 25-10**Left pane** 2-3**List display format** 2-6**List toolbar button** 2-15**Locking**

- login 14-3

Locking logins 13-4, 14-3**Lock management configuration**

- parameters A-15

lock promotion HWM configuration

- parameter A-5

lock promotion LWM configuration

- parameter A-5

lock promotion PCT configuration

- parameter A-5

Locks *Denied Rat* 27-7**lock shared memory configuration**

- parameter A-5

Lock status

- filter on 5-20

Lock type

- filter on 5-21

log audit logon failure configuration

- parameter A-5

log audit logon success configuration

- parameter A-5

Log buffer size, changing 9-11**Log files**

- tuning with monitors 27-10

Login

- automatic activation of roles 13-9
- property sheet 13-3
- removing role activation 13-9

Login Alias/Users tab 13-10**Logins**

- access 12-1
- aliases 12-2
- aliasing to users 14-9
- assigning aliases 13-9
- assigning roles to 13-7, 15-5
- Component Integration Services 6-14
- creating 13-1
- databases owned 13-7
- default login mapping 6-12
- deleting 13-2, 13-4
- deleting aliases 13-11
- displaying 13-2
- explicit mapping 6-13
- locking 14-3
- locking and unlocking 13-4
- mapping 6-2
 - removing 6-14
- mapping for remote server 6-10
- modifying 13-3
- password, changing 13-6
- removing alias to user 14-11
- removing Component Integration Services mapping 6-16
- removing role assignment 15-6
- scope 25-18
- sharing 13-9

Logon dialog box 4-2**Log only caches** 7-5**Logs**

- allocating space 9-2
- data 8-13
- maintaining 8-3
- separating 8-13
- storing 8-2
- transaction 8-2, 8-13

logsegment segment 11-1

M

Maintaining logs 8-3

Managing

resources 8-1

Mapping

default logins 6-12

logins 6-2

logins explicitly 6-13

logins for Component Integration
Services 6-14

remote server logins 6-10

Maps, allocation 22-28

master database

backing up 9-4, 10-1, 10-5

creating devices 8-3

Master device 8-12

max async i/o's per engine configuration

parameter A-5

max async i/o's per server configuration

parameter A-5

max cis remote connection configuration

parameter A-5

max cis remote servers configuration

parameter A-5

max engine freelocks configuration

parameter A-5

Maximizing performance 8-14

Maximum index rows per page 24-5

Maximum table rows per page 24-8

max network packet size configuration

parameter A-5

max number of network listeners configuration

parameter A-5

max online engines configuration

parameter A-5

max parallel degree configuration

parameter A-5

max roles enabled per user configuration

parameter A-6

max scan parallel degree configuration

parameter A-6

max SQL text monitored configuration

parameter A-5

Membership exclusivity 15-9

memory alignment boundary configuration

parameter A-6

Memory pools 7-5

Memory use configuration

parameters A-15

Memory Utilization Monitor 26-30

Menus 2-12

monitor 5-10

pop-up 2-13

shortcut 2-13

Messages

Backup Server 10-8, 10-10

Meta-data caches

configuration parameters

A-16

Mirroring

adding 8-9

benefits 8-8

costs 8-8

disabling 8-10

disks 8-3, 8-7

recovery benefits 8-13

re-enabling 8-11

Mnemonic 2-11

model database 8-2, 9-3, 10-1, 20-1

Modifying

default databases 13-3

default languages 13-3

logins 13-3

procedure 17-3

mon_rpc_connect stored procedure 5-4, 5-5

Monitor Client Library 5-2

Historical Server and 5-3

Monitoring

non-sa-role 5-4, 5-6

with non-sa_role 5-6

Monitoring strategies 27-6

log files 27-10

recognizing problems 27-7

tempdb and 27-10

Monitors

- Adaptive Server name and 5-28
 - closing 5-7
 - column width on 5-8
 - configuring 5-3
 - fonts on 5-8
 - menus on 5-10
 - opening 5-6
 - privileges required 5-4, 5-5
 - rates in 26-3
 - refreshing 5-13
 - sample interval 5-11
 - sample or cumulative data 5-11
 - statistic type 5-11
 - status bar on 5-9
 - summary of 26-2
 - troubleshooting 5-27
 - Monitor Server 5-1, 5-3
 - names 5-5
 - number of connections 5-27
 - sp_sysmon and 27-11
 - starting 5-4
 - Monitor Server tab 5-5
 - Monitor session length
 - benchmarks and 27-4, 27-6
 - Monitors folder 5-4
 - configuring 5-5
 - Monitor Viewer 5-1, 5-3
 - Monitor window 5-8
 - Most Recently Used replacement
 - strategy 22-14, 24-5
 - MRU replacement strategy 22-14, 24-5
 - msg confidentiality reqd configuration
 - parameter A-6
 - msg integrity req configuration
 - parameter A-6
 - msg origin checks req configuration
 - parameter A-6
 - msg out-of-seq checks req configuration
 - parameter A-6
 - msg replay detection reqd configuration
 - parameter A-6
 - Multiprocessor machines 25-3
 - Multiprocessor settings. *See* max online engines and min online engines A-5
 - Multivolume backup and restore 10-10
 - Mutual exclusivity 12-3
 - role 15-9
- ## N
- Named caches
 - See also* Cache
 - binding objects to 7-10
 - changing size of 7-4
 - changing type of 7-4
 - creating 7-2
 - creating a binding 7-11
 - data and log 7-5
 - defining buffer pools 7-5
 - deleting 7-3
 - deleting a binding 7-13
 - displaying 7-3
 - log only 7-5
 - status 7-2, 7-10
 - wash buffers 7-9
 - Navigating
 - to objects 1-3
 - window lists 2-4
 - Network Activity Monitor 26-31
 - Network communications configuration
 - parameters A-16
 - Network I/O Rate* 27-7
 - New Check Constraint Properties dialog box 22-20
 - New Database Owner dialog box 9-6
 - New Default Properties dialog box 16-2
 - New User-defined Datatype dialog box 20-2
 - No Checkpoint on Recovery database option 9-14
 - No Free Space Accounting database option 9-15
 - Non-clustered indexes
 - checking consistency 22-27
 - Nonclustered indexes 22-18, 24-1
 - non-sa_role monitoring 5-4, 5-6
 - No Nulls 20-3
 - Nonunique indexes 24-1

Null values 16-5, 20-3
 number of alarms configuration
 parameter A-6
 number of aux scan descriptors configuration
 parameter A-6
Number of Connections 27-7
 Number of connections
 to Monitor Server 5-27
Number of connections 27-4
 number of devices configuration
 parameter A-6
 number of index trips configuration
 parameter A-6
 number of languages in cache configuration
 parameter A-6
 number of locks configuration
 parameter A-7
 number of mailboxes configuration
 parameter A-7
 number of messages configuration
 parameter A-7
 number of oam trips configuration
 parameter A-7
 number of open databases configuration
 parameter A-7
 number of open indexes configuration
 parameter A-7
 number of open objects configuration
 parameter A-7
 number of preallocated extents configuration
 parameter A-7
 number of remote connections configuration
 parameter A-7
 number of remote logins configuration
 parameter A-7
 number of remote sites configuration
 parameter A-7
Number of Selects 27-7
 number of sort buffers configuration
 parameter A-7
 number of user connections configuration
 parameter A-7
 number of worker processes configuration
 parameter A-7

Number Sleeping on Locks 27-7

O

o/s file descriptors configuration
 parameter A-8
 OAM 22-28
 Object Allocation Map 22-28
 Object Lock Status Monitor 26-37
 Object Page I/O Monitor 26-42
 tempdb and log files in 27-10
 Object permissions
 granting and revoking 14-5
 roles 15-15
 Objects
 changing properties 3-3
 creating 3-1
 deleting 3-4
 dialog boxes 3-2
 displaying ownership 14-5
 execution 25-2
 group permissions 14-16
 icons 1-2, 2-3
 navigating to 1-3
 owned by user 14-5
 owner permissions 12-5
 owners 19-1
 permissions 12-6, 14-16
 related 3-4
 segments on 11-6
 selecting 2-3
 user permissions 14-5
 Object tree 2-3, 2-4
 expanding 2-4
 OmniConnect 23-2
 Online engines. *See* max online engines and
 min online engines A-5
 Online help xxxvi, 3-5
 Only Usable By Owner 10-9
 open index hash spinlock ratio configuration
 parameter A-8
 open index spinlock ratio configuration
 parameter A-8
 Opening

- code editor 2-17
- monitors 5-6
- Process Current SQL Statement
 - Monitor 5-6
- property sheet of related object 3-4
- property sheets 3-2
- open object spinlock ratio configuration
 - parameter A-8
- Open Server 5-1
- oper_role 12-4
- Operating system resources
 - configuration parameters A-16
- Operator permissions 12-4
- Options
 - dbcc 9-19
 - default 9-12
- Overriding
 - user permissions 14-2
- Owners
 - procedures 17-1
- Ownership
 - displaying 14-5
 - permissions and 12-4
 - transferring databases 9-6
 - triggers 19-1
- P**
- page lock spinlock ratio configuration
 - parameter A-8
- Pages
 - allocating 22-28, 24-9
 - free space and thresholds 11-6
- page utilization percent configuration
 - parameter A-8
- Panes
 - window 2-3
- Parameters
 - adding to extended stored
 - procedures 17-11
 - extended stored procedure 17-12
 - functional groups 4-11
 - setting Adaptive Server 4-10
- Parameters tab
 - Processes property sheet 4-13
- partition groups configuration
 - parameter A-8
- partition spinlock ratio configuration
 - parameter A-8
- Password
 - changing 13-6
 - role 15-4
 - trusted 6-14
- Path name of device 10-3
- Pausing
 - Adaptive Server 4-7
- Pending status 7-2
- Pending value 4-12
- Performance 25-1
 - benefits 8-8
 - costs 8-8
 - maximizing 8-14
 - segment benefits 11-1
- Performance Summary Monitor 26-45
 - daily monitoring and 27-7
- Performance Trends Monitor 26-52
 - setting up 5-26
- perform disk i/o on engine 0 configuration
 - parameter A-8
- permission cache entries configuration
 - parameter A-8
- Permissions
 - access 12-1, 12-13
 - Adaptive Server configuration
 - variables, modifying 4-10
 - command 12-6, 12-9
 - command for group 14-18
 - creating referential constraints 12-7, 14-7, 14-17, 15-16
 - damaged databases, deleting 9-4
 - database owners 12-4
 - databases, backing up 10-6
 - databases, creating backup tasks 10-6
 - default 12-9
 - deleting a row 12-7, 14-7, 14-17, 15-16
 - deleting a view 21-7, 22-17
 - dialogs 12-13
 - dump devices, deleting 10-3

- executing a procedure 12-7
- extended stored procedure 17-11
- granting and revoking object 14-16
- group 12-5
- groups 14-16
- group versus user 14-2
- inserting a row 12-7, 14-7, 14-17, 15-15, 21-7, 22-17
- modifying 17-6, 21-6, 22-16
- object 12-6, 14-5, 14-16
- object owners 12-5
- Operators 12-4
- overriding 14-2
- procedure 17-6
- Public group 12-9
- referential 12-7
- referential constraint referencing a view 21-7, 22-17
- roles 12-2, 15-15, 15-17
- selecting a table or view 12-7, 14-7, 14-17, 15-15
- selecting a view 21-7, 22-17
- sharing 12-2, 13-9, 14-12
- shortcuts 22-17
- summary 12-9
- System Administrators 12-3
- System Security Officers 12-4
- table 22-16
- tabs 12-13
- transfers and 12-9
- updating a row 12-7, 14-7, 14-17, 15-15, 21-7, 22-17
- view 21-6
- with cascade 12-6
- with grant 12-6
- Physical memory configuration parameters A-17
- Planning
 - backup 10-1
 - backup and recovery 10-1
 - recovery 10-1, 10-8
- Playback 5-3
- Pools, buffer 7-5
- Pop-up menus 2-13
- Precedence
 - for execution classes 25-6
 - resolving conflicts 25-8
- Precedence rule 25-7
- Precision 20-2
- Preference
 - execution 25-2
- Prefetch, large buffer 24-5
- Presorted data 24-4
- Primary key 22-9
 - specifying 22-10
- print deadlock information configuration parameter A-9
- print recovery information configuration parameter A-9
- Priority
 - setting for process 4-14
- Priority run queue 25-4
- Privileges. *See* Permissions.
- Procedure cache 7-1
- procedure cache percent configuration parameter A-9
- Procedure Hit Percentage* 27-7
- Procedures
 - changing 17-3
 - changing on thresholds 11-10
 - creating 17-1
 - displaying dependencies 17-5
 - executing 17-6
 - ownership 17-1
 - permissions 17-6
 - properties 17-4
 - property sheet 17-4
 - recompiling 22-26, 24-11
 - sp_estspace 9-2
- Process Activity Monitor 26-54
 - stopped processes and 27-8
- Process Current SQL Statement Monitor 26-59
 - opening 5-6
 - tuning statements with 27-9
- Processes
 - cumulative data for 26-4
 - displaying 4-12

- engine group and priority 4-14
 - killing 4-15
 - Parameters tab 4-13
 - researching stopped 27-8
 - Process ID
 - filter on 5-22
 - Processors configuration
 - parameters A-17
 - Process state
 - filter on 5-23
 - Properties
 - Adaptive Server 4-9
 - cache 7-3
 - changing 3-3
 - columns 22-30
 - default 16-3
 - displaying database 9-4
 - displaying dump device 10-4
 - displaying segment 11-3
 - engine group 25-10
 - execution class 25-16
 - foreign key 22-23
 - group 14-13
 - index 24-6
 - login 13-2
 - procedure 17-4
 - remote server 6-5
 - role 15-3
 - rule 18-4
 - toolbar button 2-15
 - trigger 19-4
 - unique constraints 22-25
 - user 14-3
 - user-defined datatype 20-4
 - view 21-3
 - Property sheet
 - cache 7-3
 - closing 3-3
 - Database Device 8-5
 - opening 3-2
 - tables
 - Displaying
 - table properties 22-12
 - tabs 3-3
 - Property sheets
 - Column 22-30
 - Database 9-5
 - Default 16-3
 - Dump Device 10-4
 - Group 14-13
 - Index 24-6
 - Login 13-3
 - Procedure 17-4
 - Remote Server 6-5
 - Rule 18-4
 - Segment 11-3
 - Server Properties 4-9
 - Table 22-12
 - Trigger 19-4
 - User 14-4
 - View 21-3
 - Proxy table
 - completing in table editor 23-5
 - creating 23-3
 - creating new local and remote 23-1
 - definition of 23-1
 - deleting 23-6
 - map to existing table 23-1
 - new columns 23-5
 - new table versus existing table 23-6
 - verifying 23-6
 - Public
 - group permissions 12-9
 - membership 14-2
- ## R
- Rates, in monitors 26-3
 - Read Only Database database
 - option 9-14
 - Read Only database option 10-9
 - Read Only status
 - security 6-7
 - Recompile 17-2
 - Recompiling
 - procedures and triggers 22-26, 24-11
 - Recovery
 - clean 8-12

- configuration parameters A-13
- database 10-8
- mirroring benefits 8-8, 8-13
- planning 10-1, 10-8
- separating logs 8-13
- strategies 8-3
- transaction log 10-8
- recovery interval 9-17
- recovery interval in minutes configuration
 - parameter A-9
- Reducing database size 9-7
- Re-enabling disk mirroring 8-11
- Reference
 - permission 12-7
 - type 22-22
- references permission 14-7, 14-17, 15-16, 21-7, 22-17
- Referential constraint permissions 12-7
- Referential constraints 22-2, 22-3
 - description of 22-21
- Referential integrity
 - enforcing 19-1
- Referential integrity, enforcing 22-21
- Refreshing
 - data display on monitors 5-13
 - windows and dialog boxes 2-16
- Refresh time 5-9
- Reindex
 - table 22-29
- reindex 22-29
- Remote login mapping 6-10
- Remote procedure calls (RPCs) 6-1
- Remote Server
 - property sheet 6-5
- Remote server
 - default login mapping 6-12
 - explicit login mapping 6-13
 - mapping logins for Component Integration Services 6-14
 - proxy table 23-1
 - server class 6-6
- remote server pre-read packets configuration
 - parameter A-9
- Remote Server properties
 - Default Login Mapping tab 6-12
 - Specific Login Mapping tab 6-13
- Remote servers
 - creating 6-4
 - deleting 6-5
 - displaying 6-5
 - interfaces file and 6-2, 6-3
 - login management 6-10
 - setting up for 6-2
 - trusted password 6-14
- Remote table
 - proxy table 23-1
- Removing
 - alias to login 14-11
 - automatic role activation 13-9, 15-8
 - Component Integration Services login mapping 6-16
 - role assignments 15-6
 - specific login mapping 6-14
 - users from groups 14-15
- Rep agent thread administration
 - configuration parameters A-17
- Repairing
 - databases 9-22
- Replacement strategy, MRU 24-5
- Requires Restart 4-10
- Resetting configuration
 - parameters 4-11
- Resolving
 - precedence conflicts 25-8
- Resources
 - restricting use of 25-4
- Restore
 - stopping 10-10
- Restoring
 - databases 10-9
- Restoring a database
 - from multiple volumes 10-10
- Retention time 10-6
- Retrieving data 24-1
- Revoking
 - access 12-13
 - command permissions 14-8

- command permissions for
 - group 14-18
 - object permissions 14-6
 - object permissions for a group 14-16
 - permissions 13-2
 - procedure permissions 17-6
 - table permissions 22-16
 - view permissions 21-6
 - with cascade 12-6
- Rewind 10-6, 10-9
- Right pane 2-3
- Role
- activating automatically 13-9, 15-8
 - properties 15-3
 - removing automatic activation 13-9, 15-8
- Roles 12-3, 13-4
- access 12-2, 12-13
 - activating and deactivating 15-7
 - Adaptive Server configuration
 - variables, modifying 4-10
 - assigning to a login 15-5
 - assigning to login 13-7
 - command permissions 15-17
 - configuring Adaptive Server 4-10
 - containing 15-11
 - creating 15-2
 - damaged databases, deleting 9-4
 - databases, backing up 10-6
 - databases, creating backup tasks 10-6
 - deactivating 15-8
 - deleting 15-2
 - dump devices, deleting 10-3
 - example of hierarchy 15-13
 - hierarchy 12-3, 15-11
 - mutual exclusivity 15-9
 - oper_role 12-4
 - password 15-4
 - permissions on objects 15-15
 - properties 15-3
 - removing assignment to login 15-6
 - resetting Configuration options 4-10
 - sso_role 12-4
 - subordinate 15-11
 - System Administrator 12-3
 - System Administrators 4-10
 - System Security Officer 6-10
 - System Security Officers 4-10
- Rows per page 22-24, 24-3
- RPC
- handling 23-3
 - configuring 6-1
 - security model 6-7, 6-8
- RPCs 6-2
- Rules
- as constraints 22-2
 - binding, considerations for 18-6
 - binding to columns 18-6, 22-31
 - binding to user-defined datatypes 18-7
 - creating 18-1
 - defining values 18-1
 - deleting 18-3
 - displaying dependencies 18-5
 - expression in definition of 18-3
 - properties 18-4
 - property sheet 18-4
 - unbinding from column 18-9
 - unbinding from columns 22-31
 - unbinding from user-defined datatypes 18-9
 - ways of creating 22-2
- runnable process search count configuration parameter A-9
- Running sp_estspace 9-2
- S**
- sa_role 12-3
 - Sample data on monitors 5-11
 - Sample interval 5-11
 - benchmarks and 27-6
 - variations in 5-12
 - Sample rates 26-3
 - Scale 20-2, 22-7
 - Scheduling back ups 10-5
 - Scope 10-6, 10-9
 - All 25-18

- application 25-19
- execution class 25-6
- login 25-18
- specifying 25-18
- Scope rule 25-8
- secure default login configuration
 - parameter A-9
- Security
 - access to Adaptive Server 12-1
 - configuration parameters A-18
 - read only status 6-7
 - RPC 6-8
- Security options
 - Adaptive Server 6-7
- Segment Contains tab 11-6
- Segment property sheet 11-3
- Segments
 - changing thresholds 11-10
 - creating 11-2
 - creating thresholds 11-8
 - databases 8-6
 - definition 11-1
 - deleting 11-2
 - deleting thresholds 11-11
 - displaying properties 11-3
 - extending 11-4
 - mapping of devices 8-6
 - objects on 11-6
 - performance benefits 11-1
 - placing tables on 22-18
 - storing databases 11-1
 - thresholds on 11-6
 - transaction log on 11-7
- Segment Thresholds tab 11-8
- Select a Creator dialog box 18-2
- Selecting
 - display formats 2-6
 - object 2-3
- select on syscomments text configuration
 - parameter A-9
- Select permission 12-7
- select permission 14-7, 14-17, 15-15, 21-7, 22-17
- Separating
 - disk usage 8-14
 - logs 8-13
 - performance benefits 8-14
- Server class 6-6
- Server Configuration dialog box 4-11
- Server CPU Percent* 27-7
- Server-level objects
 - creating 3-1
- Servers
 - connecting to 4-2
 - displaying properties 4-9
- Server user ID 13-4
- Session rates 26-3
- Session start time 5-9
- Set Database Online command 9-22
- Setting
 - cache strategy 22-14
 - configuration parameters 4-11
 - default login mapping 6-12
 - execution attributes for process 4-14
 - explicit login mapping 6-13
 - filters 5-13
 - maximum rows per page 24-8
 - Performance Trends Monitor 5-26
 - role activation 13-9, 15-8
 - RPC security model 6-8
 - sample interval 5-11
 - security options 6-7
 - server class 6-6
 - statistic type 5-11
 - table permissions 22-16
- Setting a database online 9-22
- Setting up
 - users 14-1
- Shared memory 5-2
- shared memory starting address configuration
 - parameter A-9
- Sharing
 - logins 12-2, 13-9
 - permissions 12-2, 13-9, 14-12
- Shortcut menus 2-13
- Shortcuts xxxiv
 - drag-and-drop 2-8
 - keyboard 2-11

- updating permissions 22-17
- Showing
 - monitor status bar 5-9
- Shutting down
 - Adaptive Server 4-8
- Single-user mode 10-9
- Single-User Mode database option 9-14
- Site handler 6-2
- Size
 - estimating database 9-1
 - increasing database 9-7
 - reducing database 9-7
- Size/Prec 22-7
- size of auto identity column configuration
 - parameter A-9
- Small icon
 - toolbar button 2-15
- Small Icon display format 2-6
- Sorting data 24-4
- sp_estspace procedure 9-2
 - estimating space requirements 9-2
- sp_sysmon 5-29
- Space
 - allocating on segments (tables) 22-18
- Specific login mapping 6-10
- Specific Login Mapping tab 6-13
- Specifying
 - primary key 22-10
- sql.ini file 4-1
- sql.ini file 5-3, 5-5
- SQL Perfmon Integration configuration
 - parameter A-1
- sql server clock tick length configuration
 - parameter A-9
- sso_role 12-4
- stack guard size configuration
 - parameter A-9
- stack size configuration parameter A-9
- Standard toolbar 2-14
- Starting
 - Adaptive Server 4-7
 - Monitor Server 5-4
- Starting address 8-4
- start mail session configuration
 - parameter A-9
- startserver 4-7
- Startup
 - connecting at 4-5
- Statistics
 - updating 24-11
- Statistic type
 - benchmarks and 27-6
- Statistic type, on monitors 5-11
- Statistic type setting 5-9
- Status
 - active and pending 7-2
- Status bar
 - on monitors 5-9
 - showing and hiding 5-9
 - turning on or off 2-8
- Stopped process
 - researching causes 27-8
- Stopping
 - Adaptive Server 4-8
 - backups 10-7
 - restore 10-10
- Storage
 - displaying allocation 9-7
- Stored Procedure Activity
 - Monitor 26-64
- Stored procedures
 - See also* Extended stored procedures
 - See also* Procedures
 - deleting 17-2
 - monitor 26-64
 - tuning with monitors 27-9
- Storing
 - allocating space 9-2
 - considerations 9-2
 - costs 8-8
 - databases 8-2
 - default devices 8-2
 - increasing space 9-7
 - logs 8-2
 - segment benefits 11-1
 - strategies 8-12
- Strategies

- monitoring 27-6, 27-7
- tuning 27-10
- Strategy
 - cache 24-5
 - recovery 8-3
 - setting cache 22-14
 - storing 8-12
- suid. *See* Server user ID.
- Summary of constraints 22-3
- Summary of monitors 26-2
- suspend audit when device full configuration
 - parameter A-10
- syb_sendmsg port number configuration
 - parameter A-10
- Sybase Central 5-3
 - window 2-2
 - refreshing 2-16
 - windows and dialog boxes of 2-1
- sybssystemprocs database 8-2
- Syntax highlighting 2-18
- sysdevices table 8-3
- System Administration Guide* 10-1, A-1
- System Administrators
 - changing database defaults 9-12
 - configuring Adaptive Server 4-10
 - issuing checkpoints 9-17
 - managing resources 8-1
 - permissions 12-3
 - transferring database ownership 9-6
- System roles 12-3
- System Security Officer 13-1
 - configuring Adaptive Server 4-10
 - defining mutual exclusivity 15-9
 - permissions 12-4
- System Security Officers
 - remote login mapping 6-10
- system segment 11-1
- System tables 22-1
 - checking consistency 9-21
 - updates to A-2
- systemwide password expiration
 - configuration parameter A-10

T

- Tabbing through dialog boxes 2-11
- Table
 - check constraint 22-19
 - verifying proxy 23-6
- tablealloc 22-28
- Table editor 22-4, 22-6
 - opening
 - Opening
 - table editor 22-5
 - proxy tables 23-5
- table lock spinlock ratio configuration
 - parameter A-10
- Table Properties tab 3-3
- Tables
 - adding columns 22-13
 - allocating pages 22-28
 - allocations 22-28
 - binding defaults to columns 22-32
 - binding rules to 18-6
 - binding rules to columns 22-31
 - binding to a named cache 7-10
 - cache replacement strategy 22-14
 - check constraints 22-9
 - checking consistency 22-26
 - checking overall consistency 22-27
 - column definitions 22-7
 - column properties 22-30
 - creating 22-2, 22-6
 - deleting 22-11
 - deleting check constraints 22-32
 - dependencies 22-14
 - displaying dependencies 22-14
 - estimating space requirements 9-2
 - fixing text 22-29
 - maximum rows per page 24-8
 - on segments 11-6
 - permissions 22-16
 - placing on segments 22-18
 - property sheet 22-12
 - proxy 23-1
 - recompiling 24-11
 - referential constraints 22-21
 - reindexing 22-29

- shortcuts 22-17
 - statistics 24-11
 - sysdevices* 8-3
 - system 22-1
 - unbinding defaults form
 - columns 22-32
 - unbinding rules 18-9
 - unbinding rules form columns 22-31
 - unique constraints 22-24
 - user 22-1
 - Tables *See also* Proxy table
 - Tabs
 - access 12-13
 - buttons 3-3
 - granting permissions 12-13
 - property sheets 3-3
 - tape retention in days configuration
 - parameter A-10
 - Tasks, administration 1-2
 - tcp no delay configuration parameter A-10
 - Technical Support, contacting xxxvii
 - tempdb*
 - tuning with monitors* 27-10
 - tempdb* database 8-2
 - Threshold filter 5-15
 - Thresholds
 - creating 11-8
 - deleting 11-11
 - hysteresis value 11-7
 - last chance 11-7
 - modifying 11-10
 - purpose of 11-6
 - time slice configuration parameter A-10
 - Title bar, on monitors 5-8
 - Toolbar
 - code editor 2-18
 - creating objects 1-3
 - hiding 2-14
 - standard 2-14
 - turning on or off 2-8
 - Tools menu 2-13
 - ToolTip 2-16
 - Top n filter 5-14
 - total data cache size configuration
 - parameter A-10
 - total memory configuration
 - parameter A-10
 - Transaction Activity Monitor 26-67
 - daily monitoring with 27-7
 - Transaction log 8-2, 8-13
 - allocating space for 9-2
 - backing up 9-13
 - checkpoints 9-17
 - dumping 9-13
 - recovering 10-8
 - segments and thresholds 11-7
 - Transaction Rate* 27-7
 - Transact-SQL commands 3-5
 - Transferring database ownership 9-6
 - Tree
 - object 2-4
 - Triggers
 - code 19-3
 - commands 19-1
 - creating 19-2
 - deleting 19-2
 - displaying dependencies 19-5
 - enforcing referential integrity 19-1
 - nested A-1
 - object owners 19-1
 - properties 19-4
 - property sheet 19-4
 - recompiling 22-26, 24-11
 - Troubleshooting 3-5
 - monitor connections 5-27
 - Truncate Transaction Log on Checkpoint
 - database option 9-13
 - Trusted password 6-14
 - Tuning
 - application 25-1
 - Tuning strategies 27-10
 - Type 10-3
- ## U
- Unbinding
 - default from column 16-8

- default from columns 22-32
 - default from user-defined datatype 16-8
 - object from execution class 25-19
 - objects from a named cache 7-13
 - rule from column 18-9
 - rule from columns 22-31
 - rule from user-defined datatype 18-9
- unified login required configuration
 - parameter A-10
- Unique 24-2
 - constraints
 - properties 22-25
 - indexes 24-1, 24-2
 - values 24-2
- Unique constraints
 - creating 22-24
- Unlocking logins 13-4
- Unsuccessful connections 4-6
- Update permission 12-7
- update permission 14-7, 14-17, 15-15, 21-7, 22-17
- Updating
 - permissions
 - procedures 17-6
 - tables 22-16
 - views 21-6
 - statistics 24-11
 - Sybase Central windows 2-16
- upgrade version configuration
 - parameter A-10
- Usable by Database Owner Only
 - database option 9-14
- User
 - processes
 - displaying 4-12
 - property sheet 14-4
 - table 22-1
- User Datatype Properties tab 20-4
- User-defined datatypes
 - binding defaults 16-7
 - binding rules 18-7
 - creating 20-1
 - deleting 20-3
 - displaying dependencies 20-5
 - properties 20-4
 - unbinding default from 16-8
 - unbinding rule from 18-9
- User-defined roles 12-3
 - See also* Roles
 - creating 15-2
 - deleting 15-2
 - properties 15-3
- User environment configuration
 - parameters A-18
- user log cache size configuration
 - parameter A-10
- user log cache spinlock ratio configuration
 - parameter A-10
- Users
 - adding to group 14-14
 - aliasing to login 14-9
 - assigning aliases 14-4
 - command permissions 14-8
 - creating 14-2
 - deleting 14-3
 - displaying 14-3
 - displaying objects owned 14-5
 - granting object permissions 14-5, 14-16
 - guest 9-3, 12-1
 - objects owned 14-5
 - permissions
 - procedures 17-6
 - tables 22-16
 - properties 14-3
 - removing alias to login 14-11
 - removing from groups 14-15
 - setting up 14-1
- Uses
 - permissions
 - views 21-6
- use security services configuration
 - parameter A-10
- Using online help xxxvi
- Utilities 2-3
 - bcp 9-13
 - dbcc 9-18, 22-26

writetext 9-13

X

xp_cmdshell context configuration
parameter A-11

V

Values

- default 16-1
- index keys 24-2
- null 16-5, 20-3
- rules defining 18-1
- unique 24-2

Verifying

- proxy table 23-6

Verify Table command 23-7

View menu 2-12

Views

- adding columns 21-1
- Create dialog box 21-1
- creating 21-1
- displaying dependencies 21-5
- permissions 12-10, 21-6
- properties 21-3
- property sheet 21-3

W

Wash area 7-9

Wildcard, in filters 5-18

Window

- changing size of pane 2-3
- container icons 2-3
- customizing the display of 2-6
- menus and toolbars 2-12
- moving through 2-4
- object icons 2-3
- panes 2-3

With cascade 12-6, 13-2, 21-6

With grant 12-6, 21-6

With override 9-3

Workload 27-4

writetext utility 9-13

